UNIVERSITY OF COPENHAGEN

Ph.D. thesis

Anders Simonsen

Sensitive electro-optical transduction

Through resonant electro- and optomechanics

Advisor: Eugene S. Polzik

Handed in: September 16, 2019

Abstract

This thesis presents a scheme that transduce an electrical signal into and optical one, while simultaneously amplifying the signal. The transduction uses a coupling between an electrical and mechanical resonator, and experiments show that the coupling make the electrical signal dominate over thermal noise in the mechanical resonator, making the transduction scheme work as a low-noise preamplifier. When operated at a Direct Current (DC) bias, the transduction reached an optical noise-temperature of $4\,\mathrm{K}$ in a bandwidth of $3\,\mathrm{kHz}$, and the noise-temperature projects to reach $90\,\mathrm{mK}$ with better optical detection, a sensitivity that is competitive with conventional electronics.

The initial transduction suffered from technical limitations, mostly from its assembly. These were addressed through a micro-fabrication process, a process that produced a new generation of transducers integrating both an electromechanical system and optical cavity on a single chip. The cavity interfaced directly with a fiber-network and greatly simplified the setup that is to operate the transducer. This technical innovation enabled an experiment where the transduction replaced a conventional electronic amplifier in a detection circuit for Magnetic Resonance Imaging (MRI), thus bringing the advantages of optical signal processing to the receiver chain of MRI scanners.

Transduction with the MRI circuit used an Alternating Current (AC) biasing and had a noise-temperature of $210\,\mathrm{K}$ when detection at $32\,\mathrm{MHz}$—the Larmor frequency of $^{13}\mathrm{C}$ at $3\,\mathrm{T}$. The noise corresponds to $99\,\mathrm{pV}/\sqrt{\mathrm{Hz}}$, $113\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$, or $8\,\mathrm{fT}/\sqrt{\mathrm{Hz}}$ for the particular circuit. The signal-to-noise bandwidth was $12\,\mathrm{kHz}$. Operating the transducer in a commercial medical scanner for MRI required additional modifications that sacrificed the performance, but this degradation may be solved with straightforward improvements. The transduction nevertheless obtained an MRI image, a success that proves the transduction can become a valuable alternative to conventional electronics in MRI.

Sammenfatning

Denne afhandling præsenterer et system der kan transducere et elektrisk signal til et optisk og samtidig forstærke signalet. Systemet gør brug af en kobling mellem en elektrisk og mekanisk resonans, og eksperimenter viste at koblingen undertrykker den termiske støj fra den mekaniske resonator i forhold til selve signalet. Det gør transductionen ækvivalent til en følsom forstærker. Når systemet opererede med et DC bias, så opnåede det en en optimal støj-temperature på 4 K i en 3 kHz båndbredde, og støj-temperature kan nå helt ned til 90 mK med et forbedret optisk system. Denne følsomhed kan konkurrere med konventional elektroniske fortærkere.

De første eksperimenter med transductionen led af tekniske begrænsninger, mest relateret til fabrikationen af selve transduceren. Disse blev løst via ved udviklingen af en mikro-fabrikations proces, en proces der integrerede både det elektromekaniske system og et optisk kavitet på en chip. Kaviteten blev koblet direkte til en optisk fiber-netværk hvilket simplificerede det nødvendige setup markant. Denne tekniske innovation tillod at bruge transduction til at danne et MRI billede, i stedet for at bruge en standard elektrisk forstærker, et resultat der bringer fordelene ved optisk processering af signaler til MRI skannernes modtager-kanaler.

Transduction med MRI opstillingen brugte et AC bias med en demonstreret støj-temperatur på 210 K for detektion ved 32 MHz (Larmor-frekvensen for ^{13}C ved 3 T). Denne støj svarer til en spændings-følsomhed på 99 pV/$\sqrt{\text{Hz}}$, strøm-følsomhed på 113 pA/$\sqrt{\text{Hz}}$ og en følsomhed for magnetfelter på 8 fT/$\sqrt{\text{Hz}}$. Signal-til-støj båndbredden for denne følsomhed var 12 kHz. At anvende transduceren i en MRI scanner krævede yderligere modifikationer der kostede på teknikkens præstation, men denne forværring kan løses med enkle forbedringer. På trods af problemerne gav denne nye teknik et MRI billede - en succes der viser at transduction kan blive et værdifuldt alternativ til konventionel elektronik i MRI.

Acknowledgments

I owe a great debt of gratitude to a long list of people without whom this project would never have succeeded. First, I would like to thank my official supervisor, Eugene Polzik, for having an inspiring vision with the research and giving me both opportunity and freedom with the project, even when progress was slow and results far apart. I would also like to thank the people that have been my supervisors in all but an official capacity: Albert Schließer, whose guidance in all things optomechanics have been invaluable; and Silvan Schmid, whose cleanroom expertise guided me towards the final product.

I also must thank all of my direct collaborators, from whom I have drawn much inspiration and motivation. First, we would never have succeeded without the help of Juan Diego Sanchez and his supervisor Jan Henrik Ardenkjær-Larsen from the Hypermag group, and their hands-on knowledge of MRI. Juan has been a very involved collaborator, even when we faced some frustrating practical issues. Then there are all the helpful people I have encountered through Qubiz. They have opened my eyes to the applied and commercial side of research. A special thanks to Radu Malureanu for offering a much-needed fresh perspective when I had starred myself blind on the fabrication. I am also grateful for the theoretical support from our resident theorist Emil Zeuthen and his supervisor, Anders Sørensen, and the original effort from Jake Taylor.

Then there are all the good people from Quantop and affiliates. During my several years as a Ph.D. student, I have had the pleasure to work with many people, and I thank you all for fruitful discussions and a friendly atmosphere that permeates all the Quantop labs. I especially thank Tolga Bagci for jump-starting me into the project, and Sampo Saarinen for stepping up and helping out with the experimental efforts. Of course, I must also thank the cleanroom gang: Yeghishe Tsaturyan and Yannick Seis. Your input has been a tremendous source of inspiration, and sharing experiences with you have made my frustrations in fabrication much more bearable. I am also grateful to Jürge Appel, Jörg Helge Müller, and

iv

List of publications

- T Bagci, A Simonsen, S Schmid, L. G. Villanueva, E Zeuthen, J Appel, J. M. Taylor, A Sørensen, K Usami, A Schliesser, and E. S. Polzik (2014). "Optical detection of radio waves through a nanomechanical transducer." In: *Nature* 507.7490, pp. 81–85

- Y. Tsaturyan, A. Barg, A. Simonsen, L. G. Villanueva, S. Schmid, A. Schliesser, and E. S. Polzik (2014). "Demonstration of suppressed phonon tunneling losses in phononic bandgap shielded membrane resonators for high-Q optomechanics." In: *Optics express* 22.6, pp. 6810–21

- A. Simonsen, S. A. Saarinen, J. D. Sanchez, J. H. Ardenkjær-Larsen, A. Schliesser, and E. S. Polzik (2019b). "Sensitive optomechanical transduction of electric and magnetic signals to the optical domain". In: *Optics Express* 27.13, pp. 18561–18578

- A. Simonsen, J. D. Sanchez, S. A. Saarinen, J. H. Ardenkjær-Larsen, A. Schliesser, and E. S. Polzik (2019a). "Magnetic resonance imaging with optical preamplification and detection". In: *arxiv*, pp. 1–6

Patent

- E. S. Polzik, A. Schliesser, S. Schmid, A. S. Sørensen, J. M. Taylor, K. Usami, T. Bagci, A. Simonsen, L. G. Villanueva, E. Zeuthen, J Appel, E. S. Polzik, A. Schliesser, S. Schmid, A. S. Sørensen, J. M. Taylor, K. Usami, T. Bagci, A. Simonsen, L. G. Villanueva, E. Zeuthen, and J Appel (2014). *Optical detector and amplifier for RF-detection having a position dependent capacitor with a displaceable membrane*

Contents

List of Figures

Acronyms

AC Alternating Current.
AFM Atomic-Force Microscope.
ALD Atomic-Layer Deposition.

DC Direct Current.
DTU Technical University of Denmark.

EBPVD Electron-Beam Physical Vapor Deposition.

GRIN GRadient INdex of refraction.

ICP Inductively-Coupled Plasma.

LPCVD Low-Pressure Chemical-Vapor Deposition.

MEMS Micro-Electro-Mechanical System.
MRI Magnetic Resonance Imaging.

NBI Niels Bohr Institute.
NMR Nuclear Magnetic Resonance.

PCB Printed Circuit Board.
PECVD Plasma-Enhanced Chemical Vapor Deposition.
PSD Power Spectral Density.

RF Radio Frequency.
RMS Root-Mean-Square.

SEM Scanning Electron Microscope.
SNR Signal-to-Noise Ratio.

UV Ultra-Violet.

Chapter 1

Introduction

Electronics and optics are mature technologies that are vital to modern science and engineering, each presenting distinct and complementary advantages. Optical systems have found purchase in impactful applications such as telecommunication, sensing, and metrology due to low-loss signal transmission and performance at the quantum limit. Electrical systems pervade all aspects of the modern world due to immense flexibility and ease-of-use attained after incredible engineering effort. Unsurprisingly, merging both technologies is an active field of research, a field that promises to bring the best of both worlds together, deliver unprecedented performance in existing applications, and pave the way for brand new scientific tools and engineering solutions. Such consolidation is offered by cavity electro- and optomechanics, a technique that has garnered significant research interest in the past ten years for its ability to transduce faint electrical signals to an optical carrier—even extending to quantum-states transfer—with most research efforts concentrating on proof-of-principle experiments in laboratory environments and the elusive quantum performance. However, pertinent practical applications demand more than just a working laboratory setup; for example, compactness, ease-of-use, scalability, and transportability are essential features for viable products. So to reach the full potential of cavity electro- and optomechanics, we must investigate relevant applications and explicitly test the scheme in the environments and conditions where they might be employed.

Transducing weak electrical signals to the optical domain, where they may be transmitted and detected at the quantum-noise level of light, could benefit applications that rely on electronic sensors. Sensing weak electrical signals have many usages owing to the omnipresence of electrical machinery. Examples in-

1

clude radio astronomy, navigation, and classical and quantum communication. This thesis considers one application in particular: Magnetic Resonance Imaging (MRI), a vital and non-invasive tool used in clinical diagnostics, and basic research in medicine. MRI stands to benefit from optical detection and low-loss transmission of signals because the light is immune to electromagnetic interference from the environment and unaffected by strong magnetic fields. These features solve some key technical challenges that MRI faces; specifically, preamplifier materials, cable loss, and Radio Frequency (RF) environment noise (Kathiravan and Kanakaraj 2013). Those problems look to get worse as MRI instruments move towards stronger magnetic fields and arrays with more detection coils. The high field challenges conventional electronics(Sobol 2012), while more coils mean more preamplifiers and cables that all create interference noise (Wiggins et al. 2009).

The next chapters describe sensitive electronics-to-optics transduction through cavity electromechanics with optical readout, without direct coupling between the circuit and light—instead, the electrical and optical system both couple to a mechanical intermediary. First, it is shown that cavity electromechanics can transduce faint electrical signals onto an optical carrier, amplifying them in the process, while adding only very little noise (chapter 2). Transduction thus acts as the first low-noise amplifier in a detection chain, which means the scheme can operate as a sensor for electrical signals (chapter 3). Moving to an MRI scanner required a new fabrication process to produce improved electromechanical devices (chapter 4), devices that added an optical cavity with direct fiber-coupling for easy optical readout (chapter 5). With these improvements, the transducer performed comparably to state-of-the-art electronic amplifiers in an MRI-ready detection circuit (chapter 6), and experiments eventually demonstrated optical detection of an MRI image collected with a standard coil (chapter 7).

The text has the following structure: the remaining introduction gives a brief overview of previous literature and current state-of-the-art, followed by an introduction of the full transduction scheme. After that comes three parts dealing, respectively, with the transduction principle, fabrication, and MRI testing. Each part has two chapters. Those chapters start with a brief history of the project to illuminate the various people involved, with the remaining chapter mostly grouped in experiments and subdivided into theory (when relevant), method, result, and discussion sections.

1.1 The background

The study of cavity electro- and optomechanics started when Braginskii and Manukin (1967) suggested and later demonstrated (Braginskii et al. 1970) that radiation pressure of photons can change the dynamics of a mechanical oscillator, provided the photons circulate in a resonator which retards the response of the force. Early demonstrations used rather large mechanical elements, but the Micro-Electro-Mechanical System (MEMS) revolution in the 1980s introduced new, micro-machined mechanical resonators with a tiny mass that were much more susceptible to the radiation-pressure. These new systems initially struggled with optical performance (Aspelmeyer et al. 2014), but some implementations exhibited optomechanical effects related to photothermal forces, e.g. self-induced oscillations (Zalalutdinov et al. 2001) and cavity cooling (Metzger and Karrai 2004).

The massive impact of MEMS and micro-fabrication cannot be understated; those techniques laid the foundation for countless products and technologies, including the transducer presented in this thesis. In the 1990s, optical MEMS started to emerge as miniaturization and micro-fabrication of optical components, producing components with widespread applicability in, for example, telecommunication (optical switching and modulation through electrostatic actuation), displays, and all kinds of sensors (Solgaard et al. 2014). All these examples demonstrate both the potential benefit of integrated electronics and optics, and suggest areas that stand to benefit from improved integration.

1.1.1 Our project

Our group, Quantop at the Niels Bohr Institute (NBI), started researching cavity opto- and electromechanics following a proposal (Taylor et al. 2011) on how to extend laser cooling to electrical resonators. Prof. Eugene Polzik supervised the effort, and we started a close collaboration with Nanotech at Technical University of Denmark (DTU) (Silvan Schmid and Anja Boisen). The proposal considered a membrane-in-the-middle setup (Thompson et al. 2008; Wilson et al. 2009), so that became our system of choice. In parallel, our group implemented the optomechanical system sans electronics (Nielsen et al. 2017) and coupled the membrane to an atomic ensemble (Møller et al. 2017), ensembles that were the hitherto expertise at Quantop. The lofty ideal was to one day couple the atoms to the electrical resonator via the membrane.

When our research began, the state-of-the-art was cavity optomechanical systems with electrical actuation based on whispering-gallery optical resonators (McRae et al. 2010; Sridaran and Bhave 2011) and photonic waveguides (Winger et al. 2011). I joined the group at the end of 2012 in a combined Master's+Ph.D. program, and began to work on the experimental effort together with Tolga Bagci. Albert Schliesser joined us shortly after that, and Sampo Saarinen joined the project in 2017. Since we started, a plethora of opto-electro-mechanical systems have emerged—see Safavi-Naeini et al. (2019) for a thorough review. Of note, Higginbotham et al. (2018) have recently gotten near 50 % conversion efficiency between microwave- and optical photons, using a membrane-in-the-middle setup similar to what we originally envisioned.

Our research took a hard turn towards classical sensing applications after our first publication, Bagci et al. (2014), demonstrated that the cavity electromechanical interaction could reduce the intrinsic mechanical noise, thus enabling low-noise detection, amplification, and transduction. In 2016, we started collaborating with DTU Hypermag (Juan Diego Sanchez and Jan Henrik Ardenkjær-Larsen) to use the electro-mechano-optical transduction in an MRI scanner, a shifted focus that required a dramatic change in how we fabricated the electromechanical devices. That labor culminated with two manuscripts, Simonsen et al. (2019a,b), where the latter present our first MRI image.

1.1.2 State-of-the-art for MRI

In MRI, the typical setup uses a coil to detect the RF Nuclear Magnetic Resonance (NMR) signal generated by a sample. That coil then feeds the signal to a low-noise preamplifier whose output goes via coaxial cable to the scanner's receiver channel (Edelstein et al. 1986). These preamplifiers are the standard to beat. Since the coil already is a part of the detection circuit, cavity electromechanics follow by adding the transducer and a tuning capacitor in parallel to the coil, thus creating an electrical resonance. However, a practical circuit may not be that simple because the preamplifiers also serve other functions, for instance, decoupling nearby coils by suppressing the current on resonance.

Other people have recognized the benefits of using optical technology in MRI and made different types of optical sensors to complement the standard scanners (Taffoni et al. 2013). For example: Filograno et al. (2016) used a fiber grating to monitor the large magnetic field. Memis et al. (2008) used conventional electronics to amplify the signal at the coil but then modulated a laser-diode to get

the signal out via fiber instead of cable. And Rugar et al. (1992) used a MEMS device to detect an MRI signal and measured the motion optically, a technique that Sidles and Rugar (1993) argued can become a viable alternative to the standard coil technology. Some commercial systems even digitize the NMR signal already at the coil and send it to the receiver over optical fiber (Phillips dStream).

1.1.3   The competition

In parallel to this thesis, other groups have implemented the same membrane-based cavity-electromechanical setup. Haghighi et al. (2018) extended the principle to encompass multiple mechanical modes which increased the transduction bandwidth, while Takeda et al. (2017) and Tominaga et al. (2018) used the transduction to detect an NMR signal.

1.2   The electro-mechano-optical platform

Figure 1.1a depicts the cavity electromechanical setup, with optical readout, that we used for transduction in an MRI scanner (Simonsen et al. 2019a). It has three main parts elaborated separately below: a mechanical, electrical, and optical resonator, with resonance frequency Ω_m, Ω_{LC}, and Ω_o, respectively—as illustrated in fig. 1.1b. All resonators were at room temperature, $T_0 = 300\,\mathrm{K}$, and the large thermal noise makes a classical description of transduction sufficient. The transducer comprises the mechanical and optical resonator, and it connects in parallel to the electrical resonator.

A voltage bias induces a linear coupling between the electrical and mechanical system, a coupling that is proportional to the bias amplitude. Furthermore, the Alternating Current (AC) bias frequency must match the difference between the electrical and mechanical resonance,

$$\Omega_{ac} = \Omega_{LC} - \Omega_m,$$

to parametrically couple the systems similarly to Dougherty et al. (1996). While the figure displays the operation with an AC bias, the scheme also functions with a DC bias ($\Omega_{ac} = 0$), but only if the membrane and circuit have degenerate resonance frequencies (Bagci et al. 2014). In that case, fig. 1.1 looks identical except that the top spectrum in fig. 1.1b gets down-converted to coincide with the mechanical peak.
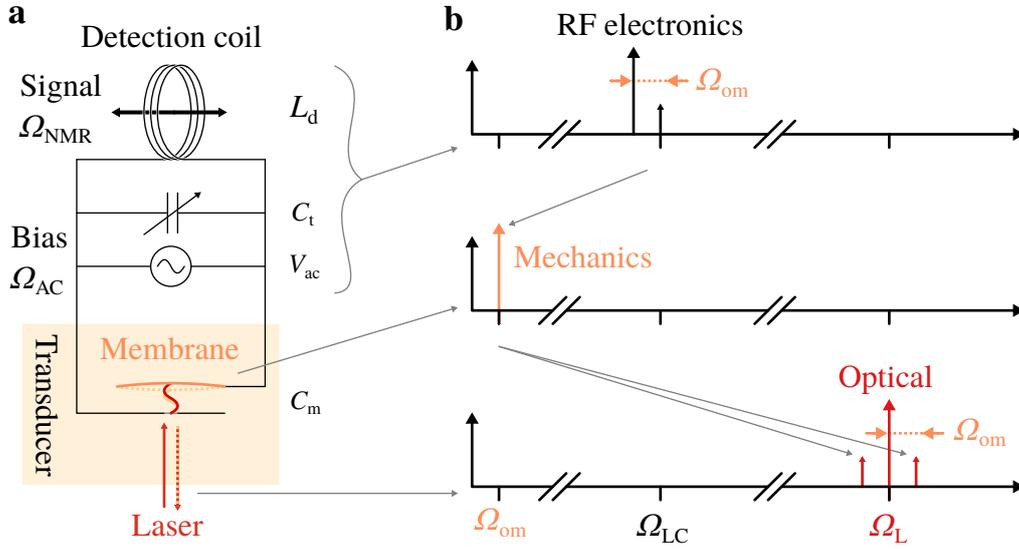
Figure 1.1: **a)** Setup schematic. A signal enters the detection coil and leaves as optical modulation. **b)** Spectral diagram showing how the signal moves in frequency through transduction. (*Adapted from Simonsen et al. (2019b).*)

1.2.1 Mechanical interface

The transduction uses a membrane under tensile stress as one electrode in a parallel-plate capacitor. In the end, our device features a small distance, around 600 nm, between the membrane and second electrode, resulting in a large electro-mechanical coupling with only a few volts of driving bias. This gap is central to transduction performance (Bagci et al. 2014; Takeda et al. 2017), and our design yield better performance over other works on room-temperature transduction with an AC bias (Simonsen et al. 2019b).

1.2.2 Circuit

The circuit consists of an electrical resonance formed between the detection coil—an inductor, L_d—and a parallel capacitor, C_T, comprised of a tuning capacitor, C_t, plus the membrane-capacitor. That resonator is a so-called *LC* tank. The circuit connects to the electromechanical capacitor through a standard integrated-circuit socket containing the device, and the resonance targets the NMR signal from ^{13}C in a 3 T magnetic field, a signal that has the Larmor frequency $\Omega_{NMR} = 32.19$ MHz. With a mechanical frequency of roughly 1 MHz, the AC bias frequency must be around 31 MHz. In DC transduction measurement, the

circuit and mechanics had degenerate resonance, $\Omega_{LC} = \Omega_{\mathrm{m}}$.

1.2.3 Optical readout

Initially, we did the transduction without a cavity and instead measured the mechanical displacement interferometrically. However, the final transducer design had the membrane act as one mirror in a cavity—the prototypical optomechanical setup (Aspelmeyer et al. 2014)—meaning a mechanical displacement modulates the reflected light. By design, that cavity yielded negligible backaction from radiation pressure, although photothermal backaction appeared instead, and we used the cavity solely to readout the mechanical displacement, in contrast to the proposal by Taylor et al. (2011) that prompted our endeavor. We deviated from the proposal in favor of an optical assembly that would be easy to move to the MRI scanner, i.e., a setup that was alignment-free, transportable, fiber-coupled, and fixed in length thus avoiding active stabilization of the cavity. All potential applications would appreciate such a compact, robust, and simple implementation.

1.2.4 Transduction process

Here is a conceptual explanation of the transduction: First, the detection coil collects a signal and convert it into a voltage in series with the coil. In the following, "signal" refers to all voltages in series with the coil, including thermal Johnson noise. Second, both signal and bias induce charges on the membrane-capacitor, thus attracting the membrane with a force proportional to the charges squared. Squaring the charges create a beat-note between signal and bias that effectively down-converts the signal as depicted in fig. 1.1b. Provided the right bias frequency, that beat-note drive the membrane near resonance where it responds strongly. The larger the bias, the more dominant the beat-note compared to mechanical thermal noise and optical readout. As the membrane moves, it changes the circuit resonance and, in turn, the charges that drive the motion. So the mechanical motion acts back onto itself—a so-called electromechanical backaction. Third, the mechanical motion changes the length of the optical cavity and its reflection of an input laser beam. That way the motion turns into an amplitude modulation of the light, effectively up-converting the signal to the optical domain and completing the transduction. By measuring the amplitude of reflection, the signal gets down-converted back to the mechanical frequency.

Part I

Electro-mechano-optical
transduction

Chapter 2

Theory

Here is the basic theory behind cavity electromechanical transduction, a classical description that assumes thermal-noise to be dominant. Transduction theory have not been the main aim of this thesis, and much of the mathematics has been covered in detail elsewhere, so the following treatise aims only to present the fundamental equations necessary to understand the experimental results and evaluate the performance of transduction. Derivations use a minimum of intermediate mathematical steps, although outlining some of the vital manipulations, and some results come directly from references when convenient. Good sources include Bagci et al. (2014), which cover DC transduction theory; Zeuthen (2015), which carefully analyses the membrane-capacitor model presented below; Zeuthen et al. (2017), which express the electromechanical interaction as an effective circuit, considering both AC and DC bias; Takeda et al. (2017), which cover AC transduction theory including an optical cavity; and Aspelmeyer et al. (2014), which gives a general introduction to cavity electro- and optomechanics.

The first section below describes the electromechanical setup, starting from the bare mechanics and then adding the capacitance and resonance circuit. The main results are the mechanical and electrical Hamiltonians used later to get the equations-of-motion, equations that further involve the effective mass, static deflection, effective membrane-capacitor distance, and the derivative of the membrane-capacitance. The second section considers the optical readout, explicitly expressing the optical output from transduction. All our experiments eventually measure this. Again, the main result is a Hamiltonian, this time for the optical interaction, but the linearization and photothermal backaction is also important. The third section presents the equations-of-motion and how to simplify them. That

simplification lumps the optical effects into an effective mechanical response and linearizes the equations around the AC or DC bias. Two additional noteworthy results are the Power Spectral Density (PSD) obtained with a thermal-noise drive—these constitute the output of practical measurements—and the relation between static displacement and mechanical frequency-shift. Finally, the fourth section expresses the intrinsic noise-temperature and bandwidth for transduction, parameters that derive from the electromechanical interaction strength characterized by the electromechanical cooperativity C.

2.1 Electromechanical platform

2.1.1 Mechanical system

The electromechanical device uses membranes for transduction. The dynamics of membranes is a standard problem in textbooks on continuum mechanics and elasticity, usually under names such as vibrating plates, disks, and drum-skins. The theory below only state the membrane motion, as derived elsewhere, and summarizes the fundamental assumptions behind the expressions.

Landau and Lifshitz (1986, ch. 11-14) calculated the dynamics of a vibrating plate by minimizing its free and potential energy, assuming the plate:

- Is a thin plate with thickness h is much smaller than its width.

- Can bend more than its thickness, but much less than its width, which means bending-induced volume changes are negligible.

- Has a dominating, in-plane tensile force $\mathcal{T} \gg 0$ that stretches the plate out isotropically along its circumference, which means bending-induced stresses are negligible.

- Has negligible in-plane motion.

A plate that meets these criteria constitutes a "membrane" in the following. With this, Landau and Lifshitz (1986) find the membrane's wave-equation:

$$\mathcal{T} \nabla^2 u_z(x, y, t) = \rho \frac{\partial^2 u_z(x, y, t)}{\partial t^2},$$ (2.1)

where ∇^2 is the vector Laplacian, ρ is the membrane-density, and $u_z(x, y, t)$ is the Lagrange strain tensor for a membrane in the x, y-plane, vibrating in the z-

direction as a function of time, t. This result is a standard wave-equation where $\mathcal{T}/\rho$ is the speed-of-sound in the membrane.

Mode expansion

Using standard separation of variables, eq. (2.1) splits into the Helmholtz equations (Riley et al. 2006, ch. 21) that describe the temporal- and spatial behavior of the strain tensor separately. Solving these equations requires knowledge of the membrane's specific boundary conditions. In this case, there are two relevant boundary-conditions: a square with side-length, L_s, and a circle with radius, L_r, that both anchors the membrane. These conditions both produce a discrete set of complete, orthogonal normal-modes of motions—i.e., standing waves with a node on the boundary. The exact expression is (Riley et al. 2006, eq. (21.55)):

$$u_z(x, y, t) = \sum_{a,b} Z_{a,b}(x, y)\, \beta_{a,b}(t)$$

$$= \sum_{a\geq 1,\, b\geq 1} \cos\left(\frac{a\,\pi\,x}{L_s}\right) \cos\left(\frac{b\,\pi\,y}{L_s}\right) \cos(i\,\Omega_{a,b}\,t), \qquad \text{(2.2 squ)}$$

$$= \sum_{a\geq 0,\, b\geq 1} J_a\left(\frac{j_{a,b}\,r}{L_r}\right) \cos(a\,\theta)\, \cos(i\,\Omega_{a,b}\,t), \qquad \text{(2.2 circ)}$$

where a and b are indices denoting a particular mode in the full set, a mode that has the spatial part $Z_{a,b}(x, y)$ and temporal part $\beta_{a,b}(t)$ oscillating at the eigenfrequencies

$$\Omega_{a,b} = \frac{\pi}{L_s}\sqrt{\frac{\mathcal{T}}{\rho}\left(a^2 + b^2\right)},, \qquad \text{(2.3 squ)}$$

$$\Omega_{a,b} = \frac{j_{a,b}}{L_r}\sqrt{\frac{\mathcal{T}}{\rho}}.. \qquad \text{(2.3 circ)}$$

The circular modes are Bessel-functions of the first kind, J_a, with $j_{a,b}$ being the b^{th} zero of J_a.

Mechanical Hamiltonian

The membrane-modes in eq. (2.2) look similar to a one-dimensional harmonic oscillator, although the motion is over a plane. That similarity develops further

by rewriting the Hamiltonian for the membrane, $\mathcal{H}_m$, like so:

$$\mathcal{H}_m = \sum_{a,b} \int_{\vec{S}} d^2\vec{r} \frac{Z_{a,b}^2}{A_{mem}} \left[\frac{m^*}{2} \left(\frac{\partial \beta_{a,b}}{\partial t} \right)^2 + \frac{m^* \Omega_{a,b}^2}{2} \beta_{a,b}^2 \right]$$

$$= \sum_{a,b} \left[\frac{p_{a,b}}{2\,m_{a,b}} + \frac{m_{a,b}\,\Omega_{a,b}^2}{2} \beta_{a,b}^2 \right], \tag{2.4}$$

where the integral is over the surface $\vec{S}$ of a membrane with area A_{mem} and physical mass

$$m^* = h\,A_{mem}\,\rho. \tag{2.5}$$

The second expression in eq. (2.4) has enforced orthogonality,

$$\int d^2\vec{r}\, Z_{a,b}\, Z_{a',b'} \propto \delta_{a,a'}\,\delta_{b,b'},$$

and introduced an effective mass, $\tilde{m}$, for each mode,

$$\tilde{m}_{a,b} = m^* \int_{\vec{S}} d^2\vec{r}\, \frac{Z_{a,b}^2}{A_{mem}}, \tag{2.6}$$

together with the conjugate momentum of $\beta_{a,b}$,

$$p_{a,b} = m_{a,b} \frac{\partial q_{a,b}}{\partial t}.$$

Now the result is a genuine Hamiltonian expressed in terms of canonically conjugate variables. It is specifically a sum of independent harmonic oscillators with each mode of motion being a unique oscillator. Out of all these, the fundamental mode[1] is the most interesting for transduction (argued in the next section).

Equation (2.6) introduces the important concept of effective mass, but the definition contains an ambiguity in how the mode-shape $Z_{a,b}$ is normalized. For example, the normalization in eq. (2.4) means that β equals the real displacement where the mode-shape $Z_{a,b}$ equals one. Incidentally, that is the point of highest deflection for both boundary-conditions, a point that appears at the center of the membrane for the fundamental modes. Alternatively, the normalization could instead set the integral in equation eq. (2.6) equal to one, making that the effective- and physical mass identical, but that would require rescaling β in subsequent analysis. See Pinard et al. (1999) and Zeuthen (2015, ch. 2 and app. A) for more rigorous discussion.

[1]Mode index (1,1) for the square boundary and (0,1) for the circular.

2.1.2 Electromechanical system

The electromechanical device forms a parallel-plate capacitor between the conductive membrane and a second electrode. Membrane motion changes the distance between the capacitor plates, thus changing the capacitance, and the effect depends on the mechanical mode-shape found in eq. (2.2). To include the mode-shape, Zeuthen (2015) developed this simple model: use the standard formula for parallel-plate capacitance in infinitesimal areas $d^2\vec{r}$ and then integrate over the membrane surface $\vec{S}$. The parallel-plate equation works well when the distance between capacitor plates is smaller than the plate-width (Carlson and Illman 1994)—as is the case here—which suggest that the model works for deflection much smaller than the plate-distance. Mathematically, the model looks like this:

$$C_{\mathrm{T}}\left(u_z(\vec{r}, t)\right) = C_{\mathrm{t}} + C_{\mathrm{m}}\left(u_z(\vec{r}, t)\right) \quad \text{with} \tag{2.7a}$$

$$C_{\mathrm{m}}\left(u_z(\vec{r}, t)\right) = \int_{\vec{S}} d^2\vec{r}\, \frac{\epsilon}{d_{\mathrm{eff}} - u_z(\vec{r}, t)} \zeta(\vec{r}), \tag{2.7b}$$

where $C_{\mathrm{T}}\left(u_z(\vec{r}, t)\right)$ denotes the total electromechanical capacitance, C_{t} denotes a constant capacitance in parallel to the membrane-capacitor,[2] and $C_{\mathrm{m}}\left(u_z(\vec{r}, t)\right)$ denotes the mechanically compliant capacitor. In the integral, $\zeta(\vec{r})$ is a function that accounts for the overlap between the top and bottom capacitor plates,[3] ϵ_0 is the permittivity of vacuum, A_{mem} is the surface area of the membrane, and d_{eff} is the distance between the capacitor plates. This distance is not the physical gap between the plates; instead, it accounts for dielectric layers between the membrane and electrode like so

$$\frac{\epsilon_0}{d_{\mathrm{eff}}} = \frac{\epsilon_1}{d_1} + \frac{\epsilon_2}{d_2} + \dots,$$

where ϵ_ν is the permittivity of layer ν.

---

[2]Either a parasitic capacitance or a tuning capacitor needed to make the *LC* circuit resonate at the correct frequency.

[3]For the actual electrode designs latter on, this overlap is a constant: either 1 or 1/4, as explaing in section section 3.1.1 on page 38

Anticipating analysis to come

For future reference, here is the k^{th} derivative of the membrane-capacitance with respect to the mode of motion:

$$\frac{\partial^k}{\partial \beta_{a,b}^k} C_{\text{T}}\left(u_z(\vec{r})\right) = \frac{\partial^k}{\partial \beta_{a,b}^k} C_{\text{m}}\left(u_z(\vec{r})\right)$$

$$\frac{\partial^k}{\partial \beta_{a,b}^k} C_{\text{T}}\left(u_z(\vec{r})\right) = \int_{\vec{S}} d^2\vec{r}\,\frac{\epsilon}{\left(d_{\text{eff}} - u_z(\vec{r})\right)^{k+1}}\, k!\, Z_{a,b}(\vec{r})^k\, \zeta(\vec{r})$$

$$\approx \frac{\epsilon\, A_{\text{mem}}}{d_{\text{eff}}^{k+1}} \int_{\vec{S}} d^2\vec{r}\, k!\, Z_{a,b}(\vec{r})^k\, \zeta(\vec{r})/A_{\text{mem}} \qquad (2.8)$$

assuming $u_z(\vec{r}) \ll d_{\text{eff}}$. When evaluated for a circular membrane, eq. (2.8) gives the first-order derivative

$$\frac{\partial}{\partial \beta_{0,1}} C_{\text{m}}\left(u_z(\vec{r})\right) = -\frac{\epsilon\, A}{d_{\text{eff}}^2}\frac{2\, J_1(j_{0,1})}{j_{0,1}} \qquad (2.9)$$

for the fundamental mode, assuming the mode overlaps completely with the capacitor electrodes ($\zeta(\vec{r}) = 1$). The corresponding second-order derivative is:

$$\frac{\partial}{\partial \beta_{0,1}}\left(u_z(\vec{r})\right) C_{\text{m}} = \frac{-d_{\text{eff}}}{J_1(j_{0,1})\, j_{0,1}}\frac{\partial^2}{\partial \beta_{0,1}^2} C_m\left(u_z(\vec{r})\right). \qquad (2.10)$$

When the equations-of-motions are derived later (section 2.3 on page 21), the first-order derivate above turns out to determine the electromechanical interaction. Evaluating the derivative (eq. (2.8)) with the membrane modes (eq. (2.2)), it is evident that the fundamental mode has the highest derivative. Therefore, the following analysis will exclusively refer to the fundamental mode. The higher-order modes have smaller derivatives because they include standing-wave nodes in the integral, nodes that entail both positive and negative displacement. The overlap function can exclude the negative contribution to the integral, but that limitation also reduces the capacitance.

Static deflection

As shown in section 2.3.3, the transduction creates a static, attractive force between membrane and electrode, causing the membrane to deflect. The deflection may be approximated by the fundamental mode. Here is why: the transduction

bias induces charges on the capacitor that attract the membrane both statically[4] and uniformly[5]. As an example, the deformation of the membrane looks like

$$Z_{\text{circ}}(\vec{r}) \propto 1 - \frac{r^2}{R^2}$$

(Voorthuyzen and Bergveld 1984) for a circular membrane, a result that can be expanded into the mode-functions $Z_{a,b}(x, y)$ through generalized Fourier synthesis (Wolf 1979, Ch. 6.2). Such a synthesis returns a Fourier coefficient for the fundamental mode close to one ($\sim$1.1), confirming that the fundamental mode approximates the deflection reasonably well. That conclusion also applies to a square membrane for a similarly uniform static force, although with a slightly different Fourier coefficient. That means subsequent analysis can encompass both the static deflection and dynamical vibration by considering only the fundamental mode.

Pull-in instability

The static deflection is proportional to the bias amplitude and eventually brings the membrane in contact with the electrode—which permanently destroys the device. A complete analysis reveals that this happens once the deflection reached one-third of d_{eff} (Bagci 2014; Chowdhury et al. 2005) because, at that point, the electrostatic force overwhelms the restorative force from the tensile stress in the membrane. This effect is the well-known pull-in instability that plagues many MEMS systems (Zhang et al. 2014).

Electrical Hamiltonian

Electro-mechano-optical transduction requires an electrical resonator formed by a capacitor, C_{T}, and inductor, L_{d}, combined in a single circuit-loop.[6] That loop

---

[4]Ignoring transients from switching on the bias. This simplification is self-consistent with Fourier-analysis (section 2.3) of the equations-of-motion.

[5]A uniform, electrostatic force implicitly assumes two things: uniform distribution of charges on the capacitor, which is valid for a small displacement of the membrane, and that contributions from the membrane-capacitor edge are negligible—and they are when the membrane-electrode distance is small (Carlson and Illman 1994).

[6]The eventual circuit involves other components, mainly related to filtering the bias, but Thevenin's theorem (Brittain and Brittain 1990) can convert them into an equivalent series impedance subsequently lumped into the total inductance and capacitance, which validates the circuit model. A well-designed circuit should have an impedance like a series LC circuit near its resonance frequency.

has the electric Hamiltonian:

$$\mathcal{H}_e(\phi, q, \beta) = \frac{\phi^2}{2\,L} + \frac{q^2}{2\,C_T(\beta)}, \tag{2.11}$$

where currents oscillating in the loop creates the charges q on the capacitor and the flux ϕ through the inductor. The electromechanical interaction derives from the mechanically compliant capacitance (section 2.3.3). Notice that this Hamiltonian is similar to a mechanical harmonic oscillator with the inductance L_d playing the role of mass, and the resonance frequency being

$$\Omega_{LC} := \frac{1}{\sqrt{L_d\,C_T}}. \tag{2.12}$$

2.2 Optical displacement detection

We measured the mechanical displacement with both optical interferometry and cavity detection. Both setups constitute weakly-coupled linear amplifiers (Clerk et al. 2010), amplifiers that have a whole quantum theory behind them to describe how the quantum noise affect the device under test through backaction. However, our setup is at room-temperature and dominated by thermal-noise, so the following analysis omits the quantum-noise backaction. The laser light is still quantized to account for the shot-noise in the final readout.

Most of our experiments measured the PSD of the mechanical motion—denoted $\mathcal{S}(\omega)$ with ω being the Fourier frequency—rather than a time-domain series measurement. So this section aims to express the optical readout in terms of the symmetrized[7], quantum-mechanical PSD, as defined by Clerk et al. (2010). Furthermore, both detection techniques probe the membrane at its center, which means the physical displacement corresponds to β for the effective mass in eq. (2.6).

2.2.1 Linearization

Three simplifications are useful for both interferometer and cavity readout: One, a unitary displacement transformation mapping the lasers coherent state[8] $|\alpha\rangle$

---

[7]We use the symmetrized PSD because they are akin to the classical PSDs obtained from real-valued measurements.

[8]An eigenstate to the annihilation operator, $\hat{a}$: $\hat{a}|\alpha\rangle = \alpha|\alpha\rangle$. Its average photon-number is $\bar{n} = \langle a^\dagger a\rangle = |\alpha|^2$.

onto a new vacuum state $|0\rangle$ and the corresponding annihilation operator, a, to (Clerk et al. 2010, App. G)

$$a \to \alpha + \delta a$$

where δa is the annihilation operator for the new vacuum. This transformation essentially turns the equations into quantum fluctuations δa around the mean, semi-classical field α with an average photon-flux of

$$\bar{n} = |\alpha|^2.$$

Two, a Taylor expansion to first order in the mechanical displacement β, an approximation that applies for the applicable limit of displacements much smaller than the laser wavelength, $\beta \ll \lambda$. Three, omitting cross-terms between β and δa as these are small compared to the accompanying cross-terms between α and β.

2.2.2 Interferometric readout

At first, we used an optical interference to measure the displacement of the membrane, thus transferring the mechanical modulation onto the optical carrier. The working principle is this: the membrane motion changes the optical path-length difference between two interfering beams and, in turn, change their interference pattern in a detectable way. The path-length difference corresponds to a phase-shift of θ between the two beams, for example,

$$\theta = \theta_0 + 2 \times \frac{2\pi}{\lambda}\beta \tag{2.13}$$

for a Michelson-Morley interferometer[9] that uses the membrane as a mirror and has a static phase-difference of θ_0. Other schemes have different path-length modulation and phase-offset. For instance, we also used a vibrometer for characterization, and that has a time-dependent phase.

Interference happens on a beamsplitter, giving two outputs described by (Gerry

---

[9]The factor of two is a property of this particular interferometric scheme, appearing because the beam has to traverse the same path twice before interfering.

and Knight 2005, Sec. 6.5)

$$n_1 = \frac{|\alpha|^2}{4} \left(1 + \cos(\theta_0) - \sin(\theta_0)\frac{4\pi}{\lambda}\beta \right)$$

$$+ \frac{i\,\alpha\,\left(1 + e^{i\theta_0}\right)}{2}\delta a_1^\dagger - \frac{i\,\alpha^*\,\left(1 + e^{-i\theta_0}\right)}{2}\delta a_1 \tag{2.14a}$$

$$n_2 = \frac{|\alpha|^2}{4} \left(1 - \cos(\theta_0) + \sin(\theta_0)\frac{4\pi}{\lambda}\beta \right)$$

$$+ \frac{\alpha\,\left(1 - e^{i\theta_0}\right)}{2}\delta a_2^\dagger + \frac{\alpha^*\,\left(1 - e^{-i\theta_0}\right)}{2}\delta a_2 \tag{2.14b}$$

after linearization. An apparatus would typically measure the difference between the outputs,

$$O = n_2 - n_1,$$

which leaves the symmetric quantum-PSD

$$\bar{S}_{OO}(\omega) = \left(\frac{4\pi}{\lambda}\right)^2 \bar{n}^2\,\bar{S}_{\beta\beta}(\omega) + \bar{n}, \tag{2.15}$$

evaluated at the point of highest sensitivity, i.e. $\theta_0 = \pi/2$ (modulo an integer of π), and simplified by the vacuum-noise commutator

$$[\delta a(t)_\nu, \delta a_{\nu'}^\dagger(t')] = \delta_{\nu,\nu'}\delta(t - t'). \tag{2.16}$$

Here, $\delta_{\nu,\nu'}$ denotes the Kronecker delta, and $\delta(t)$ is a delta-function.

Here is the point: the interferometric detection outputs a term directly proportional to the PSD of the membrane's displacement plus a constant background noise, a noise coming from the quantumness of the probing laser beam—the shot-noise of light. A complete theory includes the correlation between shot-noise and mechanical motion that arise from momentum-transfer between membrane and light, but the absence of an optical cavity makes it reasonable to neglect this effect. Moreover, the output may have additional sources of technical noise (e.g. fluctuations in laser power or frequency, or electronic noise from the detection noise), but shot-noise dominates a properly designed setup. Also, scaling the output linearly can turn it into units of mechanical displacement rather than photon-flux or detector voltage.

2.2.3 Cavity readout

Section 5.1 describes in detail the optical cavity we used for transduction experiments. That cavity had one fixed mirror while the other mirror was the membrane. The membrane is a harmonic oscillator, just like a mirror suspended on

a spring, which makes our physical setup identical to the canonical cavity optomechanical system. In contrast, many other optomechanical setups[10] only use the mirror-on-a-spring as an effective model. The membrane-mirror is a metal layer that absorbs ~5 % of the light, making the optomechanical cavity rather low-finesse and consequently putting it in the (extremely) unresolved sideband regime.

The cavity resonance is

$$\Omega_{\rm o}(\beta) = \frac{2\pi c}{l - \beta},\tag{2.17}$$

where the length of the cavity, l, is one optical wavelength. Despite being a very short cavity, that length is significantly bigger than the membrane displacement, making Taylor-expansion of the optical resonance a natural simplification. Input light to the cavity, $n_{\rm in}$, gets reflects with the reflection coefficient $R(\beta)$ that depends on the membrane displacement. That gives the linearized (section 2.2.1) cavity output:

$$
\begin{aligned}
n_{\rm out} &= R(\beta)\, n_{\rm in} \\
&\approx R(\beta = 0)\left(\bar{n}_{\rm in} + \alpha^*\,\delta a + \alpha\delta a^\dagger\right) + \left.\frac{\partial R(\beta)}{\partial \beta}\right|_{\beta=0}\bar{n}_{\rm in}\,\beta \\
&= R_0\left(\bar{n}_{\rm in} + \alpha^*\,\delta a + \alpha\delta a^\dagger\right) + R_0'\,\beta,
\end{aligned}\tag{2.18}
$$

where the last expression introduces the short-hand notation for cavity reflection, R_0, and the slope of the cavity, R_0'. The resulting PSD is

$$\bar{S}_{nn}^{\rm out} = R_0^2\,\bar{S}_{nn}^{in} + \left(R_0'\,\bar{n}_{\rm in}\right)^2 \bar{S}_{xx},\tag{2.19}$$

where $S_{nn}^{\rm out}$ is the measured PSD and S_{nn}^{in} is the PSD of the input light. Once again, the optical detection outputs a term directly proportional to the PSD of the membrane's displacement plus a constant background noise[11], and its expression neglects correlation between shot-noise and mechanical noise. We assume this simplification because the mechanical noise at room-temperature should dominate over any quantum backaction effects, and chapter 3 verify that the resulting theory works.

---

[10]Such as membrane-in-the-middle cavities, whispering-gallery resonators, and photonic crystal cavities.

[11]A more careful analysis include the vacuum-noise entering through the loss-mechanisms in the cavity, which increase the noise beyond laser shot-noise, but this may be lumped into the optical background noise.

Optomechanical dynamics

In addition to the neglected quantum-noise backaction, the optical cavity affects the dynamical behavior of the mechanical motion—an effect known as dynamical backaction. This effect originates from the optomechanical Hamiltonian, $\mathcal{H}_{\text{om}}$, i.e. (Aspelmeyer et al. 2014)

$$\mathcal{H}_{\text{om}}(a^{\dagger}, a, \beta) = \hbar \, \Omega_{\text{o}}(\beta) \left(a^{\dagger} a + \frac{1}{2} \right), \tag{2.20}$$

where a ($a^{\dagger}$) is the annihilation (creation) operator for intra-cavity photons. The equation corresponds to the interaction between the optical cavity and mechanical resonance, an interaction that arises from the radiation-pressure force,

$$F = \frac{\hbar \, \Omega_{\text{L}}}{l} \, \tilde{n}_{\text{cav}}, \tag{2.21}$$

exerted by the intra-cavity photons $\tilde{n}_{\text{cav}}$ on the moving mirror.

In addition to this effect, the membranes also absorb some of the incoming light and thus gets locally heated. We simulated the temperature increase to be negligible, but absorption gives rise to another optomechanical effect: photothermal backaction. Metzger and Karrai (2004) explain this effect as the photon-absorption changing the length of the cavity through, for instance, a mismatch of thermal expansion. The length-change decays with the thermalization time, and the delayed response causes a retarded force on the membrane that can both shift the mechanical frequency and add damping—more about that in section 2.3.1.

It is difficult to express absorption in an energy-conserving Hamiltonian. It would have to include both the internal degrees of freedom, excited by absorption, and the environment that equilibrates these degrees. Instead, Metzger et al. (2008a,b) proposed a more straightforward, classical model for the force, F_{ph}, created by the absorption, and their expression successfully described their observations of self-induced oscillation and photothermal cooling. We have also observed these effects, as explained in section 5.1.3 on page 92. De Liberato et al. (2011) and Pinard and Dantan (2008) later quantized this model to account for backaction from optical noise and arrived at

$$F_{\text{ph}} = \frac{\xi \, \hbar \, \Omega_0}{2 \, l} \int_{-\infty}^{\infty} \frac{dt'}{\tau} \exp\left(\frac{t - t'}{\tau} \right) \Theta(t' - t) \frac{a^{\dagger}(t') \, a(t')/\kappa_a}{\kappa_a} I_{\text{a}}(t'), \tag{2.22}$$

a result that they add ad hoc as a force that drives the membrane in the mechanical equations-of-motion. In this expression, κ_a is the rate with which the intra-cavity photons $a^{\dagger}(t') \, a(t')$ gets absorbed, and the Heavyside step-function $\Theta(t'-t)$

enforce causality for the absorption-induced deformation, a deformation that is assumed to relax at a rate τ—the thermal memory time. The constant ξ phenomenologically quantifies the deformation caused by absorption. Comparing eq. (2.22) to radiation pressure (eq. (2.21)) reveals that an instantaneous relaxation delivers a force equivalent to momentum transfer from absorption[12] when $\xi = 1$.

   While the photothermal might seem to lack rigor, it is justified through the timescales of the problem (De Liberato et al. 2011); the lifetime of thermal phonons are much shorter (picoseconds (Stedman and Nilsson 1966)) than thermal relaxation (milliseconds (Metzger and Karrai 2004)) and even membrane's period of oscillation (microseconds). That means all quantum effects of absorption decohere nearly immediately, and it suffices to treat the absorption as an average, semiclassical force proportional to the number of incident photons.

2.3 Equations-of-motion

Gardiner and Zoller (2004, ch. 3) have derived generic Heisenberg-Langevin equations-of-motion and expressed them through a generalized fluctuation-dissipation theorem.When applied to the electro-mechano-optical Hamiltonian (eq. (2.4) + eq. (2.11) + eq. (2.20)), their fluctuation-dissipation theorem gives:

$$\frac{\partial \beta}{\partial t} = \frac{p}{\tilde{m}}, \tag{2.23a}$$

$$\frac{\partial p}{\partial t} = -\tilde{m}\,\Omega_{\mathrm{m}}^2\,\beta - \tilde{m}\,\Gamma_{\mathrm{m}}\frac{\partial \beta}{\partial t} + F_{\mathrm{m}}(t)$$
$$\quad - \frac{\partial}{\partial \beta}\frac{q^2}{2\,C_{\mathrm{T}}(\beta)} - \hbar\left(a^\dagger a + \frac{1}{2}\right)\frac{\partial}{\partial \beta}\Omega_{\mathrm{o}}(\beta) + F_{\mathrm{ph}}(\beta), \tag{2.23b}$$

for the mechanical resonator, including the photothermal force eq. (2.22), and

$$\frac{\partial q}{\partial t} = \frac{\phi}{L_{\mathrm{d}}}, \tag{2.24a}$$

$$\frac{\partial \phi}{\partial t} = -\frac{q}{C_{\mathrm{T}}(\beta)} - L_{\mathrm{d}}\,\Gamma_{LC}\frac{\partial q}{\partial t} + F_{\mathrm{e}}(t), \tag{2.24b}$$

for the electrical resonator. The variable Γ_ν denotes the damping rate with which the system ν loses energy to its environment, while F_ν is a general driving force comprising signal, bias, and noise going into the system from its environment.

---

[12]Half the force from reflection.

To model the environment, Gardiner and Zoller (2004) assumed the system couples to a generic reservoir of harmonic oscillators, and that those oscillators are excited by thermal energy. That coupling permits excitations of energy to leak back and forth between both systems, thus generating fluctuating energy and establishing a thermal equilibrium.[13] This treatise technically applies to quantum systems where the environment noise serves to preserve quantum commutation relations. Nevertheless, the expressions extend to classical equations as a straightforward limiting case that replaces the quantum operators with their expectation-value. In the following, the classical limit applies immediately to the electrical and mechanical systems since they obey

$$k_B\, T_0 = \hbar\, \Omega_\nu, \qquad (2.25)$$

where k_B is Boltzmann's constant, $\hbar$ is Planck's constant, and T_ν (Ω_ν) is the temperature (resonance frequency) of the system denoted by ν.

When the reservoir is Markovian (memoryless[14]), the damping rate is a constant[15] while the thermal-fluctuations are a stochastic term with a delta-like auto-correlation function, i.e.

$$\langle F_\nu(t) \rangle = 0 \quad \text{and} \quad \langle F_\nu(t)\, F_{\nu'}(t') \rangle \propto \delta_{\nu,\nu'} \delta(t - t')\, 2\, k_B\, T_0\, m_\nu\, \Gamma_\nu. \qquad (2.26)$$

where angle-brackets signify an average. The fundamental underlying assumptions are that the noise does not change statistical properties in time[16] and that averaging many similar realizations equal one long realization[17].

Power spectral density

Expressing the PSDs in Fourier domain requires some finicky mathematical treatment (see for example Gillespie (1996)) to handle the Fourier transforms

---

[13]This model correspond to a system that can radiate energy to its environment and may fail for more intricate non-linear systems

[14]Meaning the reservoir do not remember the past leakage of energy to and from the system: it must lose self-correlation much shorter than one period of oscillation.

[15]The damping need not be time-independent velocity-damping. Importantly, both internal- and thermoelastic-damping have a time-dependence (Saulson 1990)). It is sufficient to have a low damping rate compared to the period of oscillation—i.e., an under-damped oscillator—where frequencies around resonance dominate the spectral response.

[16]A *stationary* process.

[17]An *ergodic* process.

of stochastic noise because such noise is generally not integrable[18]. From elaborate analysis comes these pertinent relations (adapted from Aspelmeyer et al. (2014, Sec. II.B.3)): the finite Fourier transform[19]

$$s_T(\omega) := \int_{-T/2}^{T/2} dt \, s(t) \, e^{-i\,\omega\,t}, \tag{2.27}$$

the classical PSD defined as

$$\bar{S} := \lim_{T \to \infty} \frac{|s_T(\omega)|^2}{T}, \tag{2.28}$$

the Fourier-transform of the auto-correlation function[20]

$$\bar{S} = \int_{-\infty}^{\infty} dt \langle s(0) \, s(t) \rangle \, e^{-i\,\omega\,t}, \tag{2.29}$$

and the PSD integral over all frequencies

$$s_{\mathrm{RMS}}^2 = \langle s^2 \rangle = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \, \bar{S}(\omega), \tag{2.30}$$

where s_{RMS} is the Root-Mean-Square (RMS) value of s.

Despite the problem with Fourier transforming noise, Equation (2.28) offers a strategy for theoretically handling the equations-of-motion (Haus 1984, Ch. 1): pretend the Fourier-transform exist, take the absolute-square value, normalize appropriately, and take an average. Theoretical noise-terms may then be replaced with their PSDs. The following equations adopt the square-root PSD as equivalent to the Fourier-transform of noise, a non-rigorous choice that serves to simplify expressions while retaining only the bare essentials that deliver the correct end-results. The strategy for analysis also works for estimating the PSD from experimentally measured values, but with the Fourier transform replaced by its discrete equivalent, because an apparatus samples the measurements in intervals over a finite time-window.

The PSD is symmetric around zero frequency when it derives from a variable with real values, a criterion that is met by all measurements to come. Therefore, the PSD can be considered for positive frequencies only without loss of generality, although that *single-sided* PSD must have twice the amplitude to produce the same variance and power. From here on, measured PSDs refer to the single-sided spectrum.

---

[18]An integrable function $f(t)$ must have $\int_{-\infty}^{\infty} dt |f(t)|$ finitely bounded.

[19]The typical Fourier transform is the limiting case $s_{T \to \infty}$. Notice the sign-convention in the exponential function.

[20]The Weiner-Khinchin theorem

Generalized thermal noise

General thermal PSDs (single-sided) follows by plugging the generalized, high-temperature, low-frequency thermal-noise, eq. (2.26), into eq. (2.29). That results in

$$\bar{\mathcal{S}}_{\mathrm{m}} = 4 \, k_{\mathrm{B}} \, T_0 \, \tilde{m} \, \Gamma_{\mathrm{m}} \tag{2.31a}$$

$$\bar{\mathcal{S}}_{\mathrm{e}} = 4 \, k_{\mathrm{B}} \, T_0 \, L_{\mathrm{d}} \, \Gamma_{LC} \tag{2.31b}$$

for the mechanical and electrical resonator, respectively. The first equation corresponds to the Brownian motion of a harmonic oscillator. The second equation becomes Johnson-Nyquist noise by defining

$$R_{\mathrm{d}} := L_{\mathrm{d}} \, \Gamma_{\mathrm{m}}, \tag{2.32}$$

an ohmic resistance that is in series with the *LC* resonance and corresponds to the damping of the resonator.

2.3.1 Optomechanical backaction

The interaction between cavity and membrane correspond to a dynamical backaction, effectively redefining the membrane frequency and damping rate. The effect may be derived by combining both equations in eq. (2.23) and taking the Fourier transform eq. (2.27) with $T \to \infty$, giving

$$\chi_{\mathrm{m}}(\omega)^{-1} \, \beta(\omega) = F_{\mathrm{m}}(\omega) + \tilde{G}_{\mathrm{om}}(\omega, \beta) \, |\alpha|^2, \tag{2.33}$$

with laser-fluctuation again assumed to be much smaller than the thermal noise in the membrane. The equation introduces χ_{m} as the mechanical susceptibility

$$\chi_{\mathrm{m}}(\omega) = \left[\tilde{m} \left(\Omega_{\mathrm{m}}^2 - \omega^2 + i \, \omega \, \Gamma_{\mathrm{m}} \right) \right]^{-1} \tag{2.34}$$

and $\tilde{G}_{\mathrm{om}}$ as a parameter describing the optomechanical coupling, including all optical effects, i.e.

$$\tilde{G}_{\mathrm{om}} \approx \hbar \left(\frac{\partial \Omega_{\mathrm{o}}(\beta)}{\partial \beta} + \frac{\xi}{2 \, \kappa_a \, l \, (1 + i \, \omega \, \tau)} \right) \, |\chi_{\mathrm{o}} \Omega_{\mathrm{o}}(\beta)|^2$$

assuming $|\alpha|^2 \gg 1/2$. Here, χ_{o} is the optical susceptibility, relating the intracavity operators to the input light. The first term in $\tilde{G}_{\mathrm{om}}$ corresponds to radiation pressure and the second to the photothermal force.

The optical effects may be collected into the mechanical susceptibility after two simplifications: First, replacing ω with Ω_m in $\tilde{G}_{om}$. That applies when $\Gamma_m \ll \Omega_m$ because the membrane responds the strongest to frequencies between $\Omega_m \pm \Gamma_m$. Second, separating a constant displacement from the equation-of-motion by the substitution $\beta(t) = \beta_0 + \delta\beta(t)$. That turns the mechanical susceptibility into

$$\chi_{om}(\omega) \approx [2\,\tilde{m}\,\Omega_m\,(\Omega_{om} - \omega + i\,\Gamma_{om}/2)]^{-1}, \tag{2.35}$$

with the resonance and damping rate renormalized according to

$$\Omega_{om} = \Omega_m - \mathrm{Re}\left(\left.\frac{\partial \tilde{G}_{om}(\Omega_m, \beta)}{\partial\beta}\right|_{\beta=\beta_0}\right) \quad \text{and}$$

$$\Gamma_{om} = \Gamma_m - \mathrm{Im}\left(\left.\frac{\partial \tilde{G}_{om}(\Omega_m, \beta)}{\partial\beta}\right|_{\beta=\beta_0}\right). \tag{2.36}$$

The full response looks like

$$\chi_{om}(\omega)^{-1}\,\delta\beta(\omega) = F_m(\omega) \quad \text{and} \tag{2.37}$$
$$\chi_{om}(0)^{-1}\,\beta_0 = \tilde{G}_{om}(0, \beta = \beta_0)\,|\alpha|^2.$$

Notice that the nomenclature uses the subscript 'om' for mechanical parameters that include optomechanical effects, and the subscript 'm' for the intrinsic mechanical values.

The imaginary part of $\tilde{G}_{om}$ changes the damping rate of the membrane. When radiation pressure dominates, an imaginary part can only come from χ_o, and it tends to zero in our limit of unresolved mechanical sidebands. In contrast, the photothermal force has a time-delay and can, therefore, change the damping even in the bad-cavity limit. Furthermore, the imaginary part can be both positive and negative because of the derivative of $|\chi_o|^2$, a derivative that changes sign when the laser frequency crosses the cavity resonance.

Effective optomechanical temperature

With optomechanical backaction, the membrane's PSD becomes

$$\bar{S}_{\delta\beta} = |\chi_{om}(\omega)|^2 \times 4\,k_B\,T_0\,\tilde{m}\,\Gamma_m$$
$$= \frac{k_B\,T_0}{\tilde{m}\,\Omega_{om}^2} \frac{\Gamma_m}{(\Omega_{om} - \omega)^2 + (\Gamma_{om}/2)^2}. \tag{2.38}$$

That spectral shape is a Lorentzian function, and it integrates up to the RMS displacement eq. (2.30) of

$$\langle \delta\beta^2 \rangle = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \bar{\mathcal{S}}_{\delta\beta\,\delta\beta} = \frac{k_B\,T_0}{\tilde{m}\,\Omega_{om}^2} \frac{\Gamma_m}{\Gamma_{om}} \tag{2.39}$$

around the constant offset β_0. This result inspires an effective optomechanical temperature[21] as

$$T_{om} := \frac{\Gamma_m}{\Gamma_{om}}\, T_{om}, \tag{2.40}$$

a definition that turns the RMS displacement into

$$\tilde{m}\,\Omega_{om}^2 \langle \delta\beta^2 \rangle = k_B\, T_{om}, \tag{2.41}$$

as expected from the equipartition evaluated for the mechanical harmonic oscillator, while the thermal PSD remains constant because it scales with the product between temperature and linewidth and $T_0\,\Gamma_m = T_{om}\,\Gamma_{om}$.

Evidently, the optomechanical interaction corresponds to heating or cooling a mode of mechanical motion. Such cooling can help to bring the effective mechanical temperature down until the membrane's thermal occupation becomes less than one, and quantum noise becomes prevalent. This ground-state is supposedly within reach from photothermal backaction alone (De Liberato et al. 2011; Restrepo et al. 2011), or at least nearly reachable (Pinard and Dantan 2008), for suitably optimized parameters. Oppositely, the backaction can bring the mechanical linewidth to zero at which point the system becomes unstable and start to oscillate (Metzger et al. 2008a; Zalalutdinov et al. 2001).

2.3.2 Data-analysis

Nearly all ensuing data-analysis fits a Lorentzian function to the measured PSD, thus extracting the mechanical parameters of resonance frequency Ω_{fit}, linewidth Γ_{fit}, and area under the peak RMS_{fit}, an area that corresponds to the RMS variance of the mechanical motion, given in the units of the detector output. However, the optical readout also adds shot-noise of light as constant background noise, as explained in section 2.2, so the fit includes an offset SN_{fit}. This fitting function will be referred to throughout this thesis and is therefore stated here for

---

[21]A more general definition includes a temperature from the fluctuations induced by shot-noise of light, but this is neglected here as discussed in section 2.2.3.

completeness:

$$\mathcal{S}_{\text{fit}}(\omega) = \frac{RMS_{\text{fit}}^2}{2\pi} \frac{\Gamma_{\text{fit}}}{(\Omega_{\text{fit}} - \omega)^2 + (\Gamma_{\text{fit}}/2)^2} + SN_{\text{fit}}, \tag{2.42}$$

Importantly, the fit separates the mechanical and optical noise contribution through their different spectral shape. Furthermore, this fit-function is not symmetric around zero frequency and thus presume a single-sided PSD.

Q-factor

The damping rate is a vital parameter for any harmonic oscillator, but it often gets expressed as a Q-factor defined as

$$Q := \frac{\Gamma}{\Omega}. \tag{2.43}$$

The Q-factor always reference a specific system that should be clear from context, and implicitly refer to the resonance frequency of that system, making the damping rate Γ and Q freely interchangeable.

If multiple sources limit the Q-factor, each of them corresponding to an intrinsic Q_ν, the total Q_{t} obeys the relation:

$$Q_{\text{t}}^{-1} = \sum_\nu Q_{\text{n}}^{-1}. \tag{2.44}$$

2.3.3 Electromechanical system

Since the optical backaction only affects the mechanical motion, it may be accounted for as above by rescaling the membrane parameters. In contrast, the electromechanical interaction requires more elaborate consideration. With the optomechanical rescaling, the full equations-of-motion, eqs. (2.23) and (2.24), reduce to

$$\tilde{m}\left(\Omega_{\text{om}}^2 + \frac{\partial^2}{\partial t} + \Gamma_{\text{om}}\frac{\partial}{\partial t}\right)\beta = F_{\text{m}}(t) - \frac{q^2}{2}\frac{\partial}{\partial\beta}\frac{1}{C_{\text{T}}(\beta)}, \tag{2.45a}$$

$$\left(\frac{1}{C_{\text{T}}(\beta_0)} + L_{\text{d}}\frac{\partial^2}{\partial t} + R_{\text{d}}\frac{\partial}{\partial t}\right)q \approx F_{\text{e}}(t) - q\,\delta\beta\,\frac{\partial}{\partial\,\delta\beta}\frac{1}{C_{\text{T}}(\beta_0)}\bigg|_{\beta=\beta_0}, \tag{2.45b}$$

where the last expression anticipates a future linearization by introducing the first-order expansion of $C_{\text{T}}(\beta)$ to highlight the similarity between the expressions. This approximation presumes the substitution $\beta(t) \rightarrow \beta_0 + \delta\beta(t)$ like above (section 2.3.1).

Electro-mechano-optical transduction requires an external bias to operate, a bias that enters through the driving force F_e together with both the signal to be transduced and the Johnson noise, i.e.

$$F_e \rightarrow F_e^{\text{Johnson}} + V_{\text{signal}} + V_{\text{bias}}.$$

Since delivering the bias requires a channel through which energy can enter the resonator, energy can also leak out through said channel which makes it akin to adding an ohmic loss in the *LC* circuit, an ohmic loss that adds noise according to the fluctuation-dissipation theorem. Technically speaking, the bias loads the resonance. The same can be said for the signal. A good circuit design prevents the loading from the bias by using, for example, filters with low impedance at the bias frequency and high impedance at the detection frequency, thus decoupling the bias from the relevant circuit resistance. The following analysis assumes such an ideal filter is in place, and that the *LC* circuit model works well within the linewidth of the resonance. While not practically possible, residual impedances may be lumped into the model constants. Alternatively, Takeda et al. (2017) present a theory that tracks these loss-channels to estimate the ultimate transduction performance including loading.

The next step is to linearize further around the external bias by splitting q into a bias term and a small perturbation. This substitution is slightly different for AC and DC biasing, so the next subsections treat the two situations separately. Nevertheless, both biases lead to similar results. For instance, both types of bias give rise to a static deflection of the membrane, as considered in section 2.1.2, which motivate the linearization of β already introduced. Both biases also cause switching transients until the system settles on the static equilibrium, but the treatise below neglects this effect by assuming the bias switched on in the distant past[22], thus permitting a straight Fourier transform.

DC bias

The DC bias, V_{DC}, induce a static charge on the capacitor, motivating the linearization $q \rightarrow q_0 + \delta q$. Introducing that in eq. (2.45) splits the equations-of-

---

[22]Ideally, the bias turned on at $t = -\infty$, but it practically suffices to have the bias turned on for much more than $1/\Gamma_{\text{om}}$

motion into a static and dynamic part. The electrical response is

$$\frac{q_0}{C_\text{T}(\beta_0)} = V_\text{DC}, \tag{2.46a}$$

$$\chi_\text{e}(\omega)\,\delta q \approx F_\text{e} + \delta V + \delta\beta\,\frac{V_\text{DC}}{C_\text{T}(\beta_0)}\,\frac{\partial\,C_\text{T}(\beta_0)}{\partial\,\delta\beta}\bigg|_{\beta=\beta_0}, \tag{2.46b}$$

where δV is the electrical signal to be transduced, $\chi_\text{e}(\omega)$ is the electrical susceptibility that relates to the electrical impedance as

$$Z(\omega) = \frac{\chi_\text{e}(\omega)}{i\,\omega} = i\,\omega\,L_\text{d} + R_\text{d} + \frac{1}{i\,\omega\,C_\text{T}}.$$

The mechanical motion is

$$\tilde{m}\Omega_\text{om}^2\,\beta_0 = \frac{q_0^2}{2\,C_\text{T}(\beta_0)^2}\frac{\partial\,C_\text{T}(\beta)}{\partial\,\beta}\bigg|_{\beta=\beta_0}, \tag{2.47a}$$

$$\left(\chi_\text{om}(\omega) + \frac{q_0^2}{2}\,\frac{\partial^2}{\partial\,\delta\beta^2}\,\frac{1}{C_\text{T}(\beta)}\bigg|_{\beta=\beta_0}\right)\delta\beta \approx F_\text{m}(\omega)$$

$$+ \delta q\,\frac{V_\text{DC}}{C_\text{T}(\beta_0)}\,\frac{\partial\,C_\text{T}(\beta)}{\partial\,\delta\beta}\bigg|_{\beta=\beta_0}, \tag{2.47b}$$

where the term added to the susceptibility may be absorbed into the mechanical frequency just like the optomechanical effects. Assuming the frequency-adjustment is much smaller than the mechanical resonance, the new mechanical frequency is

$$\Omega_\text{eom} = \Omega_\text{om} + \frac{q_0^2}{4\,\tilde{m}\,\Omega_\text{om}}\,\frac{\partial^2}{\partial\,\delta\beta^2}\,\frac{1}{C_\text{T}(\beta)}\bigg|_{\beta=\beta_0}. \tag{2.48}$$

This frequency-shift scales with the bias voltage squared, and it brings the membrane's resonance down. The two dynamical equations correspond to two coupled harmonic oscillators, defining the electromechanical coupling constant

$$G_\text{em}^\text{DC} := \frac{V_\text{DC}}{C_\text{T}(\beta_0)}\,\frac{\partial\,C_\text{T}(\beta)}{\partial\,\delta\beta}\bigg|_{\beta=\beta_0}, \tag{2.49}$$

a constant that gets enhanced by the DC bias.

Relating the displacement to frequency shift

Section 5.1 describes an experiment that scans the length of the optical cavity by applying a DC bias to deflect the membrane statically. The analysis relies on a relation between the deflection and the mechanical frequency-shift given here. The derivation assumes the membrane-capacitor has a large capacitance C_t in parallel, which means

$$\frac{1}{C_T(\beta)} \approx \frac{1}{C_t} + \frac{C_m(\beta)}{C_t^2}$$

with $C_m(\beta)$ defined in eq. (2.7). That reduces the scaling of the frequency-shift to the second derivative of the membrane capacitance, while the static deflection scales with the first derivative. These derivatives have the relation given in eq. (2.8), a relation that eq. (2.10) evaluated explicitly for circular membranes. When inserted into the electromechanical equation-of-motion, these two approximations give

$$\frac{\beta_0}{d_{\mathrm{eff}}} \approx -1.6\frac{\Omega_{\mathrm{om}} - \Omega_{\mathrm{eom}}}{\Omega_{\mathrm{om}}}. \tag{2.50}$$

This result also describes a square membrane accurately.

AC bias

AC biasing makes the induced charges oscillate at the bias frequency,

$$V_{\mathrm{bias}} = V_{\mathrm{AC}} \cos(\Omega_{\mathrm{ac}}\, t + \phi),$$

suggesting a linearization of the form

$$q_{\mathrm{AC}} \rightarrow \bar{q}_0\, \cos(\Omega_{\mathrm{ac}}\, t) + \delta q(t).$$

Notice the choice of phase ϕ avert a phase-shift on the charges q, a choice that simplifies the following notation. Making these substitutions in the equations-of-motion eq. (2.45), and linearizing like before, produce these steady-state equations:

$$\chi_e(\Omega_D, \beta_0)\, \bar{q}_0 = V_{\mathrm{AC}} e^{i\,\phi} \quad \text{and} \tag{2.51}$$

$$\tilde{m}\, \Omega_{\mathrm{om}}^2\, \beta_0 = \frac{\bar{q}^2/2}{2\, C_T(\beta)^2} \frac{\partial\, C_T(\beta)}{\partial\, \beta}\bigg|_{\beta=\beta_0}, \tag{2.52}$$

where the numerator $\bar{q}^2/2$ gets divided by two because it is an average off $(\bar{q}_0\, \cos(\Omega_{\mathrm{ac}}, t))^2$.

Fourier transforming transforms the induced charges as

$$q_{\text{AC}} \rightarrow \bar{q}_0/2 \left(\delta(\omega + \Omega_{\text{ac}}) + \delta(\omega - \Omega_{\text{ac}}) \right),$$

resulting in the dynamical equations:

$$\chi_{\text{e}}(\omega)\,\delta q(\omega) \approx F_{\text{e}}(\omega) + \delta V(\omega) + G_{\text{em}}^{\text{AC}} \left(\delta\beta(\omega + \Omega_D) + \delta\beta(\omega - \Omega_d) \right), \quad (2.53)$$

$$\tilde{\chi}_{\text{om}}(\omega)\,\delta\beta \approx F_{\text{th}} + G_{\text{em}}^{\text{AC}} \left(\delta q(\omega + \Omega_D) + \delta q(\omega - \Omega_d) \right), \quad (2.54)$$

where the electromechanical coupling constant $G_{\text{em}}^{\text{AC}}$ is identical to $G_{\text{em}}^{\text{DC}}$ save for a factor of two,

$$G_{\text{em}}^{\text{AC}} := \frac{\bar{q}}{2\,C_{\text{T}}(\beta)^2} \frac{\partial}{\partial\,\delta\beta} C_{\text{T}}(\beta)\Big|_{\beta=\beta_0}, \quad (2.55)$$

and where the mechanical susceptibility $\tilde{\chi}_{\text{om}}(\omega)$ now includes dynamical backaction from both optics and electronics through the frequency-shift,

$$\Omega_{\text{eom}} = \Omega_{\text{om}} + \frac{\bar{q}_0^2/2}{4\,\tilde{m}\,\Omega_{\text{om}}} \frac{\partial^2}{\partial\,\delta\beta^2} \frac{1}{C_{\text{T}}(\beta)}\Big|_{\beta=\beta_0}. \quad (2.56)$$

Both membrane and circuit only respond around their resonance frequencies and by design $\Omega_{LC} \approx \Omega_{\text{om}} + \Omega_{\text{ac}}$. That means the fast-oscillating terms $x(\omega + \Omega_{\text{ac}})$ and $q(\omega - \Omega_{\text{ac}})$ average out[23], which leaves

$$\chi_{\text{e}}(\omega)\,\delta q(\omega) \approx F_{\text{e}}(\omega) + \delta V(\omega) + G_{\text{em}}^{\text{AC}}\,\delta\beta(\omega - \Omega_d), \quad (2.57\text{a})$$

$$\tilde{\chi}_{\text{om}}(\omega)\,\delta\beta \approx F_{\text{m}}(\omega) + G_{\text{em}}^{\text{AC}}\,\delta q(\omega + \Omega_D). \quad (2.57\text{b})$$

These expressions are nearly identical to DC coupling (eqs. (2.46) and (2.47)), the difference being that the coupling constant, frequency-shift, and static displacement gets divided by two.

For completeness, here is the electromechanical coupling expressed as a coupling rate:

$$g_{\text{em}} := \frac{G_{\text{em}}}{\sqrt{L\,\Omega_e\,m\,\Omega_m}}, \quad (2.58)$$

a quantity that has a unit of angular frequency and that is typically used to characterize degenerate harmonic oscillators.

---

[23]A rotating-wave approximation.

2.4 Transduction performance

Both the opto- and electromechanical interaction correspond to two coupled harmonic oscillators. The optomechanical dynamics is in a limit of small coupling where the mechanical frequency is much smaller than the cavity frequency and linewidth, $\Omega_{om} \ll \Gamma_o, \Omega_L$. In contrast, the electromechanical dynamics have much stronger coupling, and the mechanical frequency is much larger than the electrical linewidth, $\Omega_{om} \gg \Gamma_{LC}$. As shown in section 3.1, the electromechanical system even extends into the strong-coupling regime.

The derivation below shows that the electromechanical interaction suppresses the mechanical noise relative to the electrical noise driving the membrane's motion, a motion that is hybridized into coupled electrical-mechanical dynamics. This reduction leads to an intrinsic noise-temperature, a central figure-of-merit that enable comparison with existing technologies offering amplification of weak electronic signals. Another vital parameter is the bandwidth of transduction. Both parameters derive from the electromechanical cooperativity, and the cooperativity is easily obtained from fits to measured PSDs.

The following sections start by defining the cooperativity and then expressing the noise-temperature and bandwidth. The analysis assumes an AC bias but is easily extended to DC by this substitution:

$$\Omega_{ac} \to 0, \quad \text{and} \quad \Omega_{LC} = \Omega_{om},$$

and by increasing the electromechanical coupling, frequency-shift, and static deflection by two as explained above.

2.4.1 Electromechanical cooperativity

Solving eq. (2.57a) gives the following driven equations-of-motion for the membrane and circuit, respectively:

$$\left(\chi_{om}^{eff}(\omega)\right)^{-1} \delta\beta = F_{th}(\omega) + G_{em}\,\chi_e(\omega + \Omega_{ac})F_e(\omega + \Omega_{ac}) \tag{2.59a}$$

$$\left(\chi_e^{eff}(\omega)\right)^{-1} \delta q = F_e(\omega) + G_{em}\,\tilde{\chi}_{om}(\omega - \Omega_D)F_m(\omega - \Omega_D), \tag{2.59b}$$

where F_e includes both Johnson noise and signal. The left-hand sides introduce the effective susceptibilities

$$\left(\chi_{\mathrm{om}}^{\mathrm{eff}}(\omega)\right)^{-1} := \tilde{\chi}_{\mathrm{om}}(\omega)^{-1} - G_{\mathrm{em}}^2\, \chi_e(\omega + \Omega_{\mathrm{ac}}), \tag{2.60a}$$

$$\left(\chi_e^{\mathrm{eff}}(\omega)\right)^{-1} := \chi_e(\omega)^{-1} - G_{\mathrm{em}}^2\, \tilde{\chi}_{\mathrm{om}}(\omega - \Omega_D). \tag{2.60b}$$

These two equations have a simple result at the resonance condition $\Omega_{\mathrm{ac}} + \Omega_{\mathrm{om}} = \Omega_{LC}$, namely

$$\left(\chi_{\mathrm{om}}^{\mathrm{eff}}(\Omega_{\mathrm{om}}\right)^{-1} = i\,\tilde{m}\,\Omega_{\mathrm{eom}}\,\Gamma_{\mathrm{om}}\,(1 + C) \quad \text{and} \tag{2.61a}$$

$$\left(\chi_e^{\mathrm{eff}}(\Omega_{LC})\right)^{-1} = i\,L_d\,\Omega_{LC}\,\Gamma_{LC}\,(1 + C) \tag{2.61b}$$

where C is the *electromechanical cooperativity*,

$$C := \frac{G_{\mathrm{em}}^2}{\tilde{m}\,L_d\,\Omega_{\mathrm{eom}}\,\Omega_{LC}\,\Gamma_{\mathrm{om}}\,\Gamma_{LC}}, \tag{2.62}$$

a unit-less quantity characterizing the electromechanical interaction strength.

The electrical resonator is much broader than the mechanical one in our parameter regime, $\Gamma_e \gg \Gamma_m$. That means the substitution

$$\chi_e(\omega + \Omega_{\mathrm{ac}}) \to \chi_e(\Omega_{LC})$$

work well over the membrane's entire resonance peak, reducing the electromechanical interaction to a broadening of the mechanical linewidth—until that broadening becomes comparable with Γ_{LC}. Therefore, the cooperativity can be estimated from the broadening of the membrane linewidth using

$$\Gamma_{\mathrm{eom}} = \Gamma_{\mathrm{om}}(1 + C). \tag{2.63}$$

On the other hand, the electrical susceptibility only sees the membrane near the resonance condition because the mechanical susceptibility goes to zero faster than the *LC* response. Specifically, the electromechanical interaction reduces the current-flow at the resonance condition where the electrical susceptibility increases. This effect is known as electromechanically-induced transparency (Agarwal and Huang 2010; Weis et al. 2010).

2.4.2 Noise temperature

Using the effective mechanical susceptibility eq. (2.61a), the cavity readout eq. (2.19) becomes:

$$\mathcal{S}_{nn}^{\text{out}}(\omega) = R_0\,\bar{n} + R_0'\bar{n}^2\,|\tilde{\chi}_{\text{om}}^{\text{eff}}(\omega)|^2\left(\bar{\mathcal{S}}_{\text{m}} + |G_{\text{em}}\,\chi_{\text{e}}(\omega + \Omega_D)|^2\bar{\mathcal{S}}_{\text{e}}\right) \quad (2.64)$$

As discussed in section 2.2.2, a similar expression applies to the interferometric readout, although it has different proportionality constant. Here, S_{e} includes both the signal and Johnson noise. The optical output is proportional to the input signal $\bar{\mathcal{S}}_{\text{e}}$. In voltage units, the output looks like

$$\begin{aligned}
\bar{\mathcal{S}}_{\text{e}}^{\text{out}}(\omega) &= \frac{\mathcal{S}_{nn}^{\text{out}}(\omega)}{R_0'\,\bar{n}^2\,|\tilde{\chi}_{\text{om}}^{\text{eff}}(\omega)\,G_{\text{em}}\,\chi_{\text{e}}(\omega + \Omega_D)|^2}\\
&= \bar{\mathcal{S}}_{\text{e}} + \frac{\bar{\mathcal{S}}_{\text{m}}}{|G_{\text{em}}\,\chi_{\text{e}}(\omega + \Omega_D)|^2}\\
&\quad + \left(\frac{R_0}{R_0'}\right)^2\frac{1}{\bar{n}\,|\tilde{\chi}_{\text{om}}^{\text{eff}}(\omega)\,G_{\text{em}}\,\chi_{\text{e}}(\omega + \Omega_D)|^2}
\end{aligned} \quad (2.65)$$

$$\Rightarrow \bar{\mathcal{S}}_{\text{e}}^{\text{out}}(\Omega_{LC}) = S_{\text{sig}}(\Omega_{LC})$$
$$+ 4\,k_{\text{B}}\,L_{\text{d}}\,\Gamma_{LC}\left(T_0 + \frac{T_{\text{om}}}{C}\frac{\Omega_e}{\Omega_m} + \frac{T_o\,(1 + C)^2}{C}\frac{\Omega_e}{\Omega_m}\right). \quad (2.66)$$

The second expression uses these simplifications: evaluating the output at the resonance condition

$$\Omega_{LC} = \Omega_{\text{ac}} + \Omega_{\text{eom}}$$

where the transducer is most sensitive, using eq. (2.61); inserting the thermal noises from eq. (2.26) plus a signal $S_{\text{sig}}(\omega)$ entering through $\bar{\mathcal{S}}_{\text{e}}$; and introducing the effective optical temperature

$$T_{\text{o}} := \tilde{m}\,\Gamma_{\text{om}}\,\Omega_{\text{eom}}^2\,(k_{\text{B}}\,\bar{n})^{-1}(R_0/R_0')^2.$$

The readout eq. (2.66) shows that, on resonance, the mechanical and optical noise adds to the electrical noise, effectively increasing the noise-temperature by

$$T_{\text{N}} = T_0 + \frac{\Omega_e}{\Omega_m}\left(\frac{T_{\text{om}}}{C} + \frac{T_{\text{o}}(1 + C)^2}{C}\right), \quad (2.67)$$

a noise-temperature that decreases with cooperativity until the optical noise starts to dominate. In other words, transduction increases the temperature of the

total ohmic resistance in the circuit, irrespective of the resistor's value. In contrast, electronic amplifiers require a specific load impedance[24], and their specified noise-temperature often refer to this impedance.

Note that the noise-temperature eq. (2.67) includes the ratio between electrical and mechanical frequency, which accounts for the difference in thermal occupation at the phonon and photon frequencies. Furthermore, considering that $C \propto 1/\Gamma_{\mathrm{om}}$ and

$$\Gamma_{\mathrm{om}}\, T_{\mathrm{om}} = \Gamma_{\mathrm{m}}\, T_{\mathrm{m}}$$

the noise-temperature must be unaffected by optomechanical heating and cooling.

2.4.3 Optimal cooperativity

There is an optimal cooperativity, C_{opt}, where the membrane and laser shot-noise reaches a minimum, specifically

$$C_{\mathrm{opt}} := \sqrt{1 + \frac{T_m}{T_o}} = \sqrt{1 + \frac{k_B\, T_{\mathrm{om}}}{\tilde{m}\,\Gamma_{\mathrm{om}}\,\Omega_{\mathrm{eom}}^2}\left(\frac{R_0'}{R_0}\right)^2 \bar{n}} = \sqrt{1 + \frac{\bar{S}_{\delta\beta}(\Omega_{\mathrm{eom}})}{\bar{S}_{\delta\beta}^{\mathrm{opt}}}}, \quad (2.68)$$

where $\bar{S}_{\delta\beta}(\Omega_{\mathrm{eom}})$ is the peak of the thermally driven mechanics (2.38) without any electromechanical coupling, and $\bar{S}_{\delta\beta}^{\mathrm{opt}}$ is the optical noise converted into units of mechanical displacement. However, the actual conversion factor cancels out in the ratio between measured peak and background noise.

In later experiments, the optimal cooperativity will be significantly larger than one, $C_{\mathrm{opt}} \gg 1$, prompting further simplifications

$$T_{\mathrm{N}} \approx T_0 + \frac{\Omega_{LC}}{\Omega_{\mathrm{eom}}}\left(\frac{T_{\mathrm{om}}}{C} + T_{\mathrm{o}}\, C\right) \quad \text{and} \quad C_{\mathrm{opt}} \approx \sqrt{\frac{\bar{S}_{\delta\beta}(\Omega_{\mathrm{eom}})}{\bar{S}_{\delta\beta}^{\mathrm{opt}}}}. \quad (2.69)$$

This limit makes the mechanical- and optical noise contributions equal at the optimal cooperativity C_{opt}, which means T_{o} may be eliminated from the transducers noise-temperature through this rewrite:

$$T_{\mathrm{N}} = T_0 + \frac{\Omega_{LC}}{\Omega_{\mathrm{eom}}} \frac{T_{\mathrm{om}}}{C_{\mathrm{opt}}}\left(\frac{C_{\mathrm{opt}}}{C} + \frac{C}{C_{\mathrm{opt}}}\right). \quad (2.70)$$

---

[24]A circuit would usually include some impedance-matching network between the amplifier and source. In contrast, the transduction scheme does not require any impedance transformation.

Here is a simple strategy for estimating the achievable and actual performance of the transduction: measure the thermally-driven mechanics with and without electromechanical coupling by turning the bias on and off. The former measurement give the linewidth by fitting the PSD with eq. (2.42), and the optimal cooperativity follows from eq. (2.68); and the latter measurement give the broadened linewidth, and the operating cooperativity follows from eq. (2.63). Both cooperativities then deliver the noise-temperature with eq. (2.70).

2.4.4 Bandwidth

The above equations consider only the response exactly at resonant. The mechanical response goes to zero far away from resonance, which means the transduction stops working for large detuning. In between, the mechanical response exhibit features of coupled dynamics, most prominently the mode-broadening discussed above, and ultimately normal-mode splitting (Dobrindt et al. 2008) when the coupling gets large enough. Our work mostly considers the weak-coupling limit where the mechanical response strictly decreases away from resonance, meaning the transduction is most sensitive on resonance and work within a specific bandwidth. Mode-splitting makes things more complicated from the bandwidth perspective because, in that case, the membrane resonance is no longer the most sensitive frequency. The above noise-temperature eq. (2.70) still holds for the resonance condition, but the mechanical response peaks at the hybridized modes that split around Ω_{eom}—see section 3.1.2 on page 44 and figs. 3.6 and 3.12b.

Two definitions of bandwidth are relevant for the transduction: One, the frequency-interval where the signal's power reduces by two from its peak value—the $-3\,\text{dB}$, full-width half-max bandwidth—which equals the broadened mechanical linewidth. Second, the frequency-interval where the Signal-to-Noise Ratio (SNR) reduces by two—the SNR bandwidth—a concept that is more useful for sensing purposes. For the electro-mechano-optical transduction, the SNR bandwidth is larger than the broadened mechanical linewidth because the mechanical noise reduces together with the signal response. The bandwidth instead depends on the detailed balance between shot-noise and thermal noise. At the optimal cooperativity ($C_{\text{opt}} \gg 1$) in particular, the SNR bandwidth is

$$BW = \sqrt{2}\,\Gamma_{\text{eom}}, \tag{2.71}$$

provided $\Gamma_{\text{eom}} \ll \Omega_{\text{eom}}$.

Chapter 3

DC transduction

This chapter describes our first experiments with the electro-mechano-optical transduction. These experiments used a DC bias and interferometric optical readout—the optical cavity came later. At first, experiments took place at DTU and mostly used a strong drive-tone to measure the system response (section 3.1). However, we moved the setup to NBI after an interferometer had been assembled there to characterize membranes. Characterization done at NBI included measuring the intrinsic transducer noise, a measurement requiring considerable noise reduction (section 3.2). All experiments had the transducer and detection coil at room temperature.

This treatise aims to be a brief discussion of the setup and results. More details are given in Bagci (2014), Bagci et al. (2014), and Simonsen (2014), and preliminary testing of membrane-capacitors was presented in Bagci (2014) and Schmid et al. (2014). In parallel with our experimental efforts, a theoretical framework has been developed mainly by Emil Zeuthen, Jake Taylor, Albert Schliesser, and Anders Sørensen—see Zeuthen (2015) and Zeuthen et al. (2017).

3.1 Electromechanical interaction

3.1.1 Methods

Here is an account of the three subsystems that together constitute the transduction scheme, presented in this order: the electromechanical device, optical characterization, and electrical circuit, with special attention paid to the optimal circuit design and the issues with the membrane-capacitor assembly.

Membrane-capacitor chip

Figures 3.1a and 3.1b show microscope images of the membrane and electrode, respectively. The chips were both fabricated by our collaborators at DTU Nanolab. Together, they form the transducer used in the experiments below. The membrane was made from highly-stressed, stoichiometric Si_3N_4 silicon nitride—henceforth referred to as 'nitride'—deposited on a silicon carrier chip. As a membrane material, nitride offers low optical loss at relevant wavelengths and excellent mechanical Q-factor (Wilson et al. 2009). While it is possible to couple a nitride membrane through dielectric forces, Schmid et al. (2013) investigated this scheme and found the electromechanical interaction works much better with a conductive layer on the membrane. Best was adding graphene onto the nitride, but that is a considerable fabrication challenge. Alternatively, aluminum worked nearly as well due to its low density (for a metal) while being far easier to deposit on the membrane, and it works as a reflector for the optical displacement readout, making it a convenient compromise. Figure 3.1a displays the typical dimension and layout of both the aluminum and nitride layers forming the membrane, a combination that typically yielded a mechanical frequency of ~700 kHz, an effective-mass of ~45 ng, and a mechanical linewidth of 2 Hz to 20 Hz.

The aluminum layer has a particular pattern rather than covering the entire membrane, as shown in fig. 3.1a. For one, the metal layer has a hole in the center where the fundamental mode has its largest vibrational amplitude. This feature anticipated future experiments where the membrane would go into an optical cavity as in the membrane-in-the-middle configuration (Møller et al. 2017;
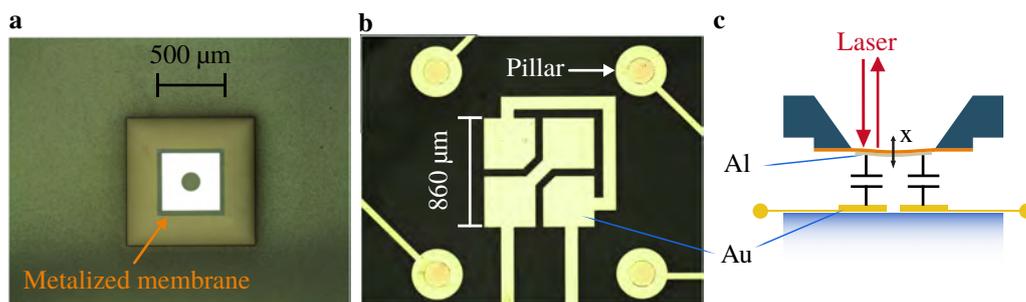
Figure 3.1: **a)** Nitride membrane in the center square, but the nitride is transparent. The white area in the square is the aluminum metallization. **b)** A four-corner, planar electrode at the center, surrounded by pillars shaped as small circles. **c)** Cross-sectional drawing of the assembled transducer with the membrane hovering above the electrodes, causing the indicated capacitive coupling. (*Adapted from Bagci et al. (2014).*)

Nielsen et al. 2017; Wilson et al. 2009). In the transduction measurements considered here, it would have been better to omit this hole. Furthermore, the aluminum does not cover the edge of the membrane because Yu et al. (2012) proved such a design preserves the nitride's excellent mechanical Q-factor, although that essentially means the metal cannot have a direct electrical connection to the actual circuit. Instead, we coupled the membrane capacitively to a planar electrode on a different chip by sandwiching the two chips together. The electrode consists of two segments[1]—see fig. 3.1b—that each creates a capacitance with part of the aluminum on the membrane, and those capacitors add in series as illustrated in fig. 3.1c. Best case scenario, the full capacitance is four times smaller than if it all had been one big capacitor because each segment divides the membrane area in two, and adding two identical capacitors in series divides the total capacitance by another factor of two.

Notice the electrode die in fig. 3.1b contains pillars that we hoped would support the membrane-chip, thus defining the distance between chips, when the membrane got placed on top of the electrode. However, that turned out to not be the case. By measuring the gap through the bias-induced frequency-shift[2], and even through directly through white-light interferometry[3], we learned that the gap was always higher than the pillars themselves—often several tens of micrometers whereas the pillars only were $0.5\,\mu m$ to $2\,\mu m$. After considerable effort, the problem got improved significantly but not completely eradicated. Since then, Haghighi et al. (2018) and Takeda et al. (2017) have both reported a membrane-electrode gap beyond their target design, thus verifying this technical difficulty. So at least the problem appears to be reproducible.

Figure 3.2 shows the final assembly, including the electrical connection to the electrode on the bottom chip implemented through spring-loaded, gold-probes pressed into contact with the metal layer that forms the planar capacitor. To assemble the device, we placed and moved the membrane by hand while monitoring under a microscope that the aluminum overlapped nicely with the electrode.

---

[1]The actual design splits the electrode into four smaller units, but they are connected two-by-two and therefore behaves like two segments. The particular choice of segmentation controls the coupling the mechanical motion and can thus discriminate between different modes of motion.

[2]I.e. evaluating eq. (2.47a) on page 29 with eq. (2.48) where eq. (2.8) on page 14 relates the first- and second derivative.

[3]We used a commercial vibrometer from Polytec (MSA-500) with this capability. The technique images the bottom electrode through the transparent nitride, while also imaging the top electrode. That results in a surface plot with a step going from aluminum to electrode where the step-height corresponds to the gap size.
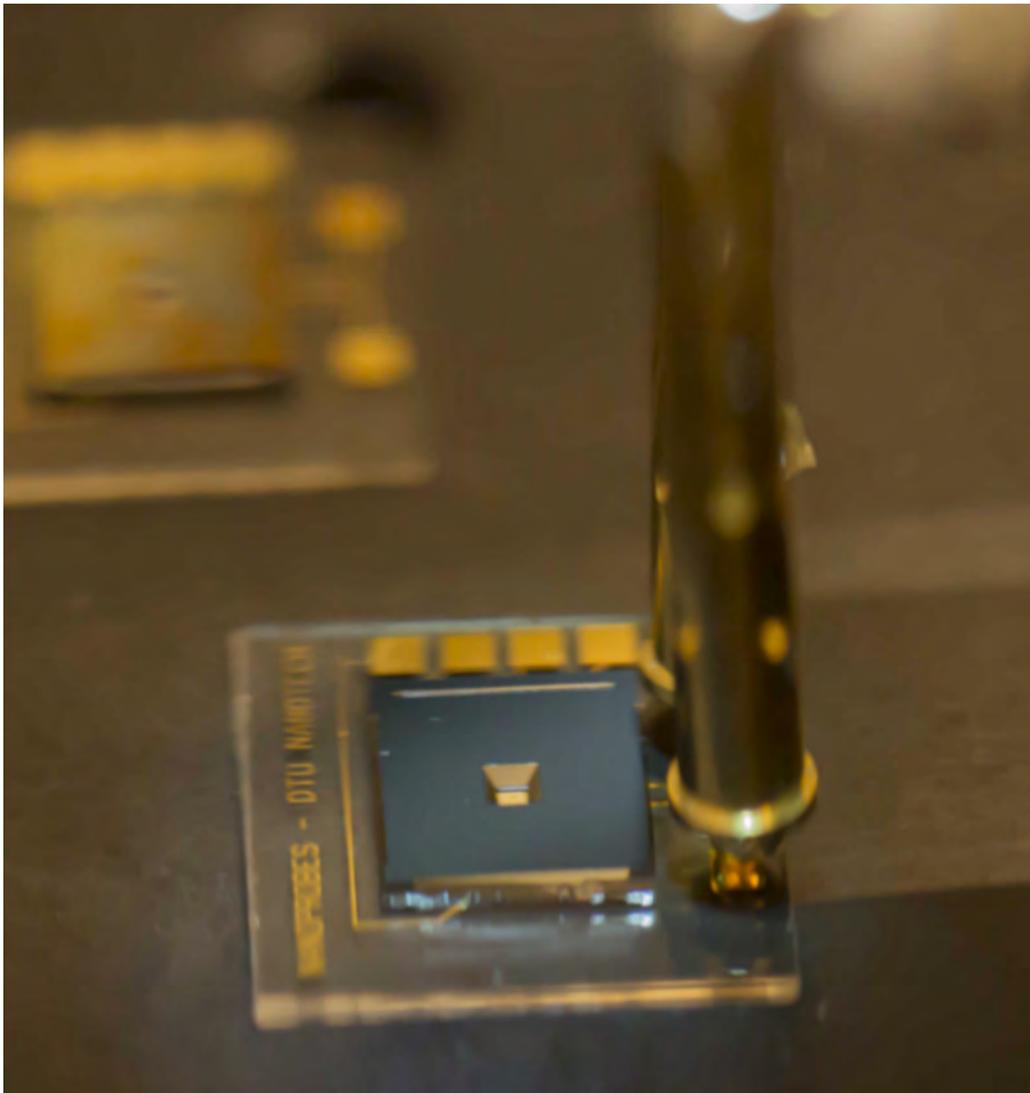
Figure 3.2: Photo of the membrane chip on top of the electrode chip, when mounted in the vacuum chamber and connected electrically. (*Reprinted from Bagci (2014).*)

Then it was transferred carefully to a vacuum chamber containing suitable electrical feed-throughs that permit connection to the gold-probes from outside the chamber. Throughout all transduction measurement, the vacuum level was well below 1×10^{-5} mbar.

Optical vibrometer

In the very first experiments, we converted the electrically-driven membrane-displacement into an optical signal by way of a commercial Doppler vibrometer (Polytec MSA-500), thus rounding out the transduction chain. Figure 3.3a presents the fundamental optical scheme behind the vibrometry. The instrument operates as a black box that provides spectra of the membrane motion, converted directly into units of displacement with an optical readout noise near $100 \, \text{fm}/\sqrt{\text{Hz}}$. The instrument also has software for analyzing the spectra, including a simple tool for fitting the resonance peak, as shown in fig. 3.3b. We used this tool[4] when extracting mechanical Qs or frequencies with the vibrometer in nearly all such measurements reported in this and subsequent chapters.

Furthermore, the vibrometer has an analog output proportional to the membrane motion. This output could be analyzed by external equipment, although turning the output into units of volts without retaining the calibration in displacement.

Aligning the vibrometer required moving the transducer into the laser beam

---

[4]After averaging the spectra twenty to thirty times.
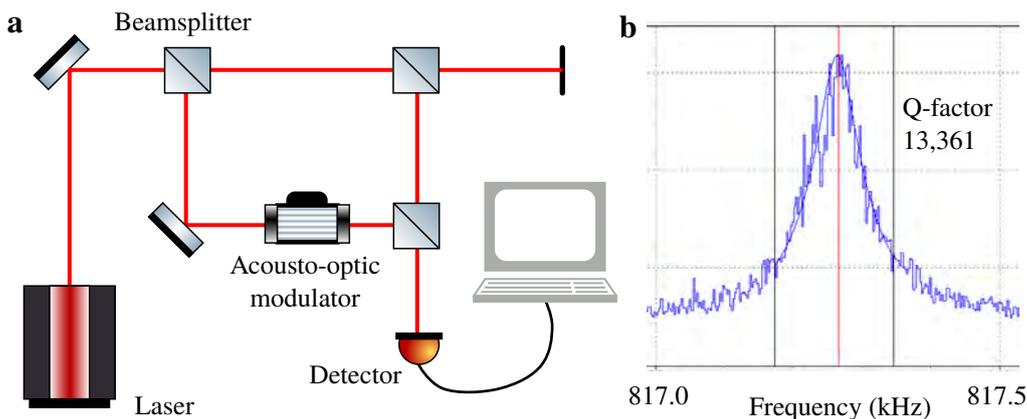
Figure 3.3: **a)** Principle setup behind Doppler vibrometry. **b)** Screen-shot example of recorded spectrum and fit using the instruments own software tool.

such that it gets reflected by the aluminum metallization on the membrane, and
then having the reflection travel back into the vibrometer. In practice, this mostly
meant aligning the focal-plane of the laser by turning a knob. Aligning the tilt of
the transducer was nearly never necessary, provided the sample was roughly flat
relative to the optical table. The vibrometer's laser passed through a view-port in
the vacuum chamber that housed the transducer, seemingly without affecting the
measurement adversely.

Electrical circuit

High electromechanical cooperativity, C, benefits the transduction (eq. (2.67) on
page 34), meaning the inductance L_d must be small, according to eq. (2.62) on
page 33. However, this is misleading; the circuit also must hit a given resonance
frequency,

$$\Omega_{LC} = \frac{1}{\sqrt{L_d\,C_T}},$$

which means a lower inductance necessitates a higher capacitance and vice versa—
in this case, that target frequency equals the membrane resonance somewhere
between 0.7 MHz to 0.9 MHz. Adding capacitance to the transducer reduces the
cooperativity, according to eqs. (2.49) and (2.55), and that reduction scales faster
than the inductance: C_t^4 compared to L_d. That means the optimal circuit has the
highest possible inductance and Q-factor.

When making a high-Q, LC tank circuit, the inductor is the most crucial com-
ponent, so this project started by investigating the best way to make them. Bagci
(2014) covers this conundrum in vivid detail. Here it suffices to say that the best
results used an inductor created from Litz-wire wound around a ferrite core[5].
The coil's inductance was typically ~500 μH and its intrinsic Q as high as 600.
Adding more circuit-components reduces the overall Q slightly, but the coil con-
tributed thee most ohmic loss until the transducer was connected. Then the tank
Q dropped to 100 to 130. Partly due to a pure resistivity measured across the
planar electrode, probably because the metal layer is thin. However, that value
does not account for the full Q reduction. Other possible sources of LC loading
are contact-losses, ohmic resistance from wires go to the transducer, and dielec-
tric losses from the nearby silicon die that hosts the membrane.

The transducer itself only has small capacitance, <0.5 pF, compared to the
~70 pF required to reach the frequency-target with the high-Q inductors. The

---

[5]Ferrite 61 (mixture of Zinc and Nickel) from Amidon.
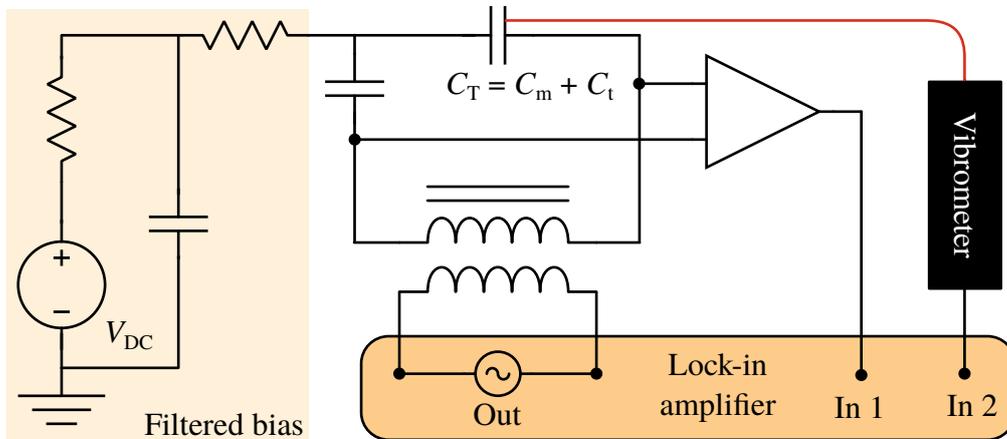
Figure 3.4: Principle circuit schematic. The depicted capacitance encompass all contributions: the transducer, tuning capacitor, and parasitic effects.

remaining capacitance, C_t, came from two sources: a tunable capacitor used to tune the exact LC resonance to match the fundamental membrane mode; and parasitic capacitances from, for example, the circuit board and cables running to and from the transducer. As argued above, the total capacitance should be as small as possible to give the highest cooperativity, and the final values were already at the limit of pure parasitic contributions despite our best effort to reduce those as much as possible.

Lastly, the transduction mandates a DC bias across the transducer. This bias came from a DC source connected to the LC tank as illustrated by the simplified schematic in fig. 3.4. This connection further included a low-pass filter for the DC bias and a current-blocking capacitor. The filter reduces noise leaking from the DC source into the electrical circuit, while the current-blockade prevent the bias from short-circuiting through the inductor. The blocking capacitor was much larger than the transducer plus tunings capacitors, giving a negligible impedance at the mechanical frequency.

To characterize the electromechanical interaction, we injected a signal into the system and measured how it reacted using the setup in fig. 3.4. The injection employed a small coil placed near the primary inductor to induce a voltage in series with the LC resonance through the mutual inductance between the two coils. The coupling was sufficiently small that the drive-coil does not shift the LC resonance frequency considerably, and the drive was strong enough to overwhelm all noises in the system. This drive signal both produces a voltage across the transducer and drives the mechanical motion, as described by the theory (sec-

tion 2.3.3). These responses were detected concurrently—the voltage detected with an operational amplifier in parallel to the transducer, and the mechanical displacement detected with the vibrometer—and sent to a lock-in amplifier for simultaneous recording, a lock-in that swept the signal- and detection-frequency in unison and gathered the ensuing frequency response.

3.1.2 Result and discussion

Thorough investigation revealed several culprits behind the distance issue. For example, when glued to a carrier, the electrode chips would curve significantly. Another inconsistent but prevalent culprit appeared to be dust because it often helped to clean the membrane and electrode with an air-gun, especially if the distance was abnormally large. Pressing the chips together would also help. In the end, the most notable improvement came by etching $50\,\mu m$ into most of the silicon chip, leaving only the original surface right around the membrane where the chip and pillars (should) make contact. With that, we got close to the designed gap of $1\,\mu m$ for some assemblies. Conversely, the samples also occasionally had the membrane in physical contact with the electrode. The outcome never got reproducible for this iteration of membranes and chips.

Figures 3.5a and 3.5b present the driven response of both the electrical and mechanical system, respectively. For these measurements, the distance issue was still not as good as it would get, which necessitated a high DC bias voltage to reach high electromechanical interaction. Figure 3.6 shows these responses for increasing DC bias, which increases the electromechanical interaction and shifts the mechanical frequency. For each measurement, the mechanical and electrical resonance was matched by tuning the circuit capacitance. As expected from theory, the voltage dips near the mechanical resonance, a feature that constitutes electromechanically-induced transparency (Agarwal and Huang 2010). Furthermore, fig. 3.5b proves that both the mechanical linewidth and resonance scales quadratically with bias as equation eqs. (2.48) and (2.63) claims.

The electromechanical coupling-rate plotted in fig. 3.6a scales linearly with bias as expected from eq. (2.58) on page 31. At a bias voltage in the strong-coupling regime, the coupled equations-of-motion exhibit an avoided crossing (Gröblacher et al. 2009) as demonstrated in fig. 3.6b. Here, the membrane frequency gets swept across the *LC* resonance using the frequency shift induced by the DC bias. Below the big scan, figs. 3.6c and 3.6d display a selected cross-section of the surface scan. These responses fit both electrical and mechan-
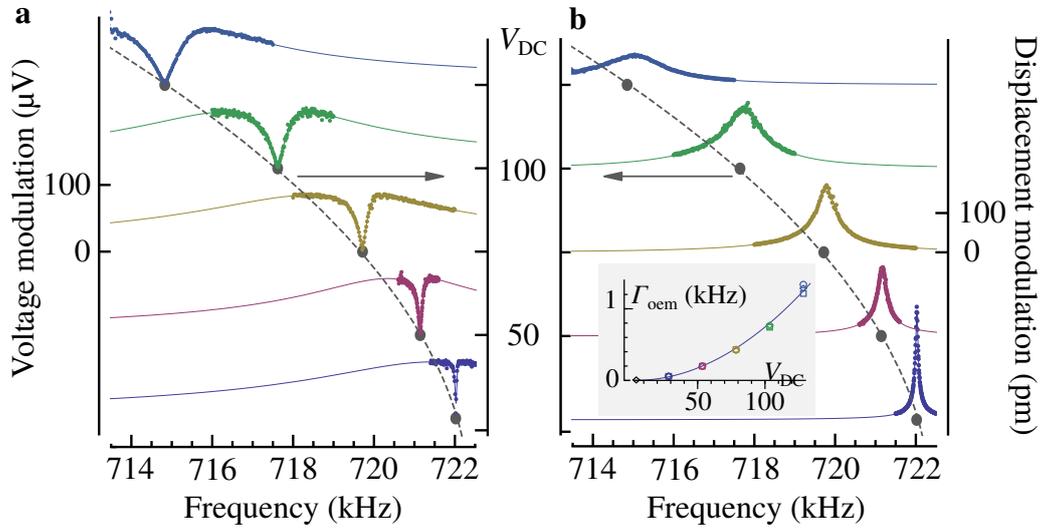
Figure 3.5: **a)** Voltage induced by the lock-in drive, measured in parallel to the transducer as in fig. 3.4. **b)** Mechanical displacement induced by the lock-in drive measured optically with the interferometer. Inset shows the fitted linewidth versus bias voltage. Parameters were: $\Gamma_{\mathrm{m}} = 2.3\,\mathrm{Hz}$, $\Gamma_{LC} = 5.5\,\mathrm{Hz}$, $\tilde{m} = 24\,\mathrm{ng}$, and $d_{\mathrm{eff}} = 5.5\,\mu\mathrm{m}$. (*Reprinted from Bagci et al. (2014).*)

ical equation-of-motion eq. (2.59), like the responses in fig. 3.5, and the fit-parameters for curve-pairs agrees within 1 %. That underpins the validity of the theory in section 2.3.3.
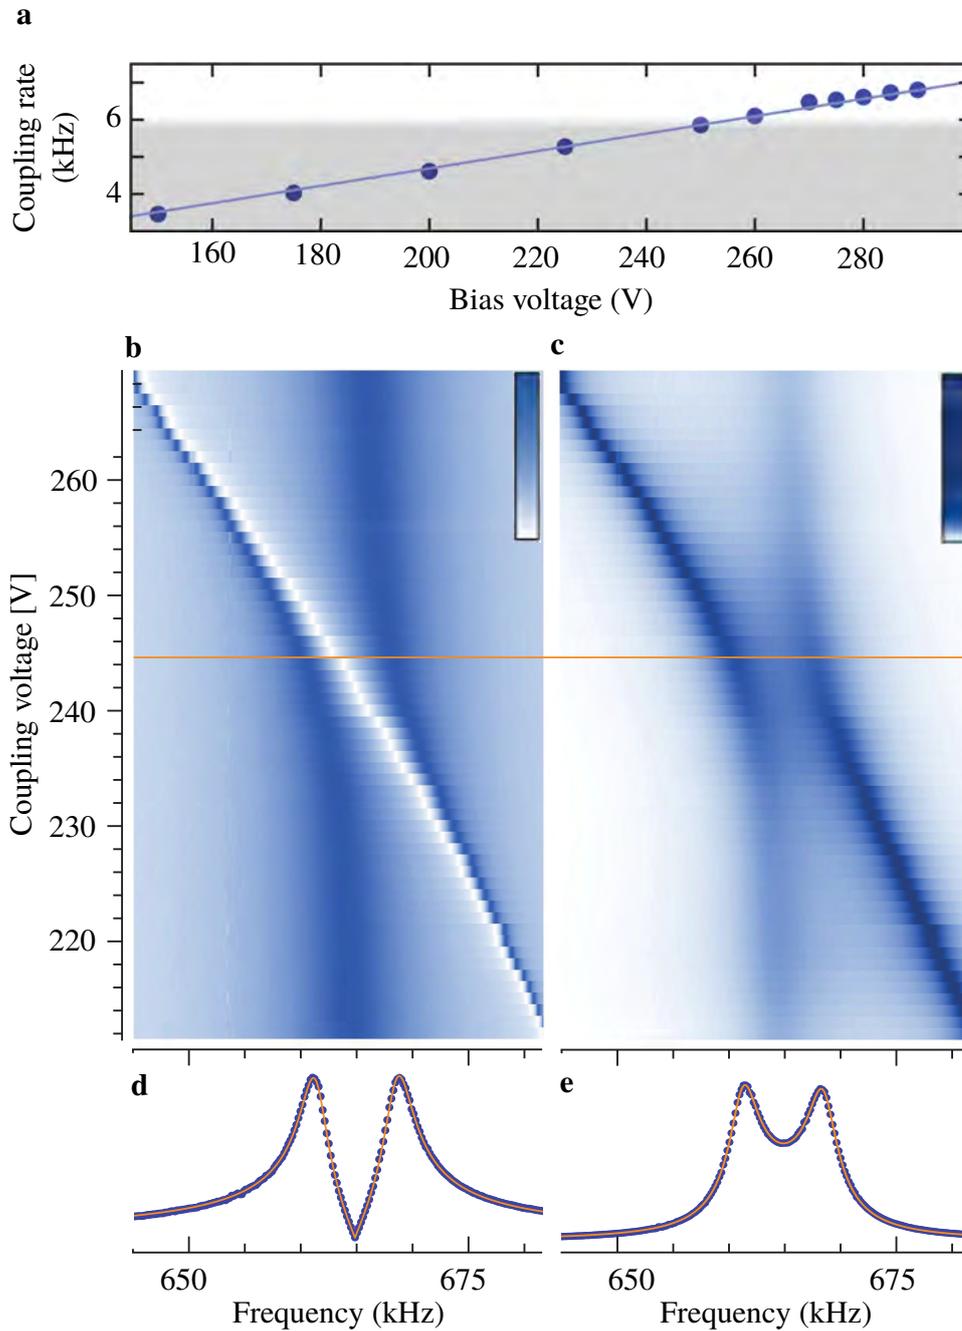
Figure 3.6: **a)** Coupling rate versus bias, entering the strong-coupling territory above the shaded area. **b+c)** Electrical (**b**) and mechanical (**c**) response at different bias voltages. The yellow line indicates the cross-sections shown below in (**d+e**); the electrical response being in (**d**) and the mechanical in (**e**). Parameters were: $\Gamma_\mathrm{m} = 3.1\,\mathrm{Hz}$, $\Gamma_{LC} = 5.9\,\mathrm{Hz}$, $\tilde{m} = 24\,\mathrm{ng}$, and $d_\mathrm{eff} = 4.5\,\mathrm{\mu m}$. (*Reprinted from Bagci et al. (2014).*)

3.2 Transduction sensitivity

Environment noise plagued the experiments presented thus far, so we had to improve the electronic setup considerably to see the intrinsic noise of the transduction. This section covers these improvements plus the Y-factor technique we used to measure the noise-temperature of the transducer after moving the experiments to NBI. This move entailed replacing the vibrometer with the interferometer described below. The membrane-capacitor remained the same as before.

Interferometer

A custom-made Mach-Zender interferometer replaced the vibrometer for optical measurements of mechanical motion done at NBI. Barg et al. (2018) have described the optical setup in full detail. Here is only a brief description: as shown in fig. 3.7, one arm of the interferometer goes to the membrane while the other goes to a reference mirror, and both objects retroreflect the light, making it interfere on two detectors—a standard homodyne detection design—with the power balanced between both detectors. When the membrane moves, it changes the optical path-length of one arm, thus shifting the interference slightly, which the balanced detector turns into a voltage signal. The voltage is directionally proportional to the displacement if said displacement is much smaller than a wavelength, and the interference pattern has equal power on both detectors. To ensure the latter condition, we locked the interference pattern at this point using a servo loop to drive a mechanical actuator (piezo) that displaces the mirror in the reference arm. This feedback was too slow to follow the membrane vibration.

The interferometer used a laser with a wavelength of 1064 nm, and it included a probe-head designed for this wavelength that both focused the laser onto the sample and imaged said sample with the laser-spot, enabling alignment on the correct membrane position. Figure 3.7 shows a schematic and photo of this probe-head. The sample was situated inside a vacuum chamber. With proper alignment, the interferometer routinely operated at the detection limit set by shot-noise of light. That limit corresponded to $1 \, \mathrm{fm}/\sqrt{\mathrm{Hz}}$ for our typical optical powers of around 1 mW shining on the membrane[6], and equal reflected power from both arms.

A calibration procedure was necessary to turn the measured modulation into actual displacement units. That procedure entailed a calibration peak with a

---

[6]About 30 % made it back to the detector.
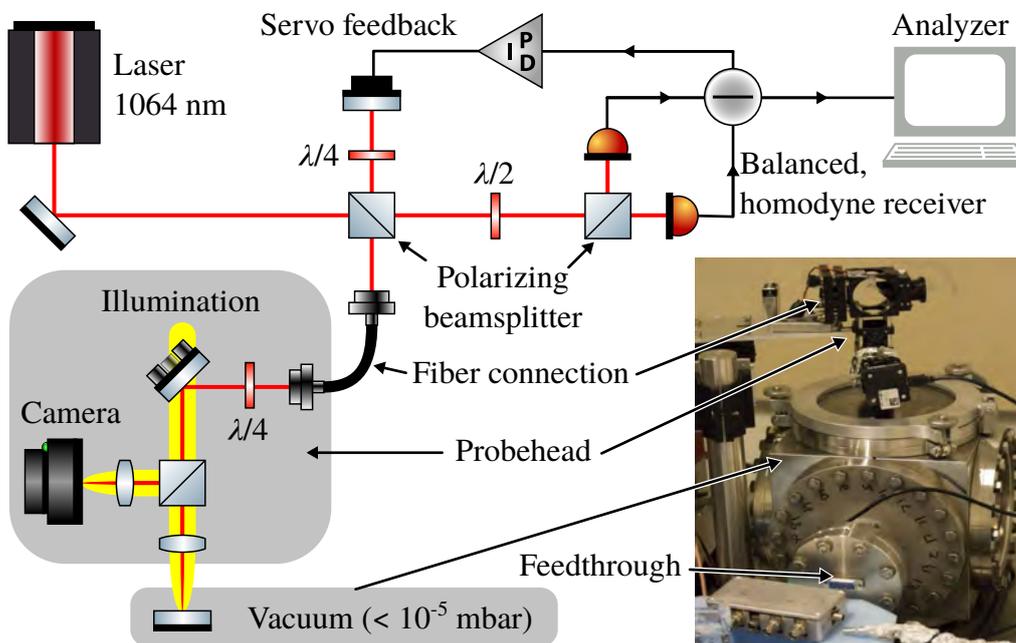
Figure 3.7: Simplified schematic of the interferometer setup and probe-head. The photo shows the probe-head mounted above the vacuum chamber.

known RMS displacement included in the signal such that the measured PSD could be normalized to yield the correct RMS value for the calibration peak. That procedure also canceled out any drift of the optical detection. The calibration peak came from driving the piezo in the reference arm at a frequency near the membrane. To convert it into displacement, we compared the driven RMS signal to the full interferometer fringe and used that the fringe corresponds to a path-length different of one full wavelength.

3.2.1 Y-factor extrapolation

Extrapolating intrinsic amplifier noise requires a dedicated setup that implements the standard Y-factor method (Horowitz and Hill 1989, ch. 7.19); the idea is to amplify known quantities of noise and then extrapolate how much readout noise remains when the amplifier has no input. That is easier said than done since the amplifier needs an input load to operate in a meaningful way, and such a load always carries intrinsic[7] noise, a noise that the Y-factor method must account

---

[7]Ultimately, it comes thermal excitations in the circuit, but electrical noise is ubiquitous and could come from many sources.
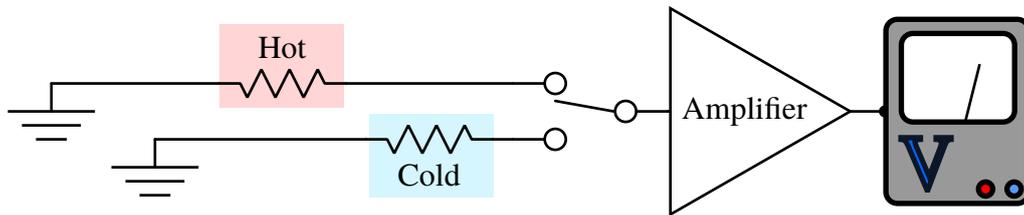
Figure 3.8: Prototypical Y-factor setup for electrical circuits.

for to extract the intrinsic noise of the amplifier. In this context, the amplifier's signal is the total output from the load, including its intrinsic noise.

Standard electrical amplifiers usually suffer from both current- and voltage noise, making the natural units of the problem an ambiguous choice. Worse, the dominant noise-source depends also on the load-impedance. Instead, the standard approach calibrates the noise-power in terms of noise-temperature, i.e., the temperature of the load impedance that delivers the same power of Johnson noise as the amplifier. Sometimes this noise-temperature gets converted into the standard, $50\,\Omega$ impedance no matter the actual load[8]. This thinking presents a direct path to estimate the actual noise-temperature, a path that fig. 3.8 illustrates: measure the amplifier output with its load resistor at two or more physical temperatures, and extrapolate to the equivalent (negative!) temperature that nullifies the total power of the noise-driven output. The extrapolation fits a straight line to the output noise, as exemplified in fig. 3.9. With this method, the units after amplification are irrelevant; it can be voltage or current for a standard amplifier, or mechanical displacement or optical modulation in case of the transducer.

Implementation

To do the Y-factor extrapolation in fig. 3.9, we used the circuit in fig. 3.10 to vary the temperature of the load resistor, R_L. Practically, this circuit was similar to the one in section 3.1.1, and we used the same small coil and electronic amplifier to register and fine-tune the LC resonance. However, after that tuning, we removed the second coil and disconnected amplifier and therefore excluded them from fig. 3.10. An even more significant change was adding a resistor in series to the main LC resonance, a resistor that served as the load-impedance for the Y-factor protocol. When cooled, the total electrical noise decreases, which

---

[8]Often, a load impedance can be transformed into any other, but it can be practically challenging to do the conversion without loss.
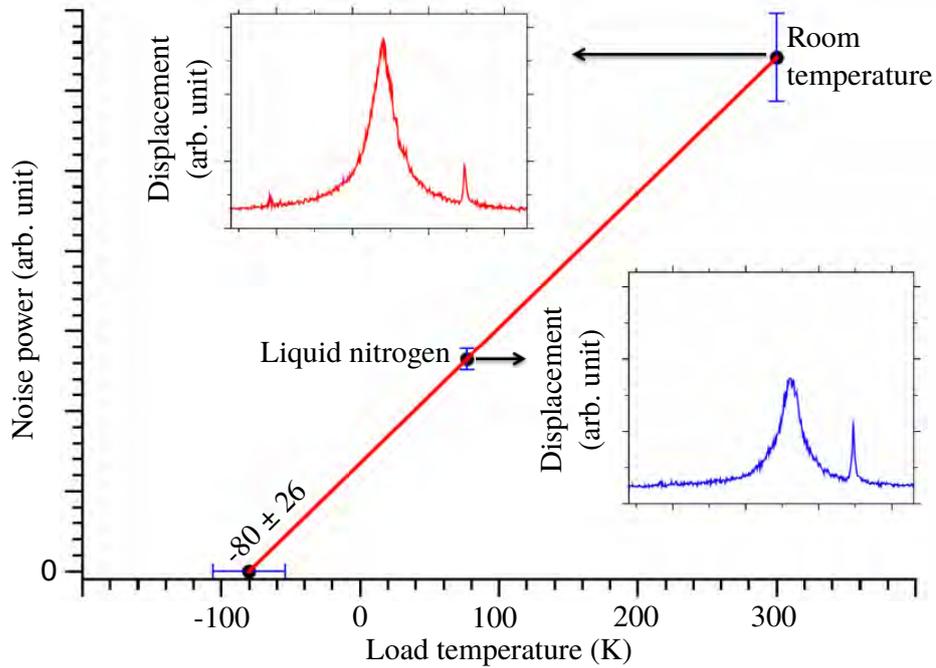
Figure 3.9: Noise-temperature extrapolation for the transducer based on the membrane's motion (the insets), measured optically. The y-axis corresponds to the noise-driven, RMS displacement, obtained by fitting a Lorentzian (eq. (2.42)) to the measured spectrum. (*Reprinted from Bagci (2014).*)
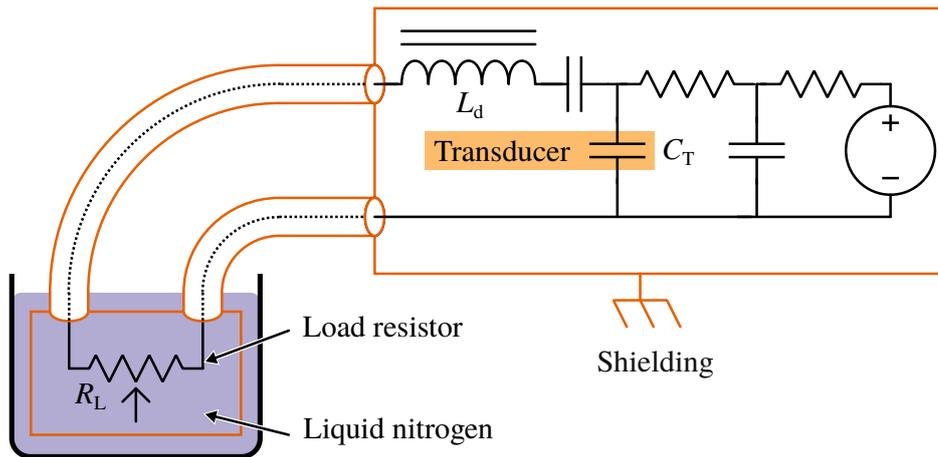
Figure 3.10: Schematic of the circuit used for the Y-factor calibration with the final shielding setup also illustrated.

mean less noise drives the membrane and its RMS displacement go down, as evidenced by fig. 3.9. The resistor and circuit connected through a coaxial cable, thus keeping the resistor some distance from the primary circuit and avoiding the thermal loading from cooling the inductor. To control its temperature, we either held the resistor at room temperature ($T_0 = 300\,$K) or dumped it into liquid nitrogen (77 K), as shown in fig. 3.10. The resistance could be varied by adding or bypassing individual resistors on a Printed Circuit Board (PCB) using jumpers.

Adding the load-resistor complicates the analysis in section 2.4 because the load changes the *LC* linewidth and, consequently, the cooperativity and mechanical broadening. When referenced to the load resistance R_L, the noise-temperature in eq. (2.70) on page 35 instead becomes (Bagci et al. 2014, Suppl. eq. (S37))

$$T_N \rightarrow \frac{R_d}{R_L} T_0 + \frac{R_d + R_L}{R_L} \left(\frac{T_m}{C} + \frac{(1+C)^2}{C} T_o \right). \tag{3.1}$$

3.2.2 Noise elimination

Initially, our experimental setup had tremendous excess noise in the circuit, which is problematic for the Y-factor extrapolation because such noise adds to the intrinsic noise from amplification. The excess noise should come as no surprise; the inductor acts as an antenna and picks up ambient electromagnetic noise, and there are many sources of RF noise in the environment at the frequency of interest (~1 MHz). This issue is well-known in radio technology. The noise-pollution stems from, for example, atmospheric effects (lightning), galactic, and human-made electronics like radio-stations and nearby electrical equipment.

Eliminating the excess noise was an iterative process of trial-and-error that continuously monitored the membrane's thermally-driven peak area for various setup configurations. In the limit of weak-electromechanical coupling, the thermally-driven mechanical area remains constant, provided both membrane and circuit are degenerate and only brings thermal noise to drive the hybrid electromechanical mode.[9] Excess noise causes the area to increase. In the end, the setup had these improvements implemented: enclosing the entire circuit in a metal shield, as shown in fig. 3.10 (that shield includes the vacuum chamber that contains the membrane-capacitor); replacing the inductor with a commercial product from Pico electronic that had a built-in shield; disconnecting all powered

---

[9]The equipartition theorem predicts the RMS motion to be eq. (2.41), a result that only depends on parameters that remain constant in the presence of electromechanical interaction.
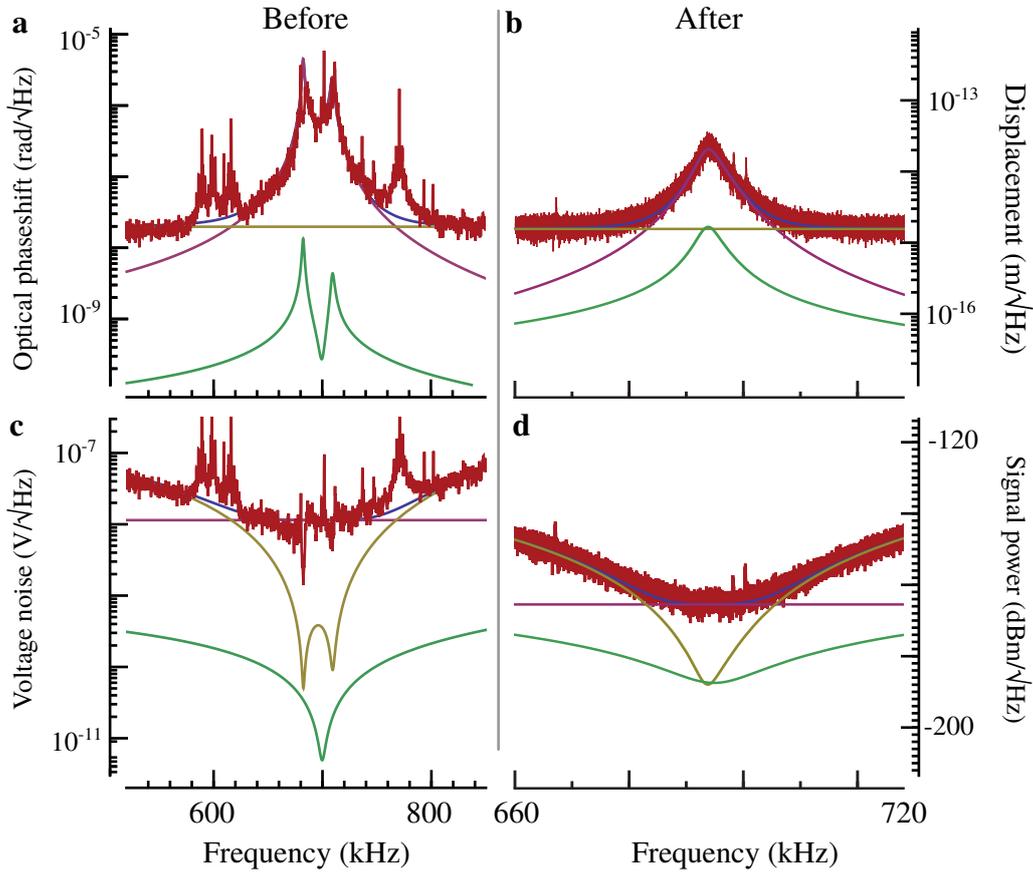
Figure 3.11: **a)** Early transduction with the setup at NBI, without a load resistor in the *LC* circuit. The mechanical motion measured with the interferometer (section 2.2.2), normalized to displacement units. **b)** The same measurement but now with the noise-reduction implemented. **c+d)** As (**a**) and (**b**) respectively, but converted into an equivalent voltage by dividing the displacement readout (eq. (2.15)) with the mechanical susceptibility (eq. (2.61a)). Model-curves come from eq. (2.59) and show the electrical noise (purple), optical noise (yellow), mechanical noise (green), and total noise (blue). Setup parameters for (**a+c**) were: $\Gamma_m = 20\,\text{Hz}$, $\Gamma_{LC} = 5.5\,\text{Hz}$, $\tilde{m} = 64\,\text{ng}$, $d_{\text{eff}} = {\sim}1\,\mu\text{m}$, and $V_{\text{DC}} = 21\,\text{V}$. For (**b+d**), the setup parameters were: $\Gamma_m = 20\,\text{Hz}$, $Q_{\text{LC}} = 47$, and $\tilde{m} = 64\,\text{ng}$. (*Adapted from Bagci et al. (2014)*.)

electronics from the *LC* tank; and meticulously grounding everything, in particular, grounding all coaxial cables[10] to the optical table. Sadly, the new inductor lowered the *LC Q* to 30 to 60 and, consequently, also the available electromechanical cooperativity, not to mention that all the parasitic capacitance added by the shielding. Further, the metal shield could not be a complete enclosure because the membrane-capacitor needed a window for the optical probing of the mechanical motion.

Despite our best efforts, some visibly noisy spectral regions persisted throughout all experiments, but these could mostly be avoided by tuning the detection frequency with the DC bias. With that, the membrane's RMS motion seemingly remained constant when turning on the electromechanical coupling, indicating that the excess noise got well below the actual thermal noise. Figure 3.11 showcase the transduced signal before and after the noise-reduction—without the load resistor—with the 'after' spectrum tuned to optimal cooperativity C_{opt}.

3.2.3   Result

Figure 3.12a shows the extrapolated noise-temperature in the shielded setup at NBI, measured with the interferometer, as a function of the load resistance with everything else kept constant. That meant the *LC* frequency had to be tweaked at each resistance value. Notably, the plotted curves are not a fit to the data; they come directly from the model (eq. (3.1)), with the electromechanical cooperativity evaluated from the mechanical broadening (eq. (2.63)).

We have seen cooperativity as high as 6800. The corresponding data is plotted in fig. 3.11a (noise-driven spectrum) and fig. 3.12b (strongly driven spectrum). That measurement was after the improvements to the membrane-capacitor assembly, requiring a modest DC bias of 21 V, but before all the necessary shielding. In the shielded measurements with the interferometer, we found an optical cooperativity 150 (fig. 3.11b), corresponding to a mechanical linewidth of 3 kHz.

---

[10]*All* the cables near the setup. Even the ones that did not connect directly to the circuit, but only connected to, for example, analysis equipment. Without grounding, additional noise somehow managed to creep in via those.
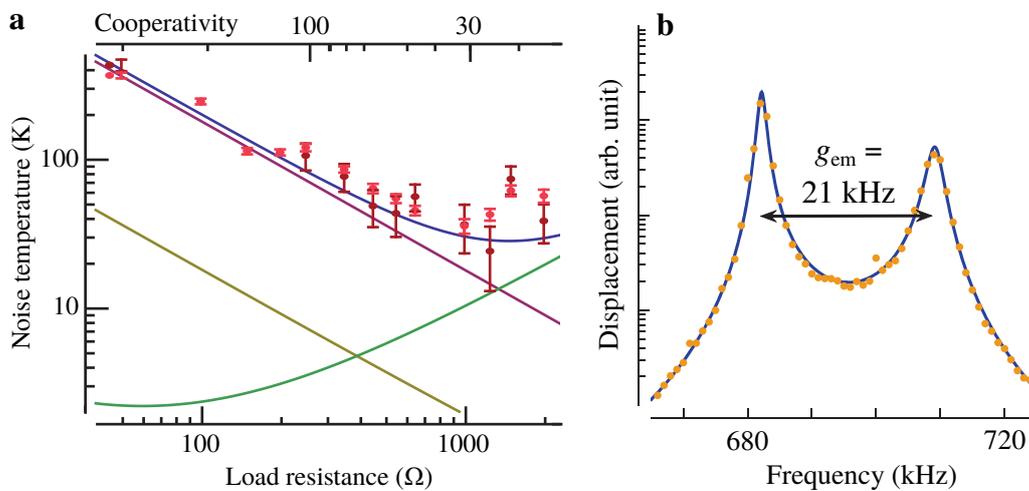
Figure 3.12: **a)** Noise-temperature, including intrinsic circuit noise, for different load resistances (and therefore different cooperativity). The data-points evaluate the Y-factor extrapolation with two different measures of output noise: the peak of the mechanical displacement (dark red) and the PSD integrated over a 10 kHz bandwidth around the resonance (light red). Models-curves come from eq. (2.66) and show the circuit noise (purple), optical noise (yellow), mechanical noise (green), and total noise (blue). Setup parameters were the same as fig. 3.11b. **b)** Driven measurement with our highest achieved electromechanical cooperativity. Setup parameters like in fig. 3.11a. (*Adapted from Bagci et al. (2014).*)

3.3 Discussion

After appropriate shielding of ambient RF noise, we saw a good agreement between the theoretically-predicted noise-temperatures (eq. (3.1)) and the values measured with the Y-factor methods, indicating that we reached the thermally limited noise level, and corroborating the agreement with experiments established by the driven measurements (section 3.1.2). Further, the improved membrane-capacitor gap significantly reduced the voltage required to achieve strong coupling.

According to eq. (2.70), the highest cooperativity corresponds to a transduction noise-temperature of $40\,\text{mK}$ ($5\,\text{pV}\,\sqrt{\text{Hz}}$) for the mechanics only, while the optimal cooperativity corresponds to $4\,\text{K}$ ($85\,\text{pV}\,\sqrt{\text{Hz}}$) including both mechanical and optical noise. For comparison, state-of-the-art low-noise amplifiers have noise-temperatures of about $20\,\text{K}$.[11] Cooling the amplifiers can bring the noise-temperature down to a few kelvins (McCulloch et al. 2017)—not considering amplifiers based on super-conductors since they only pertain to experiments in dilution refrigerators—but cooling the transducer will also reduce its noise-temperature. That means the transduction noise is very competitive with the state-of-the-art, plus it adds the benefit of up-conversion to light. However, the SNR bandwidth of transduction, $4\,\text{kHz}$, is very narrow compared to conventional electronic amplification.

---

[11]That is at least among the lowest numbers listed in specifications for commercial products, for example the Macom MAAL-011078.

Part II

Fabrication of transducers

Chapter 4

Process development

While the transduction scheme worked for the already presented membrane-capacitor device, that implementation left a lot to be desired. The main contentions were to achieve a smaller membrane-capacitor gap, because that boosts the coupling between membrane and circuit, and to eliminate issues associated with the flip-chip assembly. A more robust membrane-capacitor design is particularly advantageous for getting the transducer into an optical cavity, and cavities are required for more elaborate optomechanical experiments, not to mention they can both replace and improve upon the interferometric readout (elaborated in section 2.2). So solving these problems became the new goal of this project.

The first membrane-capacitor chips were made in a cleanroom, so it was natural to continue down that path for the new devices. Thus commenced my journey through all manner of cleanroom tests and process design, a journey that started in 2014 under the supervision of Silvan Schmid. Fabrication happened at Nanolab (then Danchip) which is a class 10-100 cleanroom at DTU. Nanolab functions both as a research- and industrial-facility for micro-fabrication. The available machines have many users and dedicated staff for maintenance and development of standard recipes.

Once our collaboration with DTU Hypermag began in 2016—with the ambition of trying the transduction in an MRI scanner—the goal-posts moved considerably. Specifically, a transportable setup became a high priority. Our new collaborators also desired a higher bandwidth, and that requirement opened up for new membrane materials, effectively resetting our hitherto work on a new device.

This chapter first elaborates on the device and fabrication criteria (section 4.1),

and potential ideas for the fabrication flow (section 4.1.2). Testing these ideas (section 4.2) culminated with the process-flow and sample-characterization described in section 4.3. The chapter excludes the optical cavity and fiber-coupling design, deferring these details to the next chapter. This split is purely conceptual; in practice, the cavity design and device fabrication go hand-in-hand. A coherent description of the process-flow, optics included, is given in Simonsen et al. (2019b).

4.1 Design considerations

Governing equations

The equations derived here aim to guide the membrane-capacitor design by expressing the transduction performance in terms of fabrication variables. The three vital transduction parameters are noise-temperature, bandwidth, and cooperativity. Section 2.4.2 proved that the optimal cooperativity, eq. (2.68) on page 35, sets the ultimate noise temperature and bandwidth for transduction, and the optimum depends on the membrane's mass, frequency, and linewidth. Controlling the linewidth is tricky[1], but both the mass and frequency scales with the membrane's area A_{mem}, height h, density ρ, and tensile stress[2] $\mathcal{T}$. Inserting eqs. (2.3) and (2.6) into eq. (2.70) gives:

$$T_{\text{N}} \approx 2\,\frac{\Omega_{LC}}{\Omega_{\text{m}}}\,\frac{T_{\text{m}}}{C_{\text{opt}}} \approx \Omega_{LC}\,\sqrt{\frac{2\,S_{xx}^{\text{light}}\,\tilde{m}\,\Gamma_{\text{m}}\,T_{\text{m}}}{k_{\text{B}}}} \propto \sqrt{A_{\text{mem}}\,h\,\rho\,\Gamma_{\text{m}}}, \qquad (4.1)$$

$$BW \approx \sqrt{2}\,\Gamma_{\text{m}}\,C_{\text{opt}} \approx 2\,\sqrt{\frac{k_{\text{B}}\,T_{\text{m}}\,\Gamma_{m}}{\tilde{m}\,\Omega_{\text{m}}^2}} \propto \sqrt{\frac{\Gamma_{\text{m}}}{h\,\mathcal{T}}} \qquad (4.2)$$

with $C = C_{\text{opt}}$. The last rewrites only emphasize the process variables. An ancillary question to answer is if the optimal cooperativity can be reached with a reasonable bias voltage. Hence, here is the cooperativity (eq. (2.68) with eqs. (2.3) and (2.6)):

$$C \propto \frac{A_{\text{mem}}\,L_\nu}{d_{\text{eff}}^4\,h\,\Gamma_{\text{m}}\,\sqrt{\mathcal{T}\,\rho}} \qquad (4.3)$$

expressed in terms of the process variables, where L_ν is the side-length (radius) for the square (circular) membrane.

---

[1]It depends on material losses and capacitor design (Yu et al. 2012), and also a coupling to phonon-modes in the substrate (Jöckel et al. 2011).

[2]Although we have only limited control over the tensile stress.

Design goals

To begin with, we aimed to achieve the lowest noise-temperature possible in an AC setup with little regard for any other figures-of-merit. However, once the focus shifted to implementation in an MRI scanner, so did the target requirements. In particular, the Hypermag group desired a detection bandwidth of 10 kHz or higher. They also set an upper limit on the AC bias (2 V) and preferred it as low as possible, essentially requiring a small distance between membrane and electrode. The MRI circuit design additionally demands the lowest possible ohmic loss, both from the membrane-capacitor and the bias. The latter necessitates decent filtering between bias and *LC* resonance and, consequently, a sufficiently high mechanical frequency—more in section 6.1.1 on page 105. We translated this constraint to a membrane frequency above 1 MHz for a realistic circuit filter.

The equations above show that sensitivity and bandwidth get better by reducing the membrane's area, density, thickness, and tensile stress. A high mechanical Q further helps the sensitivity but it is detrimental to the bandwidth. Additionally, the membrane-capacitor gap should be small to reach high electromechanical coupling. A bigger area also increases cooperativity, everything else kept constant, but that sacrifice noise-temperature. Other parameters also matter, but to a much lesser extent. The membrane, of course, has to have some tensile stress to have a resonance and meet the 1 MHz target; in fact, the eqs. (4.1) to (4.3) uses eq. (2.1), thus implicitly assuming high tensile stress.

Since the transduction eventually had to into an MRI scanner, it became a must to realize a robust and portable optical setup. That favored direct fiber-coupling between laser and membrane-capacitor and using an optical cavity instead of the interferometer. Even a poor cavity suffices to detect the mechanical motion, which made the membrane-metal a viable cavity-mirror despite its low reflectivity. We decided to define the cavity through the fabrication as well, as that would deliver a simple end-product. That solution demands an accurate placement of the membrane, with the accuracy determined by the linewidth of the cavity.

Up to this point, we had used a laser operating at $\lambda = 1064$ nm to characterize samples, as that was available in our lab at NBI. Therefore, the first transducer designs also targeted optical operation at this wavelength. This criterion was set out of convenience rather than from necessity. Unfortunately, the choice made it cumbersome to switch wavelength later on although that might have been beneficial. Switching would have entailed getting a new laser, detector, and other

optical components, not to mention changing the established fabrication parameters.

Scope

While the next section describes the fabrication processes we studied, it is equally interesting to disclose which procedures we deliberately did not consider. Concretely, we limited investigations to standard Ultra-Violet (UV)-lithography. While Nanolab does offer high-resolution lithography, i.e., deep UV or e-beam lithography, the dedicated machines are costly and time-consuming, and plenty of traffic from other users makes turn-around time between tests an issue. Regular UV-lithography is anyway a well-established, flexible process that can make many devices at once. If it is viable, it seems sensible to favor it.

Additionally, we have omitted vacuum-packaging for the first integrated membrane-capacitor devices. Getting the membrane-capacitor into vacuum is a technical challenge that must be solved in the future, but we deemed it too cumbersome for the first integrated transducers.

As the last scope, we prioritized processes available at Nanolab simply because they were the most convenient for us. While not a strict scope per se, it has had considerable ramifications for the fabrication developments.

4.1.1   Material choices

Initially, we exclusively considered nitride membranes made by Low-Pressure Chemical-Vapor Deposition (LPCVD). They can reach the highest $Q \times \Omega_{\mathrm{m}}$ product for a mechanical oscillator at room temperature (Tsaturyan et al. 2017), and a high mechanical Q-factor leads to a low noise-temperature according to eq. (4.1). For the same reason, nitride membranes were used in all other works on room-temperature transduction (Bagci et al. 2014; Haghighi et al. 2018; Takeda et al. 2017; Tominaga et al. 2018). However, since high Q-factor is detrimental to the transduction bandwidth (eq. (4.2)), we started searching for alternative materials.

While many conventional cleanroom materials might function as a membrane for the transducer, pure aluminum stands out for several reasons: it is conductive, lightweight, and have a good reflectivity at λ for a metal. Moreover, Yu et al. (2012) showed that annealing aluminum layer induces tensile stress in it and that resultant intrinsic Q was decent.[3] After testing and verifying that annealing

---

[3]Another conclusion was that metal does not degrade the Q-factor of nitride appreciably, as

induced tensile stress—specifically, one hour at $400\,°C$ yielded $350\,MPa$[4] in a thin aluminum layer deposited with Electron-Beam Physical Vapor Deposition (EBPVD)—we decided to attempt pure aluminum membrane.

Other materials considered included silicon and amorphous Al_2O_3—henceforth called alumina. Silicon is a well-established material in all cleanrooms and can be deposited with tensile stress, act as a conductor, and have a density below aluminum. However, silicon is a semiconductor, so its conductivity is no match for aluminum, and its optical properties are not ideal for $1064\,nm$. Silicon basically requires an additional metallization. Alumina emerged as an option because it turned out to be compatible with the same processes as aluminum, and Atomic-Layer Deposition (ALD) deposition creates tensile stress, specifically $350\,MPa$ when using the standard Nanolab recipe at $200\,°C$, a result that agrees with Yli-vaara et al. 2014 for similar deposition parameters.

The substrate material determines the scope of potential fabrication processes, just like the membrane material, so this choice also needs some attention. In the end, we decided to make a process work with a glass substrate rather than, for example, silicon. Using glass meant that light at the operating wavelength, $1064\,nm$, could pass through the substrate, making the fiber-coupling easier to implement. A glass wafer also isolates electrical signals very well, thus boosting the electrical performance of the final transducer compared to the same device on silicon.

4.1.2   Potential processes

This section explains our different ideas for fabrication the next-generation transducers that meet the design objectives, in the most generic terms possible, to present the thought process underlying each scheme and introduce the primary process step and technical challenges. All these ideas have one thing in common: they all must lead to a membrane somehow, and there are not many ways to fabricate those. The usual way is to deposit the membrane on a carrier material and then remove the carrier in a region thus exposing the membrane-layer. Alternatively, the membrane can be deposited on one wafer and transferred to

---

long as the metal does not cover the edge of the membrane. Therefore, we initially tried to meet this design criterion in the transducer fabrication.

[4]The stress-measurements used Stoney's formula, evaluated for the wafer-curvature measured with and without the layer in question. The curvature comes from a profilometer that sweeps a stylus over the whole wafer. We always tried to scan along the same line in both measurements, but that was difficult in our setup.

another carrier that has had regions removed before the transfer. That effectively leaves a membrane exposed in an area while holding it in place at the boundary where carrier and membrane meet. Furthermore, the membrane has to be near an electrode while being able to move freely, and that electrode must be supported somehow.

Bonding

This process idea is to make the chips separately, just like the old transducer design, add a spacer layer to one of the chips, and bond them together as a part of the cleanroom fabrication. The spacer should not be between the membrane and electrode, but instead surround them, thus creating a cavity. The membrane-electrode distance would then be the thickness of the spacer layer. It is essentially the same idea as section 3.1.1, but with a better way of joining the two chips—the wafer bonding—being the new processing step tested below (section 4.2.1).

The process seems straightforward, especially considering that making the membrane and electrode chips separately was an already established process. However, full fabrication procedure has some ambiguity in its implementation; most importantly, releasing the membranes can happen either before or after the bonding. We decided to attempt the process with membrane-release after bonding because it seemed perilous to bring released membranes near a structured wafer, and because bonding works better with less preprocessing of the wafers as discussed in section 4.2.1. Unfortunately, the process did not work even though the idea followed a hitherto established process closely. The reason was an unforeseen consequence of the chosen type of wafer bonding, and there were technical difficulties with the equipment that discouraged trying the same process with other types of bonding. Later on, we reconsidered bonding as a way to transfer layers between wafers.

Under-etching

This process idea is to deposit a stack of three layers and remove the middle one, thus freeing a membrane. The sacrificial layer must be removed by an isotropic etch. The membrane can either be the first or last deposited layer, and the electrode then has to be the other layer. That means the sacrificial determines the membrane-capacitor gap. The essential processing step is the isotropic etchant; it must remove the sacrificial layer without attacking either membrane or electrode,

and the materials most obey restrictions that the cleanroom machines have.

The tests below fall into two categories: wet- and dry-etching. Wet-etching demands special attention to drying. The wet-etching also has the membrane-layer deposited first, and the electrode on top. For dry-etching, it is the other way around. This configuration is not a requirement but just happened to be the case for processes we have considered. Processes compatible with wet-etching often ended up with both the top and bottom layer being unsupported around the membrane. That makes vibrations in the electrode an additional complication of these designs.

Note that the sacrificial layer must be exposed to the etchant somehow. Inspired by Southworth et al. (2009), we opted to make holes in the membrane through which the etchant enters, and the reaction products leave. The etch happens vertically and undercut the membrane in concentric circles around the holes. Importantly, these holes have to be far enough apart to leave a smooth surface that can reflect the probing laser light[5]. A reasonable laser spot would need a hole-spacing of $30\,\mu m$ or more, plus the diameter of the holes themselves. Since the membrane simultaneously has to be thinner than $100\,nm$, the etchant must remove the sacrificial layer much more than 150 times faster than it etches the membrane. If the etchant contacts the metal on the membrane or electrode, the same statement applies. In principle, any etch of the membrane or metal could be accounted for through design, but it would jeopardize the membrane-capacitor gap.

Layer transfer

This process idea is a variation of under-etching where the triple-layer stack gets transferred from the first wafer and onto another. By doing so, the stack can be made first, and then processed from both sides afterward. That circumvents many of the material restrictions that otherwise challenged idea of under-etching, specifically those concerning wet-etching and the choice of metallization.

4.2 Testing process-steps

This section summarizes our tests of the fabrication ideas described above. The work was not carried out in an orderly fashion; instead, it ended up following

---

[5]The holes could also form a sub-wavelength reflective grating (Kemiktarak et al. 2012), but that necessitates high-resolution lithography which is outside the scope of our fabrication.

a trial-and-error approach were each idea offered potential solutions to some problems, and each test introduced more. These problems related not only to the physics and chemistry behind the process, but also to technical issues, availability of machinery, and turnaround time. It is plausible that all the processes ideas could work, provided sufficient optimization, but the technical challenges often limit the feasibility of such optimization, thus becoming the deciding factor in process design.

All the test below employed a whole host of standard fabrication techniques—deposition, etching, cleaning, and UV-lithography—mostly using standard recipes developed by Nanolab staff. Details on these standards will be kept to a minimum as they tend to be very procedural and lengthy without adding too much insight. Instead, this treatise aims to describe only the processes without an established standard recipe and that required development.

4.2.1 Wafer-bonding

There exist several ways to bond two wafers together. The main ones available in Nanolab are direct, anodic, eutectic, and adhesive bonding. These are also some of the more common bonding techniques, although other types do exist. The sections below only present testing outcomes of direct and anodic bonding. We have not tested adhesion bonding more than once[6], and have not tested eutectic bonding at all[7]. Figure 4.1 illustrates the three tested bonding types.

Wafer bonding has two key characteristics: the strength of the bond and how well the bond covers the wafer. The strength primarily derives from the bond type but may still be dependent on the process parameters. To test the strength, we forced a scalpel between the wafers and tried to pry them apart. If the wafers break rather than coming apart, then the bond is good. Both anodic and direct bonding yielded good bonds with proper process parameters. To evaluate the coverage of the bond, we merely looked at the silicon through the glass wafer.

---

[6]The established fabrication flow—spinning a polymer on a wafer and using it as adhesive—offered only poor control over the polymer thickness, and only thick layers (>1 μm), in addition to being a cumbersome and dirty setup.

[7]Eutectic bonding was not tested because, at first, we wanted to move away from wafer bonding entirely, and when we started to reconsider bonding, we wanted to move away from silicon substrates, and eutectic bonding typically is between a silicon wafer and a gold layer. In hindsight, it might have been a better choice than anodic bonding, at least to begin with, as it should (theoretically) circumvent the issues we observed with anodic bonding and be more forgiving than direct bonding.
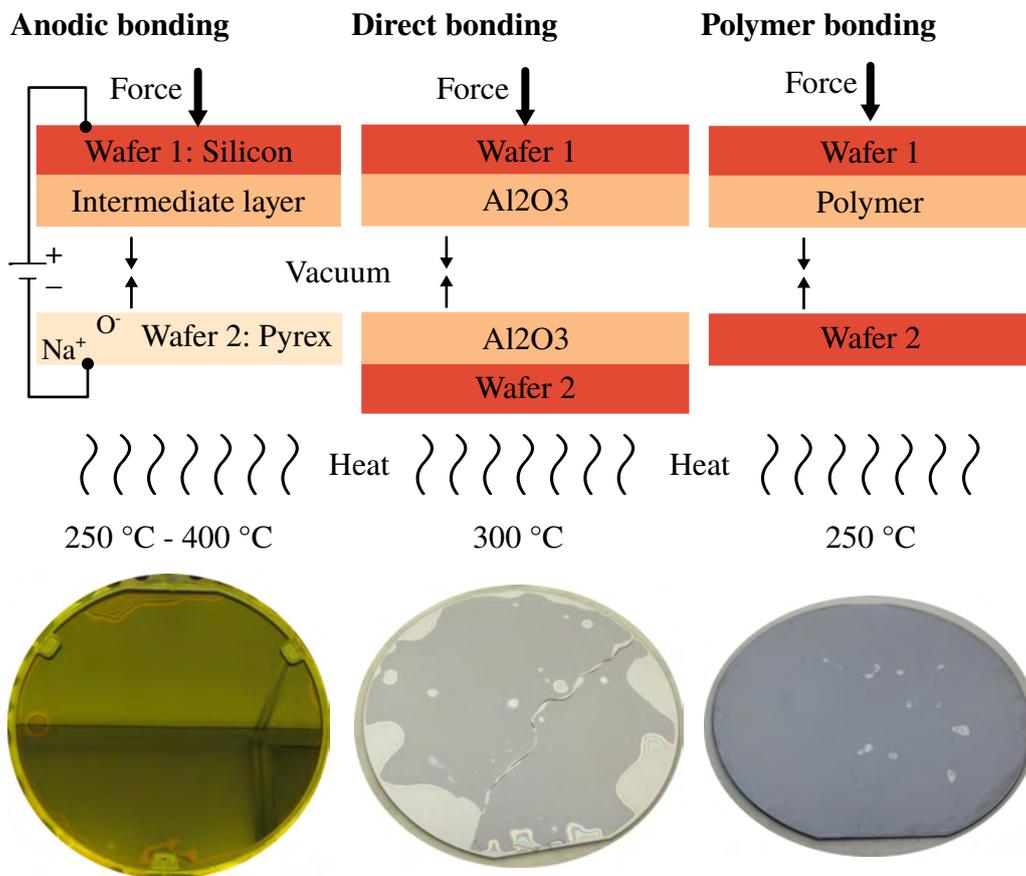
Figure 4.1: Schematic cross-sections of tested bonding types with photos after bonding between a clean silicon and glass wafer.

That is of course only possible when one wafer is transparent, but luckily that was the case for most of the bonding tests described below. If there is a region where the wafers have not made contact, the interface usually has a different color. Furthermore, the wafers often have different curvatures in such a region, which leads to Newton rings, i.e., an interference pattern in the light that passes through the (curved) glass and reflects of the (curved) silicon. Figure 4.1 shows the Newton rings after bonding plain wafers without any structure or processing.

Direct bonding

As illustrated in fig. 4.1, direct bonding (or fusion bonding) takes place between two identical and atomically clean surfaces. They are usually pressed together at low temperature (pre-bonded) and then annealed at a higher temperature. The re-

sulting bond is strong. Directly bonding silicon or fused silica wafers requires very high temperatures(Rushton et al. 2014), thus limiting material added in prior processing. Instead, we tested direct bonding between two different wafers (silicon and glass) covered in alumina, where the bond forms between the intermediate alumina layers (Sahoo et al. 2018). This technique has the advantages that it can bond wafers made from different materials, and it does not require as high a temperature as fusion bonding. However, the surfaces still have to be very clean. That limits the tolerable processing before bonding significantly, or it demands an aggressive cleaning procedure that inevitably also attacks metal on the wafers

The photo in fig. 4.1 display the two main challenges with the direct, alumina-bonding approach: the bond did not cover the entire wafers, indicated by regions of newton rings, even when the wafers came clean and unprocessed out of the box; and the glass wafer would crack after running the bonding process. The cracking was an unexpected and inconsistent quandary, and we never found a definitive explanation. Maybe it relates the mismatch in thermal expansion coefficients between the wafers. Interestingly, the best bond result—without cracks— came with the piston in the bonding machine being the same size as the wafer (this is not the default setting). Nevertheless, this process could probably lead to integrated transducers with sufficient wafer cleaning and process optimization.

Anodic bonding

Usually, anodic bonding occurs between silicon and a glass containing NaO_2. An applied voltage splits the NaO_2, pulling the sodium ions away, and pushing the oxygen towards, the bonding interface (Rushton et al. 2014). The drift of ions causes a current to run between the two wafers, and it creates a powerful, attractive force that pulls the wafers into intimate contact. Without the force, small impurities would prevent the wafers from making contact just like direct bonding. As it is, anodic bonding is more forgiving in terms of cleanliness of wafer surfaces compared to direct bonding.

Throughout initial attempts at this process, it became apparent that the wafer-bonder machine at Nanolab was problematic[8], and that the actual results of the bonding were somewhat inconsistent. Therefore, we constructed a custom setup

---

[8]At some point during the process testing, the technical staff learned that the piston in the bonder had broken. No cause was ever determined, nor was the exact time of the accident, so we cannot say which tests it affected or if it happened during our processing.
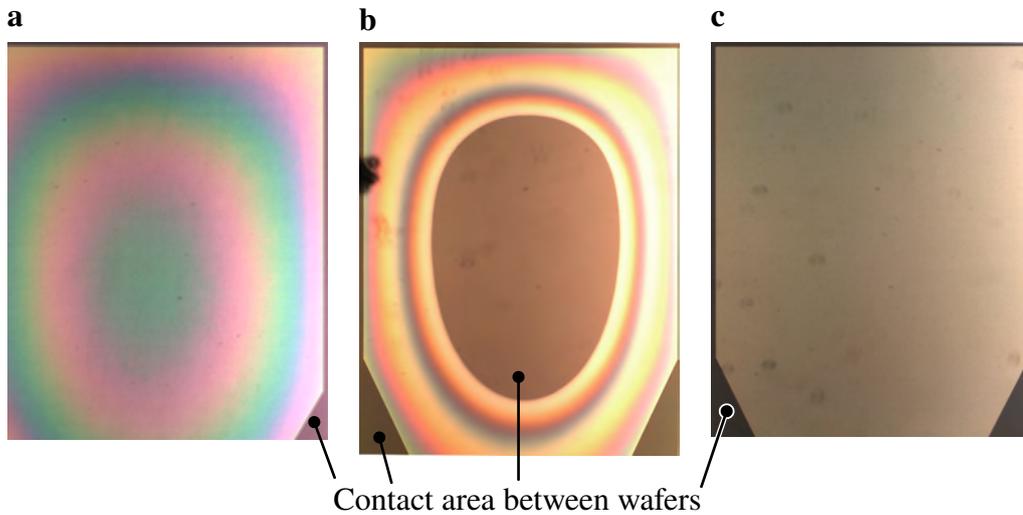
Figure 4.2: Cavities made with anodic bonding, for different bonding parameters, with the Newton rings indicating a curvature of one wafer relative to the other.

outside the cleanroom (at NBI) to do the bonding instead. These tests revealed some hurdles with the bonding process itself, mainly that the cavities between wafers mostly collapsed for bonding parameters that created strong bonds. The problem is the electrostatic force that makes anodic bonding a reliable process. That force also pulls the wafers into contact when the wafer-distance gets too small. Figure 4.2 demonstrates this effect for a cavity defined by a nitride spacer made with Plasma-Enhanced Chemical Vapor Deposition (PECVD). The figure also shows the collapse did not always happen. The issue has been investigated more thoroughly by Mao and Han (2005) in the context of microchannels, and they estimated the gap between wafers had to be >0.004 times the channel width to prevent collapse. In our case, that would be >1 μm That discourages the hole process idea. Even when the cavity resisted the collapse, the glass and silicon wafer are not necessarily parallel, as indicated by the Newton rings in fig. 4.2, which jeopardizes the whole idea of getting a well-defined distance between the membrane and electrode.

4.2.2   Wet etching

It proved challenging to find etchants to do isotropic under-etching while also meeting the objectives and materials choices defined above. For instance, Williams et al. (2003) is a relatively comprehensive source of etch-rates for standard materials and etchants, but they offered no viable etch that neither attacks membrane

or aluminum and have a decent etch-rate in a suitable sacrificial layer.[9] However, for a pure aluminum membrane, Williams et al. (2003) indicate that one particular chromium etch could work: CR-14. Since this etchant was no longer available, we mixed it ourself to try it out. While the mixture did etch chromium without an aluminum layer present, no etching occurred at all if the chromium had aluminum on top; we have not investigated this failure in more detail.

Before shifting attention to aluminum membranes, we did try with gold instead of aluminum because it is much more resistant to chemicals, and it still offers a high reflectivity at our target wavelength. However, its density is a noteworthy downside. With gold as a metal, Williams et al. (2003) predict that KOH can work for a nitride membrane using the process described in the below. By the time we abandoned it, this process had worked to an extent, but it required more optimization regarding the drying. Since then, Takeda et al. (2017) have used a gold-coated nitride membrane to do transduction, using a flip-chip assembly design similar to chapter 3, and they concluded that the mass of gold restricted their performance. In their follow-up work (Tominaga et al. 2018), they switched to aluminum and got drastic improvements.

Critical-point drying

All the isotropic etchants mention so far are liquids, making drying the samples after processing a crucial challenge. Without precautions, the evaporating liquid has surface-tension that collapses the membrane onto the electrode. We observed this effect even if the final liquid had very low surface tension, e.g. ethanol. Luckily, a solution was available in Nanolab: namely, critical-point drying. It worked as follows: we transferred the membranes to a process-chamber filled with isopropyl alcohol after etch and rinse. Isopropyl alcohol is miscible with liquid CO_2, the latter being the actual liquid used for drying. The critical-point dryer machine then fills with liquid CO_2 until all the alcohol has been replaced, and proceeds to heat the chamber and increase the pressure until CO_2 becomes a supercritical fluid. Then it releases the pressure and cools to ambient values but without crossing the liquid-to-gas phase-transition, thus avoiding condensation altogether. The wafer with membranes then comes out dry after a few hours of processing.

Even though drying wafer with the critical-point drying ended up being a reliable process, it started with a few technical challenges. First, the settings of the

---

[9] H_2O_2 heated to $50\,°\,°C$ supposedly might work, but it thickens the aluminum layer.

machine needed significant adjustment from the technical staff at Nanolab. Before that, our few test-runs ended with the wafer still wet. Second, the process left rather large, dust-like particles on the wafer[10], which were circumvented by placing the wafers with the membranes facing down. Lastly, it is precarious to transfer wafers with the fragile membranes from one liquid to another—the key is to do it fast.

Process flow

Figure 4.3 shows the cross-section of a design we made using isotropic etching with KOH to release the transducer's membrane. The process started by deposition and patterning the depicted stack, one layer at a time, on the front of the wafer. Then followed an anisotropic etch from the backside, through the silicon substrate, to release the stack. The etchant was KOH which follows the $\langle 111 \rangle$ crystal-plane of silicon. To avoid etching the front-side during the long etch through the entire substrate, we placed the wafer in a dedicated holder that protect the front from chemicals—when mounted correctly[11], that is. After etching the substrate and removing the wafer from the holder, the entire wafer went into

---

[10]Nanolab technical staff later traced this to particles from tissues used for cleaning.

[11]What 'currently' means, in this context, was never really established. Sometimes, what looked like proper mounting would nevertheless not protect the wafer sufficiently. To make matter worse, it was not possible to tell before the substrate had etched through completely, which takes several hours, making optimization of this process a tedious endeavor.
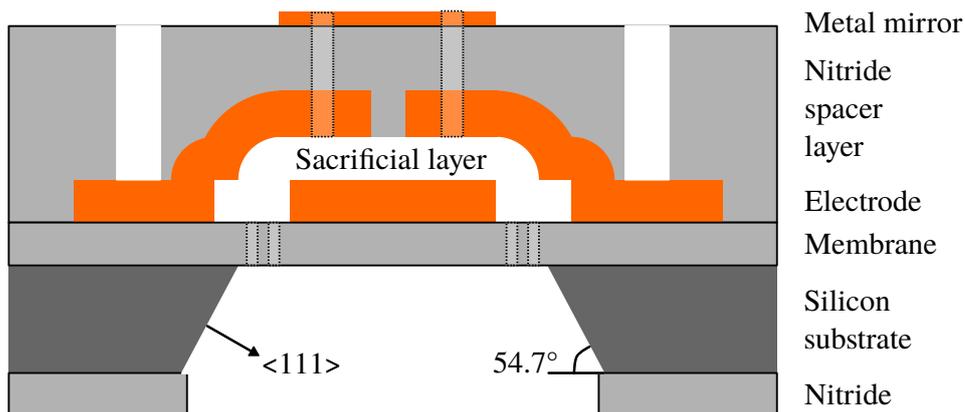
Figure 4.3: Transducer-design based on membrane-release through isotropic etching of a sacrificial layer. All metals are gold, and the top metal forms an optical cavity together with the metal on the membrane.

KOH to remove the sacrificial layer—sputtered, amorphous silicon. Afterward, the wafer remained submerged in liquid right up until the critical-point drying.

Before abandoning the process, we did observe some problems with it. For instance, the electrode layer also forms a membrane with resonances, complicating the electromechanical interaction between membrane and capacitor. Moreover, the metal electrode needed a supporting layer (PECVD, low-stress nitride), and KOH would penetrate the support and create pinholes during the final membrane release. Furthermore, with silicon as the substrate, the electrical properties were rather bad because silicon itself is a (poor) conductor, allowing current to leak between the electrodes of the planar capacitor. This issue might be solvable with high-resistivity silicon and dielectric layers between the electrode and the silicon itself, but for our first implementation, it was a problem.

4.2.3 Layer transfer

Figure 4.4 presents the aim of the layer-transfer idea, together with the outcome of a limited set of test. These test started with a sacrificial wafer (standard silicon) processed to make the nitride layer, a sacrificial layer (amorphous silicon), and aluminum metallization. Then we bonded that wafer to a glass-substrate and removed the silicon using KOH etching, thus leaving only the layer-stack behind but now exposing the backside of the stack for continued processing. As shown in fig. 4.4a, the new glass carrier had predefined cavities that eventually define the membrane and can contain protruding metallization.

Importantly, we wanted to use anodic bonding because, out of all the bonding techniques in section 4.2.1, it is the most insensitive to residue on the wafers. Such residue comes, for example, from all prior processing—in particular, the lithography required to pattern the aluminum metal—and it cannot be removed with standard cleaning like Piranha[12] because that attacks aluminum as well.

This layer-transfer technique gave rise to some technical problems shown in fig. 4.4. Figure 4.4b shows the most critical problem: that the membrane bulges outward after removing the sacrificial wafer. One explanation is that anodic bonding causes outgassing from the interface between the two wafers (Rushton et al. 2014). A possible solution would be venting channels running across the entire wafer to release the pressure, but they would have to be closed off after the transfer, preventing fluids from entering in subsequent processing. Alternatively, one could ignore the problem entirely and hope those bulging membranes

---

[12]4:1 H_2SO_4:H_2O_2
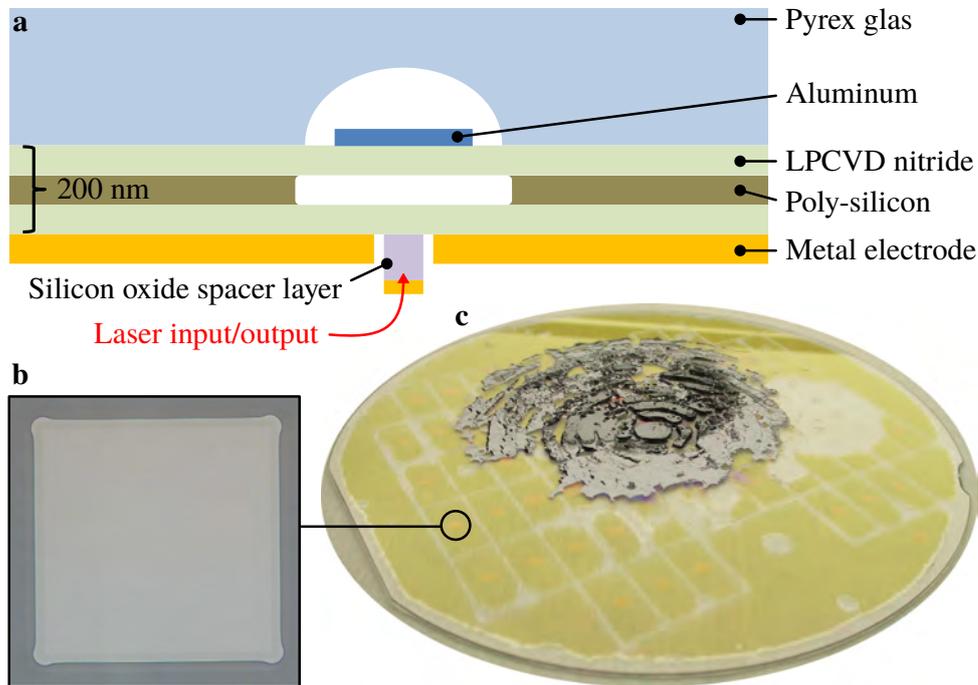
Figure 4.4: **a)** Target design for the layer-transfer process. **b)** Microscope view of a membrane transferred this way. Notice the faint Newton-rings near the edge of the membrane. Those rings indicate a bulging membrane. **c)** View of a full wafer from an initial layer-transfer test, after bonding and failed etching of the sacrificial wafer (silicon).

can survive subsequent lithography and metal deposition, and that the bulging does not degraded alignment appreciably. Two more issues were that the silicon did not always etch through everywhere and that the layer-stack did not adhere equally well over the glass wafer (fig. 4.4c). The first issue seemed to stem from a residue appearing after bonding, a residue that seemingly protected the silicon in the KOH etch. The second issue could be alleviated by adding many small holes in the glass substrate, presumably to collect the outgassing or trapped air between the wafers, not to mention catching the occasional particle contamination.

Another prohibitive technical challenge turned out to be making the cavities in the glass wafer. The predicament is that those holes have to be deep—at least a few micrometers—to not collapse during bonding, and that severely limits the choice of masking and substrate material. Conventional lithography-resist would not survive a dry-etch in glass going that deep. Instead, deep glass-etches are

typically made with strong HF acid because that offers high etch rate and decent selectivity to resist. However, HF acid also causes the resist to delaminate from the substrate, thus expanding the lithographic pattern and making it challenging to hit the target dimension for the membrane. We even saw delamination with a metal (Au/Cr) mask. A better masking material would be LPCVD silicon or nitride (Zhu et al. 2009), but these are incompatible with the pyrex glass[13] demanded by the anodic bonding.

As a different approach to make cavities, we have also tried machining holes in a glass with laser drilling, but to no avail. Those holes ended up somewhat asymmetric and surrounded by debris, debris that make the surface very uneven near the edge of the hole. Improving the hole-shape would require considerable optimization, and the debris necessitates improved masking or wafer-polishing after processing.

4.2.4 Dry-etch

Isotropic under-etching is also possible to do with a dry etch process, which avoids the critical-point drying altogether. The etch rates in Williams et al. (2003) suggest some feasible processes, but none using nitride as the membrane material. Instead, aluminum and alumina are viable membrane materials when using either an SF_6-based plasma or an HF-vapor etching process. Neither etches aluminum if implemented correctly, but they can etch nitride or oxide respectively at a high rate, meaning these materials would work as a sacrificial layer. Tests with these techniques are detailed in the next two subsections.

HF vapor etching

Nanolab does not have a dedicated machine for HF vapor etching. Instead, we tested a setup with the wafer placed in proximity to an aqueous solution with a high concentration of HF, but not submerged, as illustrated in fig. 4.5a. The vapor right above this liquid contains enough HF to etch oxide. Unfortunately, this process etched aluminum, as shown in fig. 4.5b, because the ambient vapor contains water, whereas commercial HF-vapor etchers are anhydrous. Instead, we tried out the water-free HF vapor etch through two external facilities: at VTT

---

[13]At least at Nanolab, due to cross-contamination concerns. In principle, LPCVD silicon could be deposited on pyrex. However, LPCVD nitride uses a temperature above the melting point of pyrex.
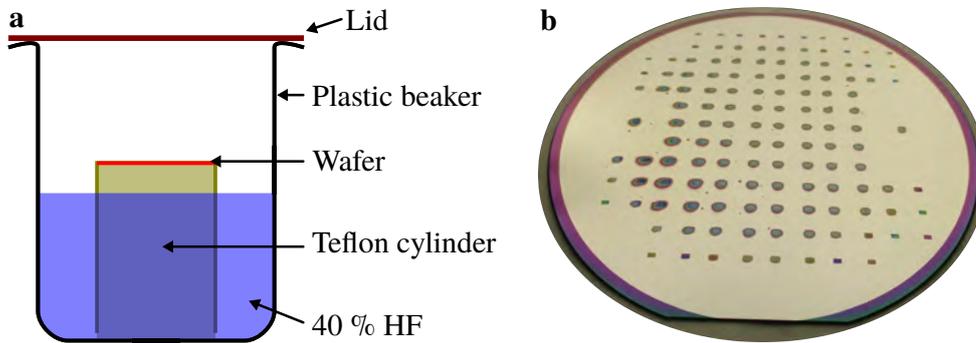
Figure 4.5: **a)** Schematic of an HF vapor etch setup that works at ambient pressure and temperature. **b)** An attempt at releasing aluminum membranes with this process where all membranes have vanished due to the etching.

(Aalto University) with the samples prepared in our lab, and at Oscar Painter's group (Caltech) with the aluminum and oxide deposited there.

Figure 4.6 (inset) shows the result of the HF vapor-etch from Finland. Notice how the aluminum remains, but it has started to curve weirdly upwards, creating a volcano-like surface around each small hole in the aluminum. This issue also occurred with the plasma-etch described next when the membranes were pure aluminum, but not for alumina membranes. Presumably, this would also be true for the HF vapor etcher, but we have not repeated this test with just an alumina membrane. Interestingly, the process run at Caltech did not reproduce this curving, but the membrane collapsed onto the electrode even though it had been annealed to induce tensile stress. Supposedly, the difference in outcome originates from the difference in layer deposition[14], but the tests have been exhaustive, making it hard to draw any decisive conclusions. The most appropriate deduction is, perhaps, that it can be challenging to move fabrication processes between facilities.

Plasma etching

In parallel with the HF vapor test, we tested an isotropic plasma-etch based on SF_6 gas, using parameters inspired by Cicak et al. (2009). That gas etches fast both silicon and PECVD-deposited nitride. Optimizing the process was rather straightforward; the highest pressure and power yielded the highest etch rate in

---

[14]The aluminum deposition at Caltech was optimized for low-temperature superconductivity and therefore met some high standards concerning, for example, purity and vacuum level Our in-house deposition is certainly less optimized in that regard.
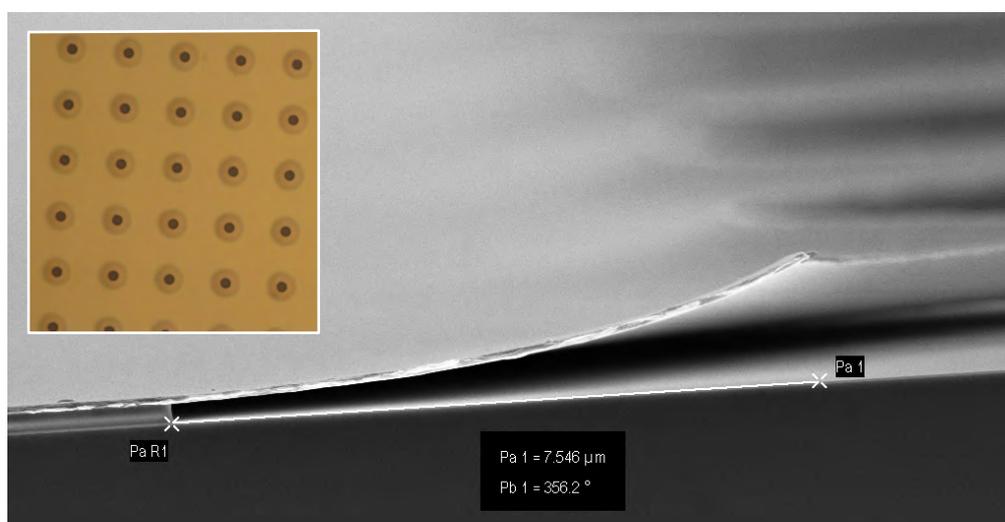
Figure 4.6: Scanning Electron Microscope (SEM) image showing the cross-section of an aluminum layer partially released with this SF_6-based plasma process. Note that the hole curves upward, also in the background. The inset shows a microscope top-view of samples processed with HF vapor etching at VTT. The color-gradient around the holes also comes from a curvature of the pure aluminum layer similar to the plasma etching.

the horizontal direction, and platen power was off to prevent ion-bombardments and resultant physical sputtering of the membrane layers. Based on that, we developed an etch-recipe with pressure and power near their maximum values, and we timed the etch to ensure that the membrane has fully released. That recipe became the work-horse for all membrane-fabrication in the remaining thesis because of its advantages: available in-house, only ~30 min etch necessary, perfect selectivity to the sacrificial layer[15], compatible with aluminum/alumina membranes and standard cleanroom materials. As a sacrificial layer, we went with PECVD nitride rather than silicon because nitride is an insulator. However, that lowers the isotropic etch-rate a little bit.

However, the process does have some technical issues. As mentioned above, the membrane curved if the layer was pure aluminum. Therefore, we have added an alumina layer under-neath the aluminum, thus giving mechanical stability to the released membrane. That solved the problem, so we have not investigated it any further.

Another issue was that the membranes collapsed if they were too large, as

---

[15]We have at least never observed any etch of aluminum or alumina for the etch required to release the membranes.
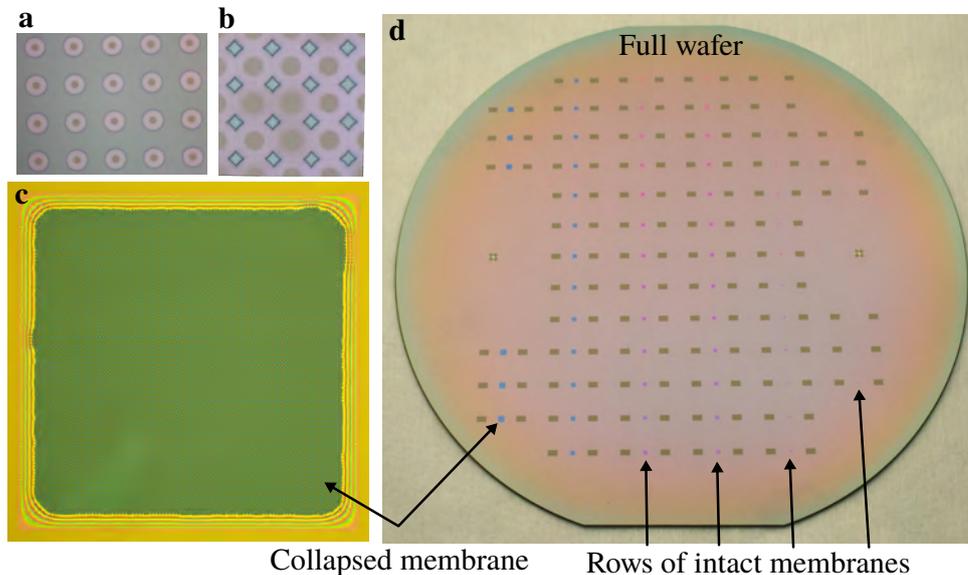
Figure 4.7: **a+b)** Partial under-etch of an alumina membrane for different hole-geometries. **c)** Fully released, and collapsed membrane. The blue color shows were the alumina membrane makes contact with the substrate. **d)** A whole wafer with the biggest membranes all collapsed, as seen by the blue color, and rows of smaller membranes still intact.

depicted in fig. 4.7 for a pure alumina membrane. The collapse brings a visible change in color, and the edges of the membrane show Newton rings as in fig. 4.2. This problem occurred when transferring the wafer out of the plasma-etcher, where it is under vacuum, and then bringing them back to ambient pressure. In fact, it happened when venting the Inductively-Coupled Plasma (ICP) load-lock. The cause seemed to be that the gap between membrane and electrode increased in pressure much slower than the immediate surrounding, which makes sense considering that any gas can only enter through the small holes used for under-etching. Venting slower could probably alleviate this issue. However, if the membranes were small enough, and the membrane-electrode distance large enough, the venting was not an issue at all, and the necessary membrane-diameters were indeed sufficiently small for the target distance. Note that the smallest membrane-electrode distance we have successfully released was 200 nm, but the cavity design in section 5.1 required 500 nm to 600 nm.

Furthermore, the aluminum had micro-cracks after additional layer deposition or annealing—see fig. 4.8b—probably caused by a combination of the tensile stress and domain boundaries in the metal. To minimize the cracks, we tried sev-
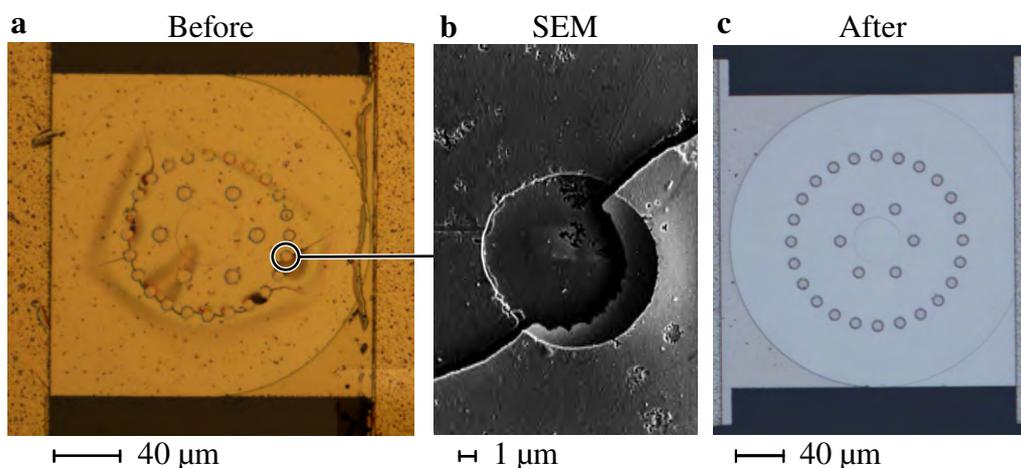
Figure 4.8: **a)** Microscope view. A 'horror picture' of the earliest incarnation of the aluminum membrane, after additional annealing of the chip and membrane release. **b)** SEM image zoomed on one of the holes in (a) showing how the membrane has broken, the poor alignment between the aluminum and alumina holes, and the many smaller cracks scattered all around. Those cracks are visible in the microscope as the many small black spots. **c)** Microscope view after process optimization.

eral things and found some improvements, but were unable to eliminate it. The two main improvements were prolonged exposure to O_2 plasma before annealing and lowering the deposition rate of aluminum. Figure 4.8 shows the final membrane-layer before and after the process optimization.

The last significant issue was that the membranes would break along the under-etch holes if annealed after release. Figure 4.8a demonstrate this effect. Therefore, we avoided exposing the samples to high temperature ($>120\,^{\circ}\,^{\circ}$C) after the isotropic etch. That restriction mostly affected the glues we used to pack the chip, demanding they cure at a low temperature.

4.2.5 Summary

Based on the test above, we decided to proceed with a process that under-etches a membrane made of aluminum and alumina. The release uses an isotropic plasma-etch based on SF_6 gas. This process uses a sacrificial layer between substrate and membrane to control the distance between them, delivering the short membrane-electrode gap required to reach a high electromechanical coupling. Furthermore, the process can produce the reflective membrane on top of a dielectric mirror, thus forming an ultra-short cavity. The sacrificial layer then also controls the

length of this cavity.

Compared to the previous nitride, these new membranes do not change the height and density appreciably, but the aluminum/alumina have notably lower stress. The reduction is about a factor of three[16], meaning the membrane's area must be smaller to reach the target frequency of 1 MHz. Both reductions improve the sensitivity and bandwidth according to eqs. (4.1) and (4.2), but the mechanical linewidth remains an open question. The linewidth is investigated in the next section.

4.3   Investigation of aluminum membranes

Disclaimer: parts of this section is adapted from a (Simonsen et al. 2019b). Some sentences and paragraphs may be very similar in language to that manuscript.

This section details the full fabrication procedure we used to make the new membranes, including mounting and packaging of the chips, together with the initial characterization of mechanical properties.

4.3.1   Method

Full fabrication flow

Simonsen et al. (2019b) describe the full fabrication flow that produced the membrane-capacitors, but it is nevertheless also reprinted here for completeness. The fabrication procedure is depicted in fig. 4.9; step-by-step. All processing used standard deposition, lithography, and etch techniques, along with standard cleaning and resist-strip procedures between each fabrication step. The process started from a standard fused silica wafer that was 500 µm thick and had a diameter of 10 cm, and produced 45 transducers in total, grouped in seven sets with identical membrane-design[17]. On top of the mirror, we built the membrane-capacitor chips as follows:

1. ALD deposition of a thin alumina layer, ~30 nm, to protect the substrate in the subsequent etch, particularly the isotropic etch in the last step. The machine could process up to five wafers at once. The first deposition included a silicon wafer at the center of the wafer-boat center; we used it as

---

[16]1.1 GPa for nitride versus 350 MPa for aluminum and alumina.

[17]Each set consisted of five to seven membranes.
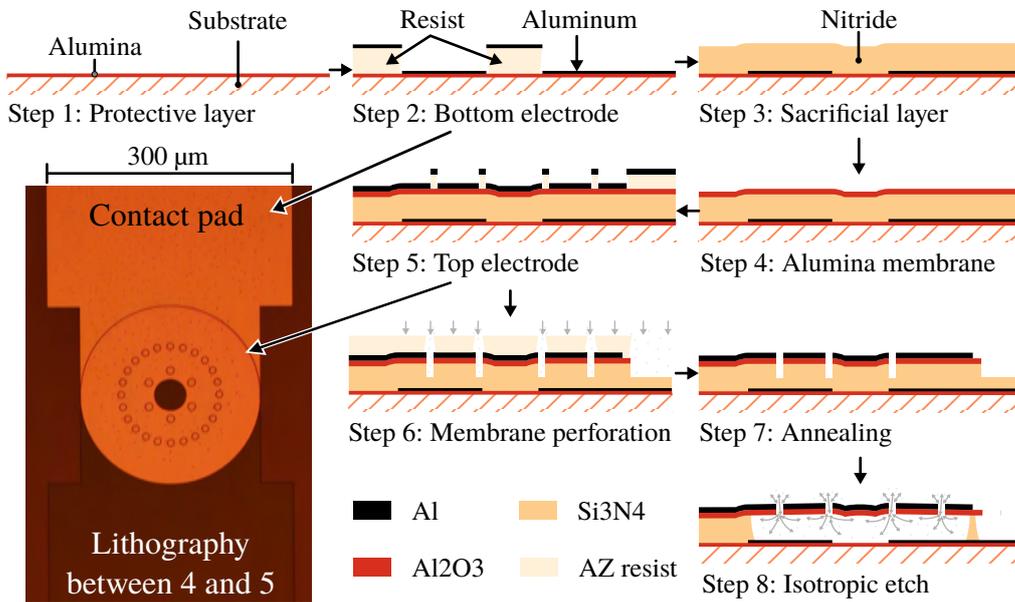
Figure 4.9: Step-by-step fabrication flow for the new membrane-capacitor. (*Adapted from Simonsen et al. (2019b)*.)

a reference of the thickness and refractive index measured with an ellipsometer. The variations in thickness over a boat were below 2 %, while the layer uniformity was below 1.5 %.

2. EBPVD of an aluminum layer, ~50 nm, to create the bottom electrode, patterned through UV-lithography and lift-off. The electrode includes a hole centered on the membrane that later defines the entry-point for the optical cavity, and it incorporated a pad allowing electrical connection to the electrode, very close to the membrane. Both features are visible in fig. 4.9.

3. PECVD deposition of sacrificial, low-stress nitride, 550 nm to 610 nm 500 to 650 nm thick, that sets the distance between top- and bottom electrode.

4. First membrane layer. ALD alumina that was 50 nm to 70 nm thick. It both fine-tunes the membrane-electrode distance and constitutes a part of the final membrane.

5. Second membrane layer. Aluminum, ~100 nm thick, made and patterned like in step 2, although with a different pattern shown in fig. 4.9. The contact pad does not overlap with the bottom electrode and, instead of a hole in the middle of the capacitor plate, there were smaller holes (from 3 μm to

10 μm in diameter) for under-etching arranged in circles around the membrane's center. The outer circle of holes determines the edge of the membrane; the membrane radius is about 30 μm larger than the outer circle's radius. The target membrane diameters were from 150 μm to 300 μm, giving frequencies of ≥1 MHz to meet a constraint from the circuit design—see section 6.1.1.

6. An anisotropic, dry-etch in an ICP transferred the perforation in the top aluminum into the alumina, hence exposing the sacrificial nitride layer for the isotropic etch. The chlorine-based etch was timed to stop in the sacrificial layer. A resist layer patterned with UV-lithography protected the top metal layer.

7. Annealed for one hour at 350 °C to induce tensile stress in the membrane electrode, as discussed above (section 4.1.1).

8. Finally, the isotropic dry-etch released the compound, alumina-aluminum membrane. Before the etching, we diced the wafer into chips that were 5 mm × 5 mm squares.

Packaging

After making the transducer chips inside the cleanroom, we brought them out, mounted them in a standard 8-pin circuit socket, and connected them to the socket-pins with wire-bonds to the two electrodes on the membrane-capacitor device—see fig. 4.10a. That way, the transducer can connect to an external through the socket pins, and the socket itself is easy to plug into a circuit. Additionally, the membrane-capacitors were protected by a 3D-printed plastic lid designed to cover the whole socket and attached with glue.

The final package also included a switch connected in parallel to the capacitor. When closed, this switch shorts both sides of the socket, which means the switch has to be open when operating the transducer. This shorting turned out to be necessary because the membranes would collapse at seemingly random times after wire-bonding. Figure 4.10b shows an example of a collapsed sample after wire-bonding. Often the transducer would collapse when transported from DTU to our lab at NBI, but occasionally they survived only to collapse at a later time for no apparent reason. This problem might occur because of electrostatic charges that build up between the floating membrane-capacitor pins—at
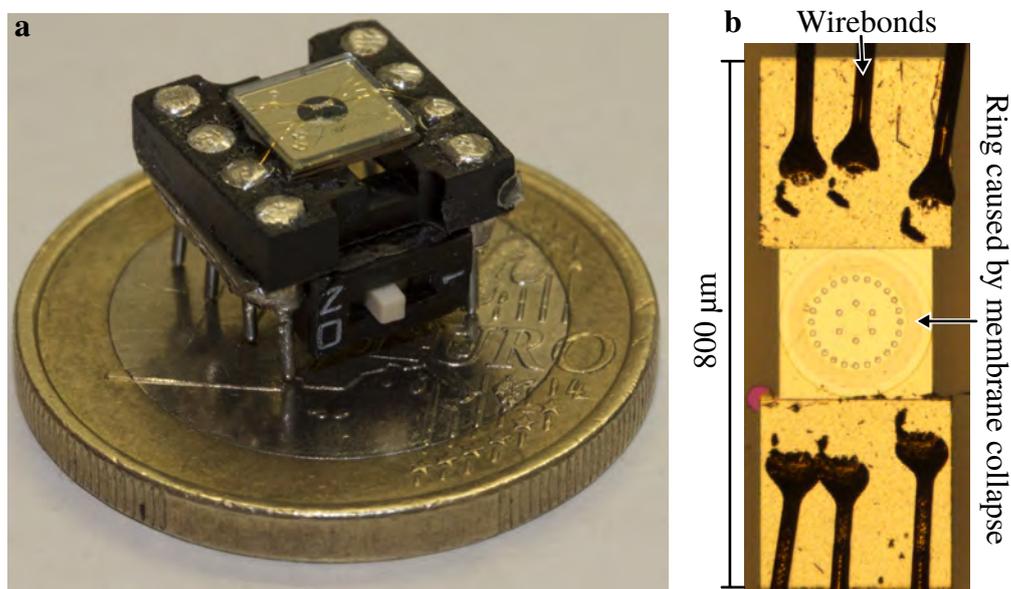
Figure 4.10: **a)** Final packaged device with the transducer chip mounted in a socket, and the membrane-electrodes wirebonded to the socket. **b)** Close-up view of a membrane that has collapsed and gotten into contact with the bottom electrode. The reflection is slightly darker where the membrane curves down-ward, which produces the characteristic ring around the smaller under-etch holes. (*Adapted from Simonsen et al. (2019b).*)

least the problem was solved by adding the short or, in fact, any circuit that electrically connects the transducer electrodes. In all the following experiments and associated sample handling, we switched the short on when connecting and disconnecting the transducer and only switched it off when another circuit loaded the socket pins.

The entire packaging process had also caused many membranes to collapse. Practically all these samples broke because of the wire-bonding, but we eventually eliminated this problem through several iterations of the assembly, leading to a packaging-procedure with an almost perfect yield. The issue with the wire-bonder arose from a powerful spark that is a part of the ball-bonding technique. This spark melts the wire into a ball before the machine presses that ball into contact with the transducer's electrode-pads and welds the ball onto the pad through a combination of force, heat, and ultrasound vibrations. If the membrane were close to the spark, then the membrane would almost always collapse. Therefore, we moved the packaged transducer away before manually creating the ball after the first wire-bond. However, once both electrodes had a connection

to the shorted pins, this was no longer strictly necessary. Afterward, we added multiple wire-bonds to both pads to reduced ohmic losses between wires and connections.

Setup for characterization

To characterize the samples and the repeatability of the fabrication process, we employed the same optical setup as in sections 3.1.1 and 3.2, i.e., the samples in $<1 \times 10^{-5}$ mbar vacuum and measured optically with either the vibrometer or interferometer at the center of the membrane, as sketched in fig. 4.11a. Like before, we obtained the fundamental resonance frequencies and Q-factor (eq. (2.43)) from fits to the thermally driven mechanical motion (eq. (2.42)), averaging the PSD thirty times and using a resolution bandwidth of 2 Hz—well below the mechanical linewidth. Through investigating several samples, we learned that these measurements, when taken with the vibrometer, had to happen fairly quickly. Prolonged exposure to the vibrometer-laser would otherwise make the mechanical frequency drift, and the measured Q-factor fluctuate significantly.

Figure 4.11a presents an additional characterization of the new transducers, one with the mechanical frequency measured as a function of an applied DC bias. These measurements used the interferometer and a low-pass filter formed between a $100\,\mathrm{k\Omega}$ resistor and $220\,\mathrm{nF}$ capacitor connected in parallel to the transducer socket through a coaxial vacuum feed-through.
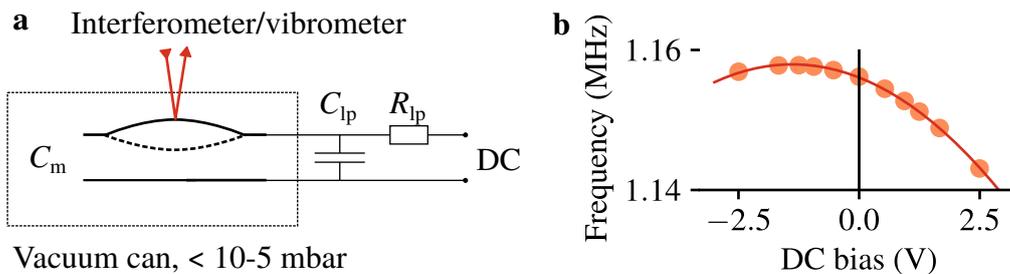
Figure 4.11: **a)** The setup used to characterize the samples, including how a DC bias connected to the transducer in measurements were this was relevant. **b)** Mechanical frequency obtained from fits to thermally driven spectra as a function of the DC bias. Sample parameters were $\tilde{m} = 3.9\,\mathrm{ng}$, frequency $\Omega_\mathrm{m} = 1.16\,\mathrm{MHz}$, and effective gap $d_\mathrm{eff} = 560\,\mathrm{nm}$. (*Adapted from Simonsen et al. (2019b)*.)

4.3.2 Results

Figure 4.11a displays how the membrane-resonance shifts with a DC bias applied in parallel to the transducer. These data fit a parabola with a second-order coefficient of $1\,\mathrm{kHz/V^2}$, but in contrast to the theory—specifically eq. (2.48)—the fit required a first-order term, roughly $3\,\mathrm{kHz/V}$, which effectively displaces the peak of the extremal point by about $1\,\mathrm{V}$. This effect was present in all investigated samples.

Beyond the DC bias measurements, we have investigated the bare mechanical properties for several samples and iterations of the fabrication design. As an example, fig. 4.12 shows the frequency and Qs of one of the first membrane releases, measured over nearly one year. Evidently, the mechanical frequency drifts down over time but reaches a plateau; the frequency decreased by about
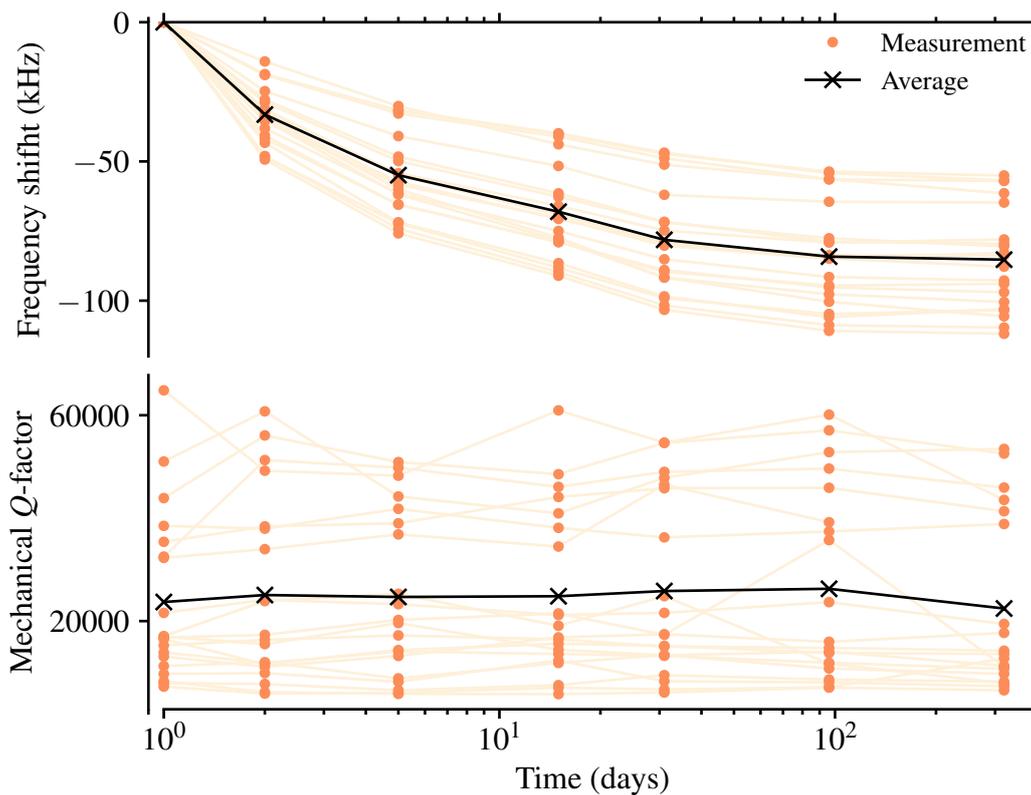
Figure 4.12: Mechanical frequency (top) and Q-factor (bottom) as a function of time for twenty membranes on the same wafer. This particular fabrication did not include the bottom electrode.

10 % during the first two weeks, but only another 2 % after nearly a year. These values were measured with the vibrometer. That means, as mentioned above, the frequencies and Qs drifts over between subsequent measurements; specifically, the resonance frequencies decreased slightly while the Q-factor fluctuated up to 50 % for its highest values and by ca. 2000 for most others.

For all tested samples, we have observed the following mechanical properties for membrane made with the fabrication procedure in section 4.3.1: the fundamental membrane resonances were between 1.0 MHz and 1.5 MHz and had a variation below 2 % for samples with identical geometry and from the same wafer. The mechanical Q-factors registered values between 5000 and 60 000 with the majority between 10 000 and 20 000—a much more pronounced spread compared to the frequencies—which fig. 4.12 also reflects. Every wafer had some samples, less than ten, with visible defects, inconsistent frequencies, or Q-factors below 10 000—and sometimes all three. Those membranes typically degraded in Q over time until no peak could be measured. Other samples possibly indicated a small decrease in Q-factor over time or from the packaging process. However, the average Qs in fig. 4.12 not feature that trend. In any case, there was usually 75 % or more samples with $Q \geq 10\,000$ after the fabrication.

4.3.3 Discussion

The mechanical Qs have a rather high spread, although most of them were 10 000 or higher. Possibly reasons for the spread include variation in film-quality due to micro-cracks along grain-boundaries and coupling between the membrane and the modes in the substrate, thus allowing phonon tunneling (Wilson-Rae 2008). The latter effect could explain why the optical probing changes the Qs in time, as argued by Jöckel et al. (2011) whose also witnessed a notable change in mechanical Q with prolonged optical probing. Ultimately, the limit on Q comes from material loss in aluminum Sosale et al. (2011) and Yu et al. (2012)

The issues of frequency drift (fig. 4.12) could become problematic for transduction in MRI. However, the drift is slow enough that a few days of measurements should be unaffected—it even plateaus in the long run—and it the effect can be canceled by changing the AC bias frequency. We believe the drift originates from stress-relaxation in aluminum because simple simulations[18] demonstrate a similar shift, assuming the stress in alumina remains fixed at ~350 MPa while the stress in aluminum starts at ~350 MPa but ends at ~200 MPa—both

---

[18]Finite-element simulations in Comsol.

these numbers are in line with Hyun et al. (2003). The issue was not reported in Bagci et al. (2014), Haghighi et al. (2018), Takeda et al. (2017), and Tominaga et al. (2018), presumably because they used nitride membrane with a very high stress that dominates over the stress in the metal, or because they did not anneal the metal or had it cross the membrane boundary.

Comparing roughly to the nitride membranes, the aluminum/alumina linewidth increase by a factor of ten, the area reduce by a factor four, the gap reduce by a factor of two or more, and the tensile stress reduce by a factor of four after the stress relaxation. The performance equations, eqs. (4.1) to (4.3), therefore predict the noise temperature and bandwidth increases by ~1.5 by ~6 respectively, while the cooperativity gets lower by ~0.4. The new device also lowers the pull-in voltage by a factor of two, according to eq. (2.47a), but it is still vastly above the actual operating voltage. These scaling refer to the performance with the old nitride membranes and the smallest membrane-capacitor distance, indicating that the new design makes it easier to reach the target bandwidth albeit with a modest reduction of the sensitivity. Furthermore, the new fabrication is a massive upgrade over the old flip-chip design. Actual transduction performance, of course, depends on the circuit, but the results in chapter 3 suggest that good performance in a sufficient bandwidth is achievable.

Unexpectedly, the frequency-shift (fig. 4.11b) did not scale with the bias squared, as previous analysis suggests (eq. (2.48) on page 29). We suspect the offset appeared because of trapped charges in the alumina layer, similar to the observation by Schmid et al. (2014) although our device did not exhibit any hysteresis with bias. A possible explanation could be that the isotropic plasma-etching leaves ions in the alumina membrane. Such trapped charges should only affect DC transduction; when added to the AC bias, they create a beat-note in the transduction theory (eq. (2.45)), but the beat-note frequency equals the bias frequency which is far away from the actual signal and mechanical resonance. As will be shown in section 7.2.4, these trapped charges lead to significant noise in MRI detection. That makes this problem a candidate for future improvements.

Chapter 5

Optical design

Disclaimer: parts of this chapter is adapted from a (Simonsen et al. 2019b).
Some sentences and paragraphs may be very similar in language to that manuscript.

With the fabrication flow in place, the next order of business is to figure out how to integrate it with an optical cavity. Here is a reiteration of section 4.1: our goal with the design was to realize a chip with a robust and alignment-free optical cavity. One could probably achieve this by inserting the membrane into an external cavity with fiber-coupling, but the most compact version has the membrane deposited directly on top of a mirror, and that is a possibility with the new fabrication process from the previous chapter. In that case, the cavity length is determined by the thickness of the sacrificial layer plus alumina, and the membrane has to act as the second mirror. This design eliminates the need to stabilizing the cavity length, but it sets steep requirements for the tolerance on the layer deposition and control of layer parameters. Section 5.1 first present a model for the layer stack comprising such a cavity, and an investigation comparing the model to data obtained with a new optical detection setup.

Fiber-coupling was the second primary objective with the new design. For that, the first attempt was simply a lens to focus and collect the light from fiber, and that proved to work decently enough for actual operation. Section 5.2 presents the full step-by-step fiber-assembly.

Simonsen et al. (2019b) describe the entire process flow, without splitting it up as done in this and the previous chapter.

5.1 Integrated optical cavity

Here is first the theory behind optical propagation through a dielectric slab, a theory that extends to multiple slabs in succession by simply multiplying corresponding propagation matrices together. The result predicts the reflection from the dielectric stack. Then comes a step-by-step fabrication flow describing how we made a stack with decent displacement sensitivity at the laser wavelength, before explaining the new optical setup and characterization of the cavity performance compared to the model.

5.1.1 Theory

This section presents a model for the optical propagation through an arbitrary stack of alternating layers. That model predicts the reflection from the transducer chip for the deposited layers and hence enabled us to design an optical cavity that resonated near the operating wavelength of 1064 nm. In essence, the model takes the refractive index of the layers used to fabricate the transducer and delivers the required thicknesses to make a cavity with an optical path-length near one full standing wave of light. The model can also estimate the cavity slope numerically. Doing so, for various layer thicknesses, led to the target designs were displacing the membrane should generate the biggest change in reflection.

When light propagates from point a to point b through a medium of refractive index n, its E- and H-fields evolves in a way presented by this propagation matrix (Brooker 2008, ch. 6):

$$\begin{pmatrix} E_a \\ H_a \end{pmatrix} = \begin{bmatrix} \cos\left(k\,l\right) & -i\frac{\eta_0}{n}\sin\left(k\,l\right) \\ -i\frac{n}{\eta_0}\sin\left(k\,l\right) & \cos\left(k\,l\right) \end{bmatrix} \begin{pmatrix} E_b \\ H_b \end{pmatrix}. \tag{5.1}$$

Here, k is the wave-vector of light, l is the propagation length through the layer, and $\eta_0 = \sqrt{\epsilon_0\mu_0}$ is a constant. In the final device, the light travels through multiple sections, each with a unique propagation matrix, which we handle simply by multiplying all these matrices together to form one general propagation matrix

$$\begin{bmatrix} m_1 & m_2 \\ m_3 & m_4 \end{bmatrix} = \begin{bmatrix} \cos\left(k_1\,l_1\right) & -i\frac{\eta_0}{n}\sin\left(k_1\,l_1\right) \\ -i\frac{n}{\eta_0}\sin\left(k_1\,l_1\right) & \cos\left(k_1\,l_1\right) \end{bmatrix} \times \ldots$$
$$\times \begin{bmatrix} \cos\left(k_N\,l_N\right) & -i\frac{\eta_0}{n}\sin\left(k_N\,l_N\right) \\ -i\frac{n}{\eta_0}\sin\left(k_N\,l_N\right) & \cos\left(k_N\,l_N\right) \end{bmatrix}. \tag{5.2}$$
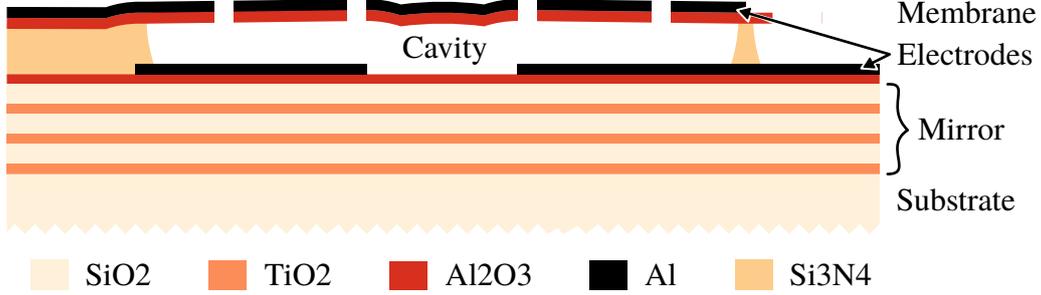
Figure 5.1: Cross-section of the transducer after membrane release showing all layers in the optical model. Layer parameters are in table 5.1. (*Adapted from Simonsen et al. (2019b).*)

The result is again a matrix relating the starting and ending E- and H-fields. Note that these fields technically depend on both incoming and outgoing light, evaluated at the interfaces that relate to the matrix. Therefore, they contain the information required to calculate the reflection from the stack. Assuming an input of one unit to interface a and no input to interface b such that light only leaves, then the reflection at interface a becomes

$$R_{\mathrm{m}} = \frac{m_1 \frac{\eta_0}{n_2} + m_2 - m_3 \frac{\eta_0^2}{n_1 n_2} - m_4 \frac{\eta_0}{n_1}}{m_1 \frac{\eta_0}{n_2} + m_2 + m_3 \frac{\eta_0^2}{n_1 n_2} + m_4 \frac{\eta_0}{n_1}} \tag{5.3}$$

for a general propagation matrix eq. (5.1), with n_a and n_b being the refractive index before and after the stack represented by the matrix.

The final cavity design started from a dielectric mirror made from alternating layers of high- and low-refractive-index material, TiO_2, and SiO_2, respectively, each with an optical path-length equal to a quarter wavelength. This coating served as input and output mirror for the integrated cavity. The coated glass-wafer also functioned as the substrate for the membrane-electrode fabrication covered in section 4.3.1, which led to the complete stack of layers illustrated in fig. 5.1.

5.1.2 Method

Table 5.1 gives the best-guess values for all individual layers in a specific fabrication run—the one used below to compare with data for a scan of the cavity-length. To fabricate this stack in the cleanroom, we followed the same procedure

Table 5.1: Thickness and refractive index of layers in the optical model (5.3) shown in Figure 5.1. The refractive index of the bottom electrode and nitride is irrelevant because the light does not travel through it. (*Adapted from Simonsen et al. (2019b).*)

Layer	Material	l_i (nm)	n_i
Substrate	SiO_2		1.45 [a]
Mirror layer (1, 3, and 5)	TiO_2	108 [f]	2.47 [b]
Mirror layer (2, 4, and 6)	SiO_2	183 [f]	1.45 [a]
Protective layer	Al_2O_3	49.7	1.65
Bottom electrode	Al	49.6	0 [d]
Spacer	Si_3N_4	605	1 [e]
Alumina	Al_2O_3	50	1.65
Top electrode	Al	100	1.03 + 9.25i [c]

[a] Malitson (1965); [b] DeVore (1951, $n(o)$); [c] McPeak et al. (2015); [d] Layer not in the optical path; [e] Layer etched away in the optical path; [f] Value assumed to be $\lambda/4/n_i$.

as described in the previous section. Below are additional details on the fabrication flow, given step-by-step, and how we characterized the layers to find the values in table 5.1:

0 A company made the first TiO_2/SiO_2 mirror coating on a fused silica wafer[1], so we used the expected refractive index and designed layer thickness. While buying wafers with a mirror coating is an expensive solution, we deemed it faster and more reliable than developing cleanroom process. The mirror design should have yielded a reflectivity near that of aluminum at 1064 nm, but an unforeseen process step removed the top layer titanium oxide layer, unfortunately, and brought the reflectivity down to about 79 %. It turned out that a mandatory wafer-cleaning procedure etches TiO_2.

1 On top of the mirror, we deposited the alumina layer to protect the mirror in the future isotropic etch of nitride. The layer's thickness and its refractive index were inferred from an ellipsometer measurement on a reference silicon-wafer from the same deposition run. This measurement had an uncertainty of about 1 nm according to quality-control data for the Atomic-Force Microscope (AFM).

2 The bottom electrode had a hole (40 μm in diameter) through its center,

---

[1] 100 mm diameter, 500 μm thick

permitting the light through and the metal-layer and into the cavity. The thickness was measured by scanning an edge of the metal near the wafer-center with an AFM. According to quality-control data for the deposition machine, this layer had a uniformity of 6 % across a wafer.

3 When making the sacrificial nitride layer, the deposition rate fluctuated between subsequent runs by 2 percent. Therefore, the target layer thickness was slightly below to needed membrane-mirror spacing, leaving a gap to be covered by the next deposition. Additionally, the thickness varied in uniformity by about 2 % (excluding the outer 1 cm of the wafer). To measure the thickness, we first etched away a small area of nitride at the center of the wafer—using the isotropic nitride etch from section 4.3.1—and then did the same scan with an AFM across the edge of the hole.

4 Given all the previous layer parameters and the cavity model eq. (5.3), the last alumina layer fine-tunes the length of the cavity to reach a chosen point on the cavity response function. The thickness was controlled by a timed process with the number of cycles determined from the previous ALD deposition.

5 Finally, the top aluminum thickness had to be thick enough to prevent transmission of light through it, a constraint demanding the aluminum be thicker than 40 nm. Otherwise, the metal would not be the last layer in the model.

Note that the membrane deforms, as shown in fig. 5.1, due to the center hole in because its surface topology gets replicated in all layers deposited on top of it. We tried to make this deformation small compared to the membrane thickness by making the bottom electrode thin. However, if the electrode gets too thin, it creates significant ohmic loss in the transducer. The final comprise was an electrode-thickness about three times thinner than the full membrane. According to initial simulation, the deformation should mostly even out at the membrane center due to the membrane's tensile stress. As discussed below (section 5.2.2), that simulation turned out to be questionable.

New fiber-optical setup

With a suitable optical setup, the integrated cavity replaces the interferometer and vibrometer because it converts the mechanical motion of its mirror into the
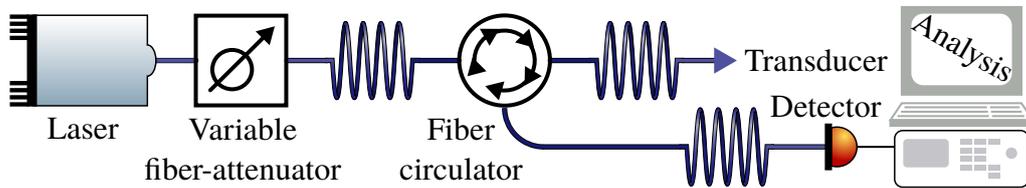
Figure 5.2: The fiber-network used with the integrated cavity for characterization and transduction.

optical modulation. However, light still needs to be sent to and from the transducer and end up in a detector. That motivated designing the new optical setup described here, a setup that had to be compact, robust, easy-to-align, and easy to transport to the MRI scanner. Basically, it had to be all fiber-components with a relatively small form-factor. With one fiber carrying light to and from the transducer, this network is compatible with the probe-head used for the interferometer, described in section 3.2 on page 47.

Figure 5.2 depicts the final fiber-network. It is a rather simple setup: light from a diode laser goes through a fiber-attenuator and into a circulator. The circulator then forwards the light to the transducer. The same circulator also distributes the light reflected from the transducer to a custom-built detector, thus preventing the back-reflection from going back into the laser. The fiber-attenuator further isolates the laser from reflections while controlling the optical power and allowing a constant driver-current to the laser.

In the remaining parts of the thesis, all measurements with the new fiber-network kept the input power to the transducer below a conservative, self-imposed value of 1 mW to prevent laser-induced degradation or damage of the membrane. At these powers, the light already shifted the mechanical frequency down considerably, around 6 kHz at 700 μW, although higher powers were okay for the nitride membranes used for DC transduction in chapter 3. In addition, from here on all measured spectra took at least twenty averages of the PSD and often repeated the same measurement several times in sequence, with each repetition used for statistics on the fitted values. To record these spectra, we mostly used a lock-in amplifier and always set its resolution bandwidth far below the mechanical linewidth. Typically, this was between 1 Hz and 7 Hz.

Cavity investigation

Figure 5.3 presents the setup used to probe the reflection from the integrated cavity, a setup that scans the length of the cavity with a DC bias. That sweep is compared to the corresponding model (eq. (5.3)) in section 5.1.3. To get the reflection, we used the fiber-network from above, combined with the probe-head from the interferometer, fig. 3.7 on page 48, to focus light on a sample and collect the reflection. This approach gave the reflected light both from the cavity and right next to it, where the reflection is very high, as indicated in fig. 5.3a. Dividing these two measurements gives the cavity's reflection with all losses in the optical path calibrated out. To change the cavity-length, we pulled the membrane closer to the electrode with a DC bias. As explained in section 2.3.3, this static displacement, β_0, relates to the mechanical frequency-shift through (eq. (2.50) on page 30)

$$\bar{x}/d_{\text{eff}} \approx -1.6 \, \Delta\Omega_m/\Omega_m, \tag{5.4}$$

and we determined this shift by fitting Lorentzian functions to thermally-driven mechanical spectra. Figure 5.3b graphs the displacement versus DC bias for the selected sample. Note that the cavity length only changes in one direction, i.e., it gets shorter. Therefore, this investigation was on a sample designed with a relatively long cavity such that the scan crosses the optical resonance frequency.

When the membrane-mirror distance gets shorter by β_0, it changes not only the overall cavity reflection, R_0, but also how much the membrane's motion modulates the reflection. As follows from section 2.2.3 on page 18, the RMS motion is proportional to the slope of the cavity, R_0'. Conversely, the slope may be estimated from the optically-detected, thermally-driven mechanical motion, i.e., the area under the Lorentzian function times the mechanical linewidth—where
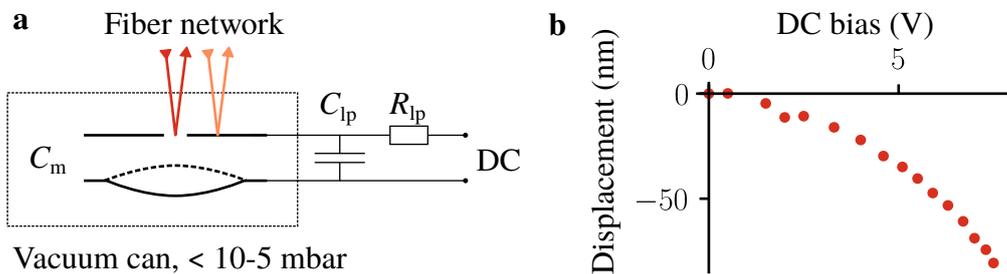
Figure 5.3: **a)** Setup for probing the reflection from the integrated cavity and sweeping the cavity length with a DC bias. **b)** Displacement versus bias. (*Adapted from Simonsen et al. (2019b).*)

the linewidth accounts for dynamical backaction as discussed in section 5.1.4 below. This estimate assumes that the bias does not add noise at the mechanical frequency. Therefore, we connected the bias to the membrane-capacitor chip through a low-pass filter, as shown in fig. 5.3, a filter that both reduces noise at the membrane frequency and prevents electromechanical coupling between bias and membrane. Instead of converting the detected signal into the units of the model—which would have required knowing the detector gain, optical losses, and the dynamical backaction—we chose to normalize both data and model to its value at 0 V bias, although that makes the resulting amplitude somewhat arbitrary.

5.1.3 Results

Figure 5.4 shows the recorded cavity scan of both the cavity reflection, R_0, and its derivative, R'_0. Also shown is the cavity-model, eq. (5.3), evaluated numerically using the values in table 5.1. The selected sample started from a cavity-length longer than the wavelength, but this was not the case for all samples. Nevertheless, other samples also changed reflection with DC bias, although the exact behavior was different. Note that the depicted data and model can agree qualitatively if the model gets a small correction: 7 nm added to the membrane-electrode distance. This added length is not from a fit but rather an estimate because fitting with this model poses significant challenges[2], and a better agreement between data and model would not improve the transducer-performance with the current design, nor offer additional, vital insights.

The cavity-scan revealed a broadening or narrowing of the mechanical linewidth that depended on the optical power at the sample. Specifically, the linewidth increased (decreased) with optical power when the cavity resonance wavelength was below (above) the laser wavelength, λ. Furthermore, in the case of linewidth narrowing, there was a threshold of optical power beyond which the membrane entered into self-sustained oscillations.

Apart from the oscillation regime, the broadening and narrowing always exhibited a constant area-linewidth product[3] for the mechanical peak, unless we deliberately defocused the setup by moving the lens. Then the area-linewidth

---

[2]This fitting has two main problems: one, there are many model parameters—even more than there are data-points—some highly correlated and all could potentially be fit-parameters; and two, the data and current model deviates significantly compared to the measurement error, suggesting the current starting-guess is too inaccurate to give a fast convergence.

[3]The area scales with the optical power squared and was therefore normalized accordingly.

would instead converge to a constant value for higher optical power, but it would deviate at low power. Finally, the mechanical linewidth broadened with DC bias, beyond what the optical effect seemingly explains. However, the area-linewidth product seemed to remain constant in the presence of this effect—otherwise, the derivative in fig. 5.4 would not match the model.
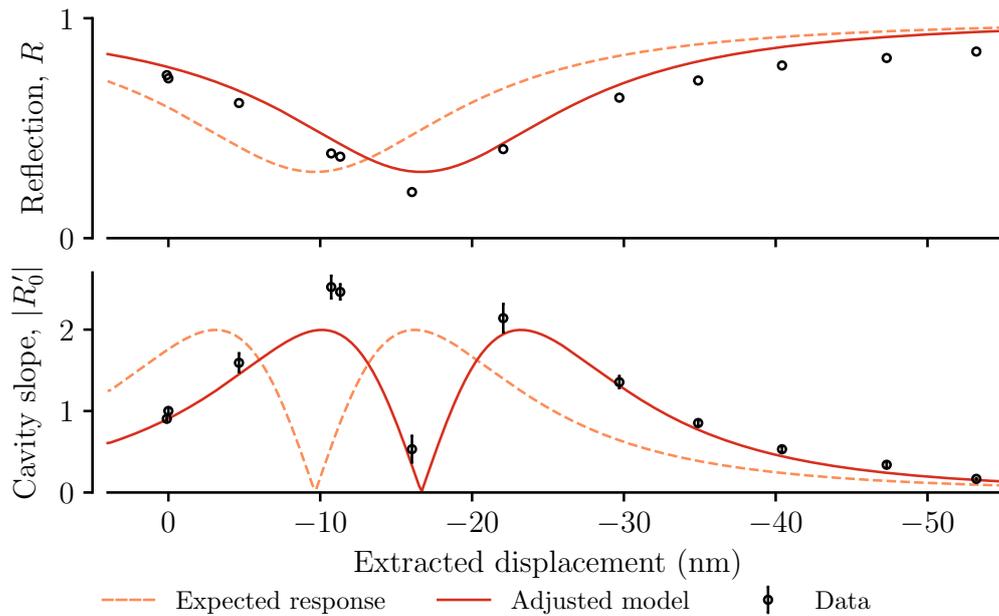
Figure 5.4: Cavity reflection (top) and its derivative (bottom) versus static membrane deflection at the membrane's center, together with the model from eq. (5.3) and its numerically-estimated derivative. Errorbars indicate one standard deviation. Throughout the measurements, the reflection remained constant, leaving no statistical estimate of the measurement error. (*Reprinted from Simonsen et al. (2019b).*)

5.1.4 Discussion

The samples exhibit clear indications of optomechanical backaction, i.e. laser-power dependent linewidth broadening, narrowing, and self-induced oscillation (Aspelmeyer et al. 2014; Metzger and Karrai 2004; Zalalutdinov et al. 2001), indicators that our samples display even though the cavities are in the (horribly) unresolved sideband regime,

$$\Gamma_o \gg \Gamma_m,$$

where radiation pressure dynamics vanish (section 2.2.3 on page 20); in fact, the cavity linewidth on fig. 5.4 is larger than $10\,\mathrm{nm}$, corresponding to about $2\,\mathrm{THz}$. That must mean the backaction comes from photothermal forces. A reasonable conclusion considering that aluminum absorbs about $5\,\%$ of light at the laser's wavelength.

Optical backaction does not change the product between the mechanical linewidth and peak area, just like we observed in measurement. The area-linewidth product must be constant because the peak-area is proportional to the RMS displacement (eq. (2.39) on page 26), a displacement that is proportional to the temperature T_{om}, and

$$T_{om}\,\Gamma_{om} = T_m\,\Gamma_m$$

as per eq. (2.40) on page 26. This observation, therefore, corroborates the theory of photothermal backaction even further, although we are unsure how the DC bias can add an effective backaction. Interestingly, the area-linewidth did not remain constant with deliberately unfocused light. This discrepancy can serve as a way to double-check the fiber-coupling after assembly.

The biggest issue with the cavity was the incorrect length, although an accuracy below $10\,\mathrm{nm}$ is a massive improvement over the DC transduction. There are two likely sources for this length: one, uncertainties in the thickness and refractive index in the cavity-model layers; and two, a deformation of the membrane as the topology from the bottom electrode gets stretched out after release. Uncertainties seem unlikely given that all layers were measured after deposition, either directly or from wafers in the same deposition run, and these measurements should have small tolerances. On the other hand, the deformation could be an issue. The original design used a simulation to verify the final geometry after release but after this discrepancy got found, we check it more carefully and found that the exact deflection depends on the parameters such as stress, Young's modulus, and the simulation boundary condition, with the deflection deviating much more than the $7\,\mathrm{nm}$ for some configurations. For Young's modulus, we

used table values for the bulk materials, but the modulus might actually be different when deposited with ALD, making the entire simulation questionable. This issue is one thing that could be improved upon for future transducer designs.

5.2   Fiber coupling

Light from the fiber lens has to hit the right point on the membrane, specifically through a 40 μm hole in the bottom electrode. That problem has three essential and mutually-dependent degrees of freedom that the assembly procedure must address: the lens placement, tilt, and focal-point. The final product should fix all these freedoms so that no further alignment is necessary when the transduction setup moves to an MRI scanner. To achieve this, we used three techniques to arrange the fiber-coupling, all shown after-the-fact in fig. 5.5: a secondary chip aligned relative to the membrane, a GRadient INdex of refraction (GRIN) lens pressed into contact with the device, and a guiding tube that fits both the lens and the ferrule of an optical fiber. The next subsections describe each step, one-by-one.
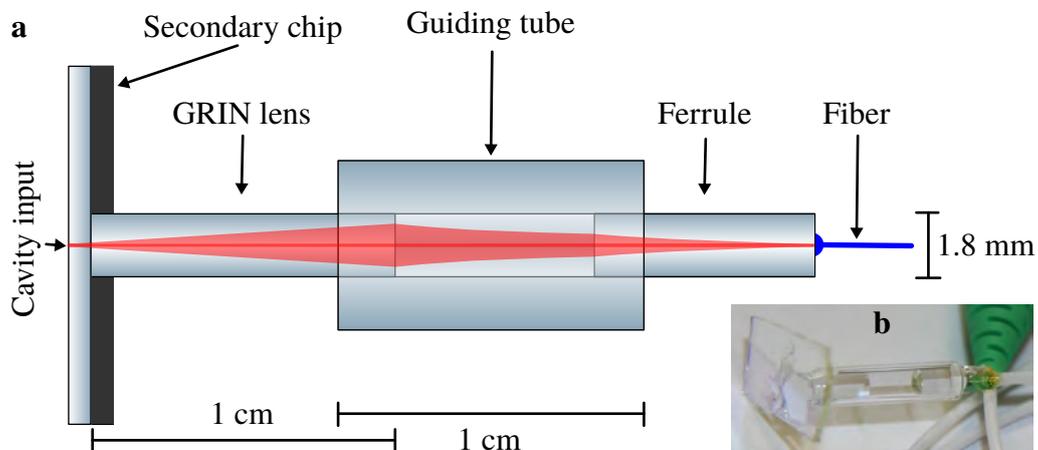
Figure 5.5: **a)** Fiber-coupling scheme to the glass-chip that holds the transducer to the far left and (**b**)) a photo of an early iteration without the secondary chip and only one electrode on the transducer chip (chosen for clarity). (*Adapted from Simonsen et al. (2019b).*)

Secondary chip

The idea here is that the secondary chip has a hole the same size as a lens used for focusing the light, such that the lens fits tightly into the hole. That means the hole can guide the position of the lens relative to the cavity. For instance, if the center of said hole coincides with the optical cavity accessed through the bottom electrode, then so will the center of a lens placed in the hole. Here is the fabrication flow for these silicon chips:

1. A silicon wafer[4] had an LPCVD nitride layer deposited on both sides. The thickness was enough to protect the silicon in the following etch.

2. The nitride layer was patterned by UV-lithography and dry-etch, thus defining circular holes the same diameter as the lens, with as little over-etching as possible.

3. The silicon wafer was then etched in KOH, an anisotropic etch that follows the silicon crystal planes (fig. 4.3 on page 69) and, therefore, left square excavations even though the original pattern was circular. The sidewall of the hole started at an angle corresponding to the {111} crystal planes but straightened to become vertical once the KOH had etched through the wafer (Kim et al. 2012). We examined this straightening with a microscope to determine the necessary etch-time. Eventually, the sidewalls would bend slightly inwards, as shown in fig. 5.6b, at which point the holes had become slightly larger than the original pattern due to under-etch.

4. The nitride on the wafer was removed after making the guiding holes because it no longer serves a purpose. However, the freely suspended nitride can rip off and stick to the chips. The etchant was buffered HF, which does not attack silicon.

5. Lastly, the silicon wafer was cleaved into individual chips that had the same size as the transducer[5].

The alignment procedure went like this:

1. The transducer-chips had an alignment layer added on their backside, a layer made from aluminum patterned with UV-lithography followed by

---

[4]100 mm diameter, ⟨100⟩, 350 µm

[5]5 mm × 5 mm

a Microscope view,
membrane-chip
(backside)

Secondary chip ———

Alignment layer
ø 1.8 mm

Cavity access

b Cross section

⊢ 20 µm
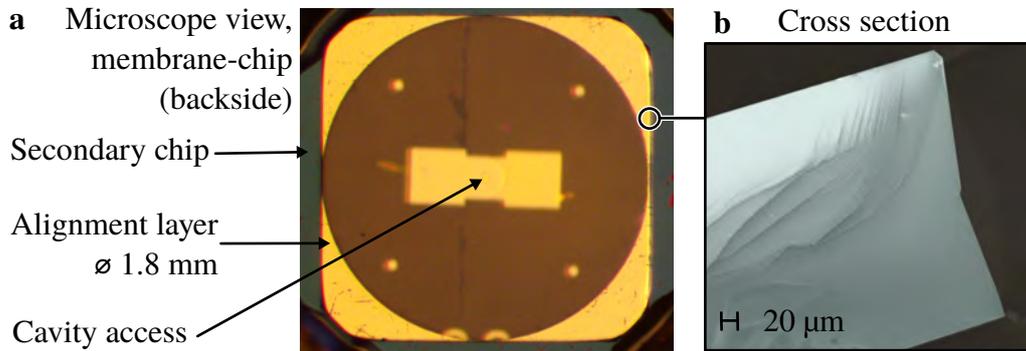
Figure 5.6: **a)** The back-side of the transducer with the secondary chip aligned and glued on. Notice the overlap between the edge of the alignment layer and the hole in the secondary chip; it indicates adequate vertical alignment but poor horizontal alignment. The cavity access is out of focus because it is on the front side of the chip, which means the view goes through the transparent substrate. **b)** False-colored SEM depicting the cross-section of the silicon chip inside the center hole.

> lift-off. This layer defined a circular hole matching the outer diameter of the lens. The center of the hole was aligned to coincide with both the membrane-center and through-hole in the bottom membrane.

2. Right before alignment, the backside of the transducer had epoxy adhesive smeared along the edge of the chip. That glue hardens in a few minutes, leaving that much time for the alignment. The full curing took about an hour.

3. Using a microscope to monitor the position of both chips[6], we carefully placed the second chip on the backside of the transducer such that the two holes align with each other, as demonstrated in fig. 5.6a. This alignment was by hand (or by tweezer, as it were) and done one sample at a time, which is a rather tedious exercise. Alternatively, the alignment could come from bonding the two full wafers together before dicing, although that requires more process optimization.

---

[6]A vital trick was to turn the chip under the microscope and look at the alignment from multiple angles. It could be that our microscope is slightly misaligned and skews the image a bit to the side.

GRIN lens

With the secondary chip in place, the next was to insert a GRIN lens into the guiding hole. The advantage of a GRIN lens, compared to regular lenses, is its flat end-surfaces. That means pressing the lens against the chip substrate ensure the two surfaces are parallel, thus aligning the tilt between the lens and membrane. The interface between lens and substrate could cause unwanted reflection if there is an air gap between the components, thus creating a change in the refractive index. At the same time, the lens had to be attached to the substrate somehow. A solution to both problems was to fill the gap with an adhesive whose refractive index matches that of glass. Specifically, we used a UV-curable glue and exposed it with UV-light from a custom-built source.

Laser focusing

The GRIN lens comes with a glass-tube for aligning the lens to a fiber terminated in a ferrule—a commercial solution from Thorlabs. The fiber-type matches the optical network in section 5.1.2. The guiding tube snugly fits both lens and ferrule. Inserting them into the tube keeps them centered and without tilt relative to each other. As a result, the only remaining degree freedom is the distance between lens and ferrule. Changing that distance moves the focal point of the lens, which is the last step needed to couple light into the cavity and extract the reflection. If the cavity is the most prominent source of reflection, then the optimal lens-position occurs when the focal point hits the plane of the cavity. At this condition, the reflection can travel back along the same optical path to reach the fiber provided the substrate-ferrule angling is correct. Figure 5.5a demonstrate what this end product of this fiber-coupling looks like for one of our early attempts that did not yet feature the second chip.

After several tests, it was clear that only monitoring the reflection led to an unreliable assembly. Therefore, we also injected a signal into the transducer and monitored the modulation-amplitude caused by the signal. This additional measurement guarantees that a motion of the membrane is detectable. Without it, a small tilt of the ferrule can displace the focal point away from the cavity where there still is a reflection but no movement. Interestingly, the modulation amplitude did not peak at the same lens-ferrule distances as the reflection did, although the peaks seemed very close (relative to our experimental control of the lens-ferrule distance). Because the modulation-amplitude directly equals displacement sensitivity, we tried to align to this peak rather than the peak reflectiv-

ity. The final fiber-alignment procedure was this:

1. The GRIN lens, attached to the transducer, had a guiding tube pushed a few millimeters down over it and attached with the UV-curable adhesive.

2. The other end of the tube had the ferrule-terminated fiber inserted, and that fiber delivered laser light to the sample.

3. The fiber also connected to the same laser setup as in fig. 5.2, which distributed laser light to the sample and the reflected light to a detector. The detected reflection amplitude went to an oscilloscope for real-time monitoring. In turned out that there were multiple reflection-peaks as a function of the lens-ferrule distance, probably from higher-order, transversal modes in the integrated cavity.

4. The modulation-amplitude was obtained through a lock-in amplifier connected directly in parallel to the transducer. It supplied the signal[7] at a frequency far away from any expected mechanical resonance[8] and demodulated the detector's output at the signal frequency to recover the amplitude of transduced modulation. The lock-in output also went to an oscilloscope for real-time monitoring.

5. Finally, careful tuning of the ferrule resulted in a position where modulated amplitude peaked, and where the reflection was high. This positioning used a micrometer stage to move the ferrule controllably back and forth, together with some hand-operated turning and angling of the fiber. In the end, the position was secured with the UV-curable adhesive. Annoyingly, the reflection and modulation amplitude would drift during the curing, probably due to thermal expansion in the optics, an expansion that could also cause a small lens-displacement as the glue hardens.

---

[7]Because of the trapped charges, this signal will modulate the transducer at the signal frequency, without any added DC offset.

[8]Note that the setup was not under vacuum, meaning air pressures had broadened the mechanical resonances to the point where they were not discernible over the optical readout noise.

5.2.1 Results

The sample characterization used the same optical setup as section 5.1.2, but this time without an optical setup for focusing. Instead, the circulator's fiber-connector linked to the fiber-coupled sample through to a custom-made fiber-feedthrough (Abraham and Cornell 1998) going into the vacuum chamber,

$$<1 \times 10^{-5} \, \text{mbar},$$

that held the sample (the same chamber as in fig. 3.7 on page 48). Measuring several samples like that showed that the fiber-alignment procedure yielded 15 % to 25 % of coupling efficiency between the incoming and returning light. This estimate relies on the modeled cavity reflection to adjust for optical losses, although the exact model-adjustment was unknown for most samples, making the numbers somewhat imprecise. Additionally, the fiber-coupled samples featured peak-to-background ratios from 23 dB at 100 μW to 34 dB near 1 mW of optical power. According to eq. (2.68) on page 35, these ratios correspond to optimal cooperativities C_{opt}—including photothermal backaction—between 14 to 50. Finally, the fully assembled samples showed a scaling of the area-linewidth product as described in section 5.1.3 above (no fiber-coupling and deliberately misaligned focus).

Since the photothermal backaction has to be known to assess the transducer performance, we determined it for the fiber-coupled sample used in the next chapter. Figure 5.7 shows our extrapolation: a straight-line fit to the measured mechanical linewidths as a function of optical power. The linewidth goes from

$$\Gamma_{\text{m}} = 124.5(6) \, \text{Hz}$$

at no optical power to

$$\Gamma_{\text{om}} = 243(4) \, \text{Hz}$$

at the operating power of 700 μW, which means the effective membrane-temperature is (eq. (2.40))

$$T_{\text{om}} = \frac{\Gamma_{\text{m}}}{\Gamma_{\text{om}}} T_0 = 152(3) \, \text{K}.$$

Figure 5.8 displays the thermally-driven membrane peak for the sample used for AC transduction in the next chapter, together with the Lorentzian fit eq. (2.42) on page 27. This sample had an expected effective mass of 2.4 ng, a fitted resonance frequency of 1.31 MHz, and a thermal RMS displacement of 3.6 pm according to eq. (2.39). This RMS displacement equals the PSD integral (eq. (2.30)
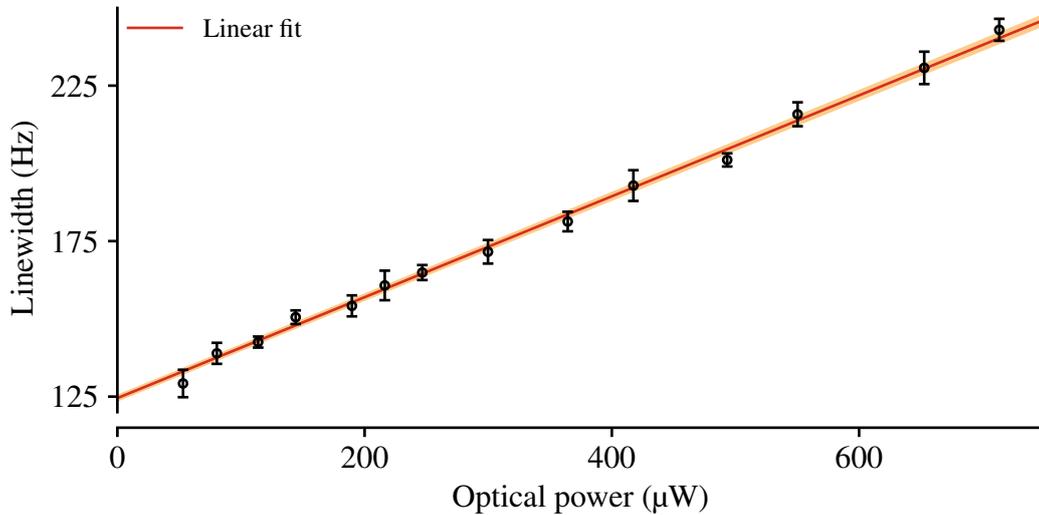
Figure 5.7: Mechanical linewidth obtained from Lorentzian fits to the thermally-driven displacement spectra. Figure 5.8 shows an example of one such fit. Each point averages over five spectra recorded in succession, with the standard-deviation used as the errorbar. (*Reprinted from Simonsen et al. (2019b).*)
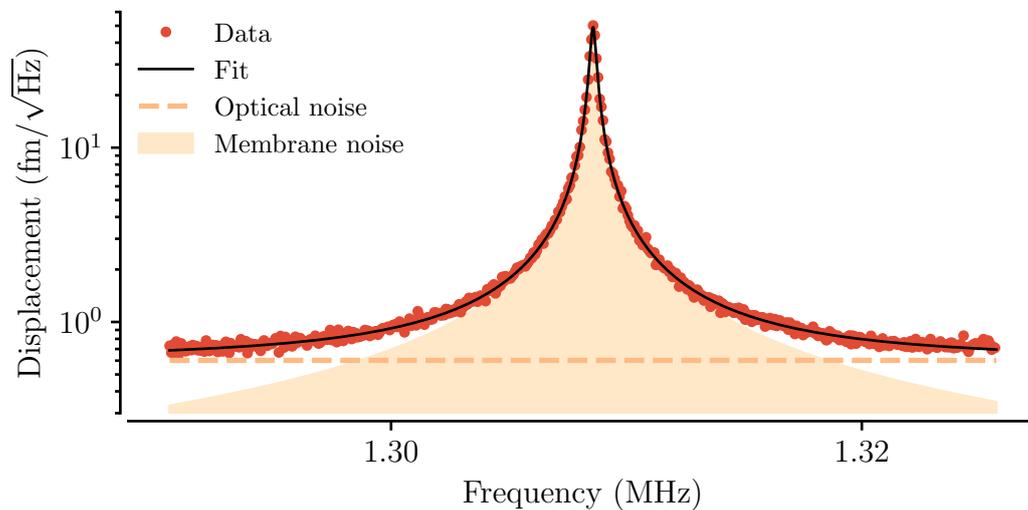
Figure 5.8: Spectrum of the thermally-driven, membrane displacement after many averages, converted into displacement units. The fit follows the Lorentzian function in eq. (2.42), consisting of the peaked membrane response plus a flat background from the optical readout.

on page 23), so we rescaled the spectrum accordingly, thus converting the y-axis into displacement. The fitted background after rescaling suggests that the optical displacement noise was $0.6\,\mathrm{fm}/\sqrt{\mathrm{Hz}}$. At the same time, the theory (eq. (2.19) on page 19) predicts this noise-floor for shot-noise of light to be $0.6\,\mathrm{fm}/\sqrt{\mathrm{Hz}}$, assuming these parameters: $R'_0 = 40/\mathrm{\mu m}$ and $R_0 = 0.34$, obtained from the cavity model eq. (5.3) in fig. 5.4; $\lambda = 1064\,\mathrm{nm}$; and an optical power of $700\,\mathrm{\mu W}$ going to the sample, but where only $5\,\%$ reach the detector due to losses in the circulator ($\sim 50\,\%$) and poor fiber-coupling.

5.2.2   Discussion

The adjusted cavity model again matches up to device performance in evaluating the optical displacement noise, and that is despite the fiber-coupled having some apparent issues with fiber-coupling efficiency, not to mention an area-linewidth scaling that is not constant. The latter indicates, according to section 5.1.3, that the focal-point lacks sufficient accuracy during alignment. These two issues are obvious candidates for future improvements.

Despite a suboptimal assembly, the selected device promises an optimal cooperativity $C_{\mathrm{opt}} \sim 80$, including the optical backaction which means this cooperativity assumes a mechanical temperature of $152\,\mathrm{K}$. When referring instead to the physical temperature, the optimal cooperativity is about 160. That is similar to the optimal performance described in chapter 3. Besides, the new device improves many other technical aspects of transduction, and it has room to improve the optical readout as well.

Part III

Transduction for magnetic resonance imaging

Chapter 6

AC transduction

Disclaimer: parts of this chapter is adapted from a (Simonsen et al. 2019b). Some sentences and paragraphs may be very similar in language to that manuscript.

This chapter explains our characterization of the new transducer-device when coupled to a circuit designed explicitly to detect the MRI signal in the scanner. These tests were done at our lab at NBI,[1] assisted by Sampo Saarinen, using a circuit designed and made by Juan Diego Sanchez. Details on the circuit fabrication are omitted here; instead, this treatise merely aims to justify the circuit layout and model (sections 6.1.1 and 6.1.2). The experimental procedure was similar to the DC transduction in chapter 3; like before, shielding the setup was important to reduce noise, and the noise-temperature extrapolation used a Y-factor-like technique. Unlike the previous Y-factor experiments, this new technique injected voltage noise into the circuit instead of cooling a resistor (section 6.1.4), a voltage noise that we estimated from an elaborate model-fit to scattering parameters from the circuit (section 6.1.2).

6.1   Method

This section explains the new circuit that implements transduction with an AC bias instead of a DC one (section 6.1.1) and uses the new transducer design. For future reference, the target NMR frequency was

$$\Omega_{\mathrm{NMR}} = 32.19\,\mathrm{MHz},$$

---

[1] The next chapter describe transduction in the MRI scanner

and the particular transducer sample had a mechanical frequency

$$\Omega_{\mathrm{m}} = 1.31\,\mathrm{MHz}.$$

A fundamental change to the circuit was placing the pick-up coil close to the main PCB and transducer. That meant we could no longer dunk the inductor in liquid nitrogen to cool it, as done in section 3.2.1. Consequently, we had to develop a new way to vary the noise-temperature of the circuit to estimate the intrinsic transduction noise. Our new approach was to inject a known voltage noise into the circuit, thus increasing the effective noise a known amount instead of decreasing it. We did so through a second coil—the probe coil—mounted to couple weakly to the detection setup. However, knowing the induced voltage required us to convert the voltage in the probe coil into an equivalent noise in series with the detection coil, and that task necessitated an accurate model of the full circuit with probe-coil and all (section 6.1.2). That model also allowed us to estimate the Johnson noise in the circuit at room temperature, which we used to convert our inject voltage noise into an equivalent added temperature. The final setup and procedure are explained in section 6.1.3 and section 6.1.4, respectively.

6.1.1 Circuit design

To move the transduction to a higher frequency, we designed a new circuit that resonated near the target frequency. Specifically, the *LC* resonance ended close to

$$\Omega_{LC} = 31.9\,\mathrm{MHz}.$$

The first step was to design a new coil for detecting RF signals. Compared to the previous inductors from section 3.1.1, this coil replaced the ferrite core with air to bring the frequency up, used thick silver wires to minimize ohmic losses, and left large spacing between coil-windings to avoid parasitic capacitance. Specifically, the coil was a flat spiral with four windings and an outer diameter of 50 mm, wound with a 1.6 mm diameter silver wire, as pictured in fig. 6.1a. The coil-inductance was estimated by simulation to be 490 nH. The surface area of the coil was $A = 5 \times 10^{-3}\,\mathrm{m}^2$, as determined by parameterizing a picture of the coil and numerically integrating over it.

Figures 6.1a and 6.1b depict the final resonance circuit as a photo and simple diagram, respectively. As shown there, the *LC* resonance appears by connecting the new detection coil, L_{d}, in parallel to both a tuning capacitor, C_{t}, and the transducer, C_{m}—the latter connected to the PCB through the shortest wires that

were practically feasible. This new circuit is similar to the one used for trans-
duction with a DC bias in the sense that it also forms an *LC* resonance, but it is
different in the way that the biased is filtered. While this filter and other circuit-
elements also affect the resonance, the final *LC* frequency mainly comes from
the detection coil and tuning capacitor. The remaining circuit only gives a small
perturbation of impedance, at least near the resonance frequency.

Had we connected the AC bias directly in parallel to the *LC* circuit, its *Q*-
factor would have degraded significantly as energy in the resonator would have
been able to leak out of this port. Technically speaking, the bias would have
loaded the *LC* resonance. Therefore, we isolated the resonator and bias by adding
a band-pass filter between them, a filter formed by the inductance T_f and capac-
itance C_f to feature a high impedance everywhere except at the bias frequency.
It connected to the *LC* resonator through a small coupling capacitor, C_c. Mak-
ing a good filter is difficult because its impedance has to change sharply with
frequency. As a practical limitation, we aimed to have the filter- and resonance-
frequency differ by more than 1 MHz. That criteria imposed a derivable con-
straint on the membrane frequency, namely that

$$\Omega_m \geq 1\,\text{MHz}.$$

Given the *LC* and mechanical resonance, we tuned the filter frequency to

$$\Omega_f = 30.6\,\text{MHz}.$$

The filter on the PCB served yet another purpose: to eliminate sideband noise
from the bias drive. This noise-cancellation is crucial to reach thermal-noise
limited transduction as any sideband noise at the detection frequency adds to the
total voltage noise. For example, Takeda et al. (2017) reported that phase-noise
from their AC bias dominated their ultimate performance. Our onboard filter
reduces sideband noise because it takes a larger voltage-drop at the detection
frequency compared to the bias frequency. A similar noise reduction comes from
the resonator itself as that enhance signals induced in the detection coil but not
the bias voltage applied in parallel. Unfortunately, the filtering offered by these
combined effects were insufficient to eliminate all phase-noise, so we sent the
AC bias through an additional, external band-pass filter before delivering it to the
detection circuit.
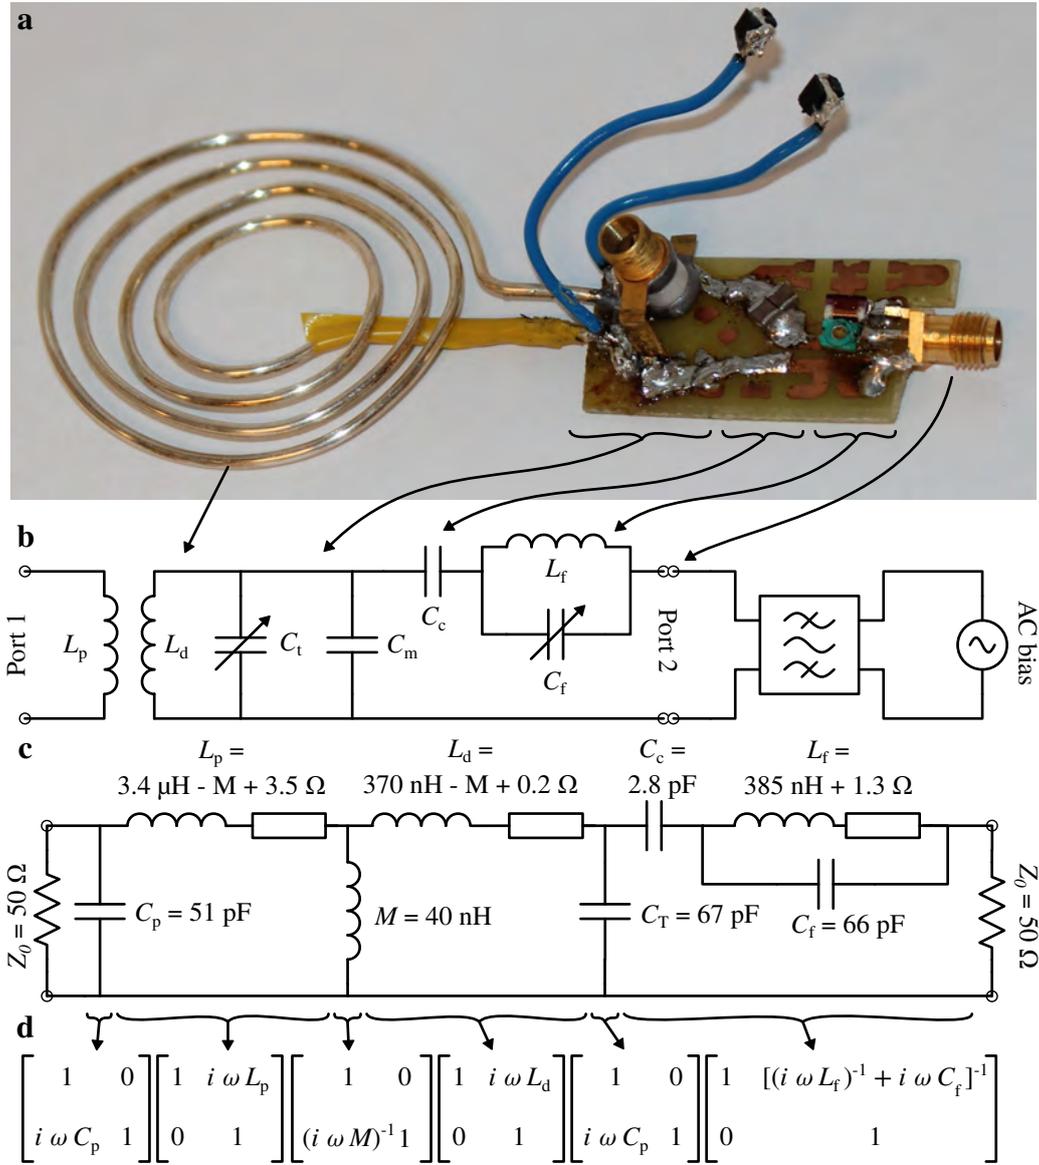
Figure 6.1: **a)** Photo of the detection circuit with the main coil on the left and a PCB with a filter, tuning capacitor, and blue wires that connect to the transducer. **b)** Simple circuit schematic for the full detection circuit, including the probe-coil T_p, external band-pass filter, and AC bias. **c)** The circuit model used to fit the real detection circuit with values obtained by fitting to scattering parameters. Both inductances from the probe and detection coil include an inherent ohmic loss, and they have the mutual inductance M. Additionally, both coils have parasitic capacitance included as C_p or lumped into C_T together with the membrane-chip C_m and tuning C_t capacitors. **d)** Complete transmission ($ABCD$) matrix that used for the fitting routine. (*Adapted from Simonsen et al. (2019b).*)

6.1.2 Circuit model

Figure 6.1c show the circuit diagram used to fit the scattering parameters measured between the probe coil (port 1) and the AC bias input (port 2). These scattering parameters, S_{kl}, correspond to the voltage output from port k, V_k, divided by the voltage input, V_l, to port l. Mathematically, this definition looks like (Pozar 2012, Eq. (4.41))

$$S_{kl} := \frac{V_k}{V_l}. \tag{6.1}$$

Usually, the in- and outputs are functions in the Fourier domain, and the scattering parameter refers to their ratio evaluated at a particular frequency. Consequently, the scattering parameters take complex values meaning they have both an amplitude and a phase. The latter is the phase-difference between in- and output. We used a network analyzer to extract the scattering parameters. Figure 6.3 presents these measurements in a frequency-window around the *LC* resonance, together with the fit we achieved with the model described below.

The model expression derived from the *ABCD* matrix formalism described by Pozar (2012, ch. 4.4). In this theory, each circuit section has a corresponding two-by-two matrix, written as

$$\begin{bmatrix} V_l \\ I_l \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} V_k \\ I_k \end{bmatrix},$$

that relates the voltage ($V_{l,k}$) and current ($I_{l,k}$) for the input (l) and output (k) port. Multiplying these matrices, in the right order, yields a new *ABCD* matrix for the entire network. Figure 6.1d shows this principle for our particular circuit. A standard matrix-transformation (Frickey 1994, table 3) converts the matrix-product back into scattering parameters. This transformation requires knowing the load impedance, both of which were equal,

$$Z_0 = 50\,\Omega,$$

in this case. Sadly, this model was not good enough to fit the circuit. It had to include two model-variables that correspond to a small phase-shift (ϕ_1 and ϕ_2) and amplitude scaling (*Amp*) of the scattering parameter. With that, the model

looked like this:

$$S11 = \frac{A + B/Z_0 - CZ_0 - D}{A + B/Z_0 + CZ_0 + D} \, e^{i\phi_1} \times amp, \tag{6.2a}$$

$$S21 = \frac{2}{A + B/Z_0 + CZ_0 + D} \, e^{i(\phi_1 + \phi_2)/2}, \quad \text{and} \tag{6.2b}$$

$$S22 = \frac{-A + B/Z_0 - CZ_0 + D}{A + B/Z_0 + CZ_0 + D} \, e^{i\phi_2}, \tag{6.2c}$$

where the A, B, C, and D are from the matrix-product in fig. 6.1d. The last modification was necessary despite our efforts to calibrate the network analyzer ports as close to the coil and circuit as possible, which should eliminate any deviations that are common between both calibration and measurement. However, that such phenomenological constants are necessary has been established in earlier works, for example by Petersan and Anlage (1998) who compares different algorithms for fitting scattering parameters from resonators. We added the two variables ad hoc but much inspired by these fitting methods. Note that these perturbations have been shown by others to be consistent with either cross-talk between ports (Leong et al. 1997), additional reflections between the ports (Deng et al. 2013), and a complex loading of the resonators (Khalil et al. 2012). The model with the added variables looks like this:

The fitting routine applied a Markov-chain Monte-Carlo algorithm (Foreman-Mackey et al. (2013) version 2.2.1 implemented in Python 3.6.3) to do linear least-squares minimization, a technique that samples the space of free fit parameters by incrementally changing the model-variables in a chain of steps. At each step, the algorithm calculates the mean-squared difference between data and model and then automatically determines from that result the next step in the model space. Importantly, this procedure fitted all scattering parameters, both their amplitude and phase, simultaneously. Unfortunately, the many free parameters made it difficult for the fit to converge without a good guess of the starting parameters. Therefore, we assessed as many parameters as possible from independent sources: the circuit components from their known or measured values, and the probe-coil parameters from a fit to a separate measurement without the detection coil nearby. The remaining parameters came from an initial fit using all these predetermined values as fixed constants. The final fitting left all parameters as free variables.

As mentioned above, it was challenging to make the fit converge. However, we have several indications that it indeed did. For example, multiple runs of the Markov routine revealed near-identical fit-parameters. Moreover, histograms
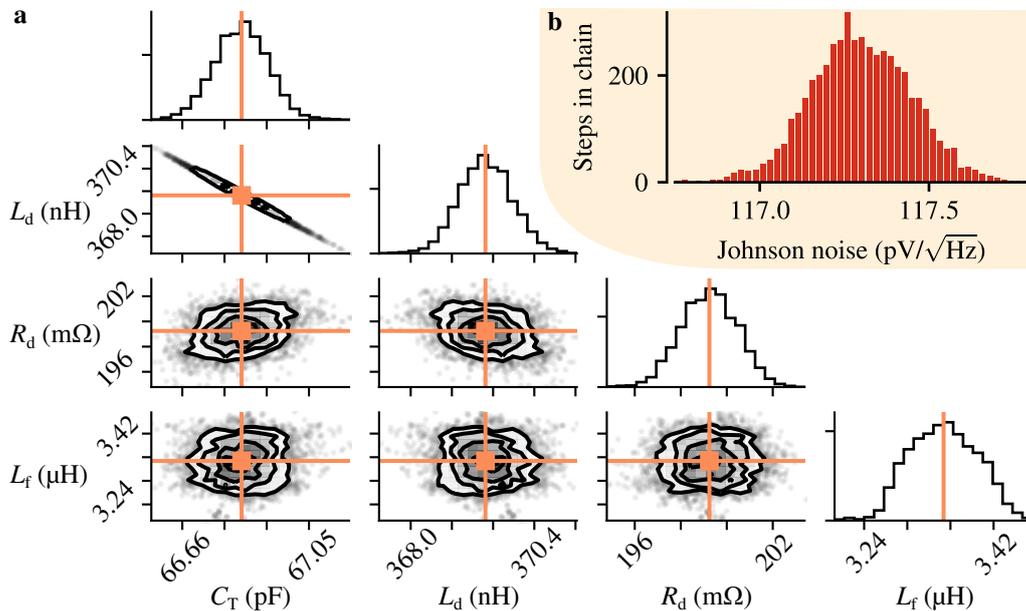
Figure 6.2: **a)** Surface maps showing the occurrences of two particular model values in the Markov-chain. The selected parameters are the inductance of the two coils, the ohmic loss in the detection coil, and the total capacitance for the *LC* resonance. **b)** The inset shows a histogram of the Johnson noise evaluated for each step in the Markov chain.

of the Markov-chains proved a bell-like correlation between all of the free parameters, as shown in fig. 6.2 for selected variables—another sign of convergence. Note that some histograms were highly asymmetric due to a strong correlation between model-parameters—most notably the inductance-capacitance pairs for the filter and *LC* resonance, with the later shown in fig. 6.2a. That observation is no surprise since both the *LC* resonance and filter frequency affects the fit strongly; however, either capacitance can cancel a change to the corresponding inductance, thus leaving the frequencies unperturbed. With a good fit-convergence, the Markov-chain that gave these histograms also represent the fitted values. For each parameter in the chain, the sampled mean and variance determine that parameter's value and error, respectively.

Figure 6.3 illustrates the excellent agreement between the final fits and data, together with the residual difference between measurements and model. To quantify the agreement, note that all residuals are within 1 % of the maximum value of the curve—with most of them far below. To account for measurement noise in the residuals, we added a random vector to the modeled scattering parameters, a
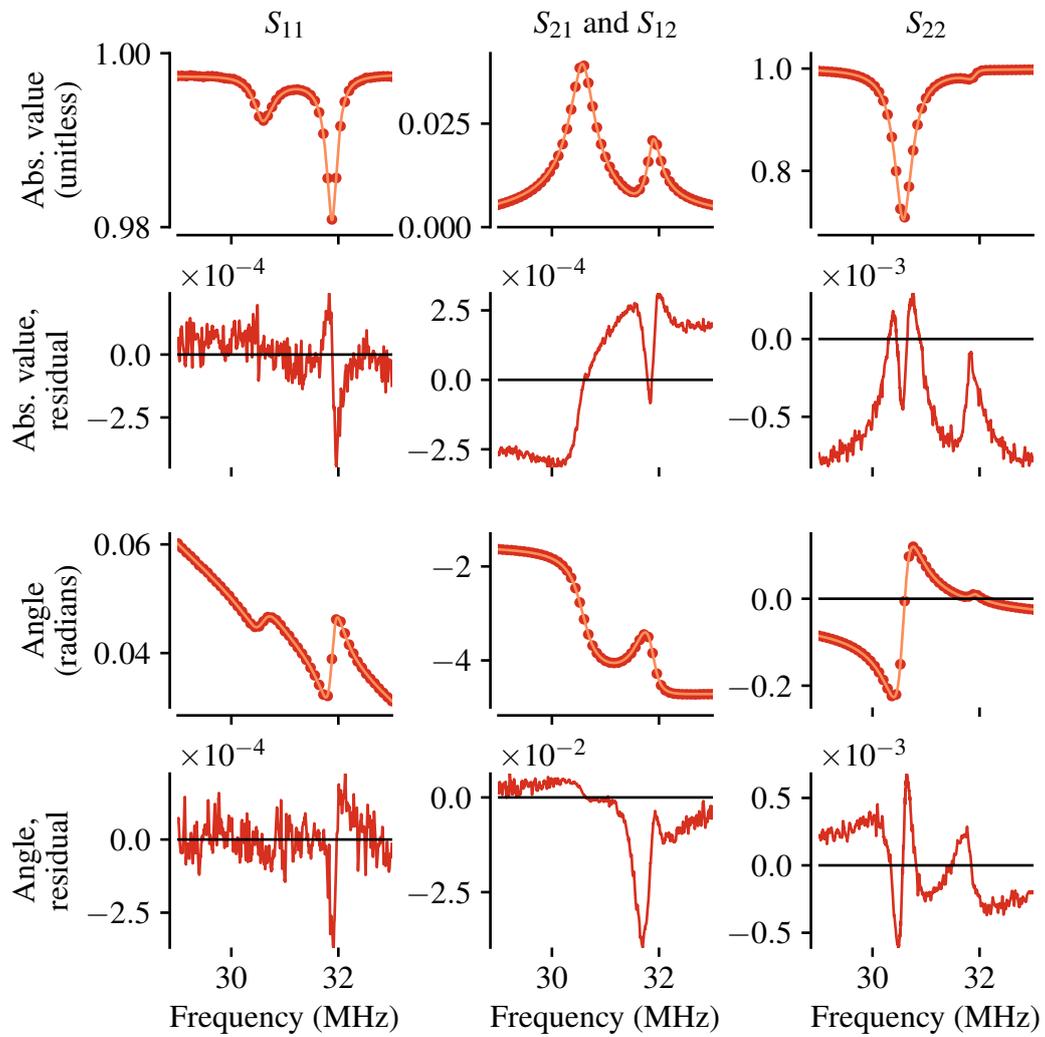
Figure 6.3: Measured (red circles) scattering parameters, plotted together with the model-fit from eq. (6.2). Only every twentieth data-point is shown. The curves directly below show the residual difference between the data and model. Port 1 and 2 are indicated in fig. 6.1b. (*Adapted from Simonsen et al. (2019b).*)

vector with entries from a normal-distribution whose width was a fit-parameter. That width turned out to be 3×10^{-4}, which is too small to explain the residuals, hence implying that the real circuit has small, systematic effects that the model cannot reproduce. Nevertheless, all fitted parameters deviated from their pre-determined values by only 1.5 % to 12 %. That agreement validates our model as a good approximation of the real circuit parameters. Lastly, both the added phase-shift and amplitude scaling turned out to be small perturbations: the phase was >11 mrad for each port, and the amplitude scaling was <0.25 % of S_{11} but negligible for S_{21} and S_{22}.

With the circuit-model shown fig. 6.1c, it is now possible to convert a voltage across the probe-coil into an equivalent voltage in series with the detection coil. That voltage then adds onto the Johnson noise, thus increasing the total voltage drive and, consequently, the effective temperature of the coil. The injected voltage should be a known quantity or at least possible to measure independently. We chose to convert the equivalent series-voltage into an added noise-temperature by referencing it to Johnson noise evaluated at room-temperature, 300 K.[2] To estimate this noise-level, we used the Markov-chain to calculate the full circuit impedance Z_{AC} in parallel with the transducer, took the real part of it, and used the standard expression for Johnson noise (eq. (2.31b) on page 24). That generated the histogram in fig. 6.2b, corresponding to a Johnson voltage of $117\,\mathrm{pV}/\sqrt{\mathrm{Hz}}$.

6.1.3 Setup

Figure 6.4 displays the detection circuit mounted inside a vacuum chamber—a domed glass tube—and enclosed in a big metal box that shields from RF ambient noise. Importantly, the vacuum chamber consisted of glass near the primary coil. Had it instead been made of metal, then it would have spoiled the coil's intrinsic Q because the coil would have radiated energy into the metal where it would dissipate. For the same reason, the surrounding metal shield was large enough to keep its sidewalls away at least 20 cm from the coil, in all directions. That distance yielded only a small effect on the coil, as determined in a separate measurement of the circuit's scattering parameter, S_{22}, with a metal-plate held at different distances from the coil.

A vacuum pump evacuated the chamber far below 1×10^{-3} mbar and pumped

---

[2]Remember that noise-temperature scales with the PSD (eq. (2.31b) on page 24), i.e. voltage squared, so the noise-voltages adds in quadrature.
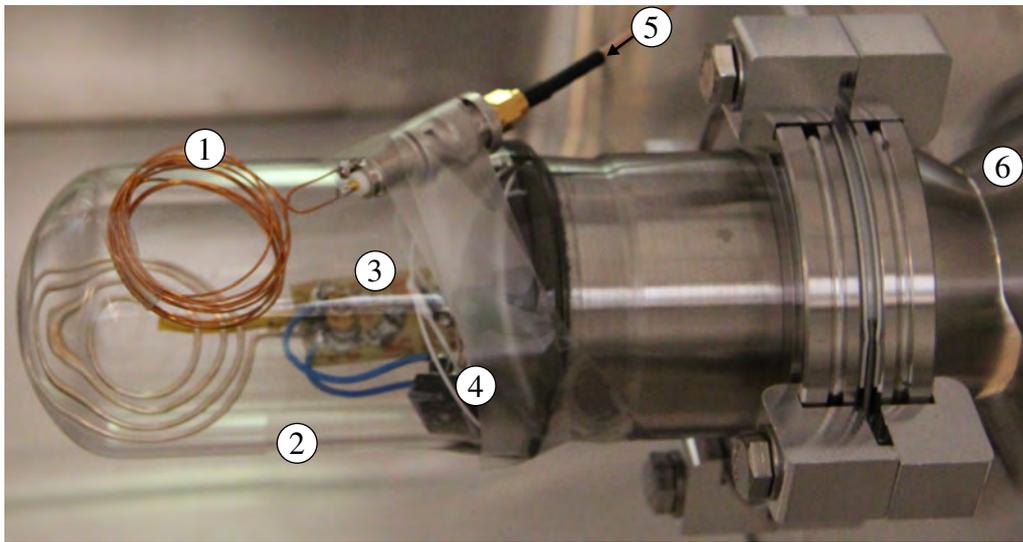
Figure 6.4: Setup for characterizing the AC transduction setup showing the probe-coil (1), glass-flange on the vacuum chamber (2), detection circuit (3), the packaged transducer (4), and the input cable to the noise-injection (5), all enclosed in a metal shield (6).

continuously through transduction experiments. The fiber entered the vacuum through a custom feedthrough (Abraham and Cornell 1998), while the bias entered through a standard feedthrough for coaxial cables. The probe-coil was outside the vacuum chamber yet still placed within a few centimeters of the detection coil, with both coils oriented to create a small mutual inductance between them. To do the alignment, we moved the probe-coil around while monitoring the probe's scattering parameters, S_{11}, until the showed small dip at the LC resonance frequency. Then we fixed the probe-coil with adhesive tape.

Without the big shielding box in place, we observed added environment noise driving the membrane, just like in section 3.2.2. However, the box removed all ambient noise without additional optimization. All cables and their connectors were merely standard coaxial ones, while all data-acquisition machines and power-supplies shared the same electrical ground, a ground connected to the optical table. The AC bias was connected directly to the external filter and vacuum-feedthrough outside the box before going to the detection circuit. Similarly, an arbitrary waveform generator injected a voltage directly into the probe-coil using a cable that entered the shielded box through a small hole drilled in the side of the box. When measuring the scattering parameters, the network analyzer replaced both the AC bias and the waveform generator but used the same cables

and connections.

The optical setup was the same as in fig. 5.2 on page 90. The final optical detection down-converts the electrical signal near the *LC* resonance, Ω_{LC}, to a frequency near the mechanical resonance, Ω_{m}. After this transduction, the output goes straight to a spectrum analyzer for post-processing.

6.1.4   Procedure

Before doing the actual Y-factor extrapolation, we tuned the AC bias to find the power that optimizes the transduction. Note that this point is not where the cavity is most sensitive because the bias changes the cavity response (section 5.1). As explained in section 2.4.2 on page 34, the optimal cooperativity equalizes the thermally-driven mechanical motion to the optical readout noise, evaluated at the mechanical resonance frequency. Finding this point was complicated because the photothermal backaction change with bias, not to mention that electronic noise also drives the mechanics with the AC bias turned; and when searching for the optimum, we did not yet know if the shielding had eliminated all ambient RF noise. Nevertheless, we found a near-optimal power after several iterations of a reduced Y-factor extrapolation that gauged the performance at each power. In the end, the optimum, sinusoidal AC bias had a root-mean-square voltage of 1 V for the membrane-sample selected for transduction.

With the AC power established, we proceeded to inject a white-noise signal, at different powers, into the probe-coil using a waveform generator, and we recorded the resulting mechanical PSD with a spectrum analyzer. Each recording averaged the spectrum twenty times. Note that the measured frequency range is around the mechanical frequency Ω_{m} because the signal has been down-converted by the interaction with the bias. A separate measurement of the white-noise gave the corresponding voltage-noise for each setting, measured a function of frequency around the *LC* resonance Ω_{LC}. Those measurements also provided the variance of the voltage-noise latter used as error bars in the Y-factor fit.

Figure 6.5 demonstrates a selection of these noise-driven measurements at the chosen AC bias power, together with their corresponding fits; these fits were a Lorentzian plus an offset like in eq. (2.42) on page 27. The figure also shows each contribution to the intrinsic transducer noise, with the mechanical and electrical contribution obtained by the Y-factor extrapolation. To derive the latter, we used the fitted peak-value of each spectrum because this corresponds to the total noise-power in the system at the most sensitive frequency, as required for the
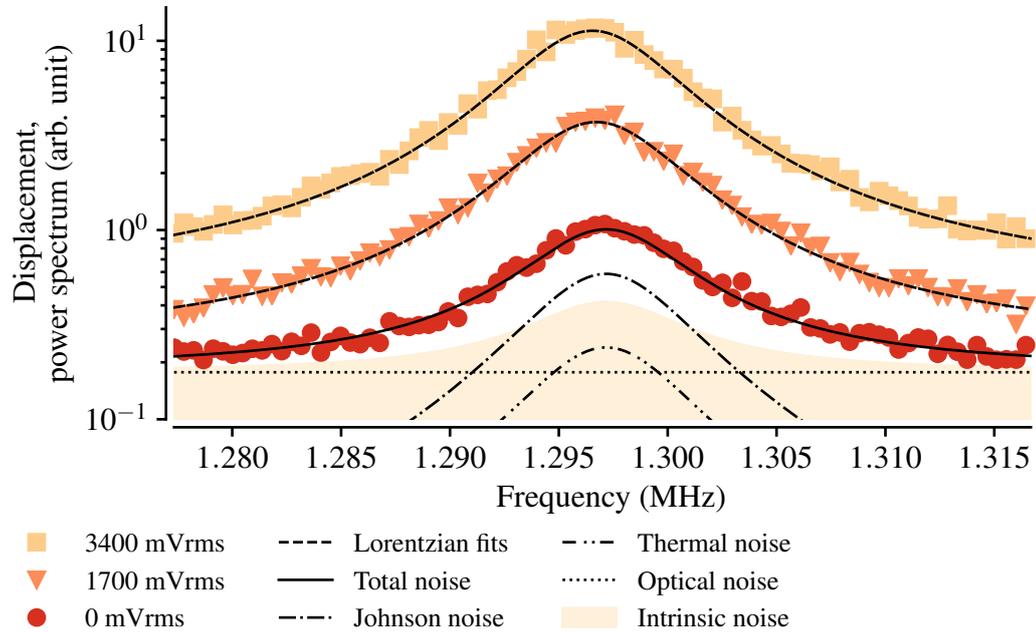
Figure 6.5: Optically measured PSDs for membrane motion driven by different levels of electrical noise, along with Lorentzian fits to the data. The lowest data-set has only Johnson noise as a drive, and the curves below indicate the individual noise-contributions that add up to give the fitted line. (*Reprinted from Simonsen et al. (2019b).*)

Y-factor method described in section 3.2.1. Each data point came from weighted means of five sequential fits that used the fitted confidence interval as weights. Similarly, the error bars on the peaks were a weighted estimate of one standard deviation between the five fits.

6.2 Results

The model predicts a Q-factor of 91.4 for the new circuit resonance. This number agrees well with an independent measurement where a double-loop probe extracted Q to be 92, a number obtained by fitting to the recorded LC resonance linewidth. Note that the model indicates intrinsic Q of the detection coil to be around 370, but the final Q gets loaded by the filter. Without the transducer connected, the double-loop probe technique results in a Q of 119. Using these numbers with eq. (2.44) on page 27 gives the intrinsic Q of the transducer chip on its

own. The result is 405.

The measured linewidth with AC bias on was

$$\Gamma_{EM} = 8260(130)\,\text{Hz}$$

according to the peak-fits in fig. 6.5. That means the electromechanical cooperativity was (eq. (2.63) on page 33)

$$C = \Gamma_m/\Gamma_{om} - 1 = 33.0(8).$$

Note that this estimate refers to the mechanical linewidth broadened by optical backaction, which means the effective temperature must be used when estimating the noise temperature. Equation (2.70) on page 35 gives

$$T_{N,mech} = \frac{T_{om}}{C}\frac{\Omega_{LC}}{\Omega_m} = 112(3)\,\text{K} \tag{6.3}$$

for only the mechanical contribution to noise temperature.

Figure 6.6 presents the Y-factor extrapolation (fig. 3.9 on page 50) for the AC biased transduction. That yields the intrinsic noise-temperature to be 210(16) K for the membrane and optical noise combined, evaluated at the mechanical resonance frequency Ω_m. That is equivalent to a noise-figure of

$$NF = 2.33(14)\,\text{dB}.$$

We can distinguish the mechanical and optical noise from the Lorentzian fit eq. (2.42), ascribing the peak and offset to the mechanical and optical noise, respectively. That gives the mechanical contribution of 120(6) K and the optical contribution 90(15) K. Note that the two noises-levels are not equal. That is because the operating bias was not tuned exactly to the optimal cooperativity. Consequently, the signal-to-noise bandwidth is slightly different from eq. (2.71) on page 36. We found it to be

$$BW = 12.3(7)\,\text{kHz}.$$

The circuit model above allows conversion between this noise-temperature and voltage- or current noise. Concretely, the circuit fit calculated the Johnson at room temperature (296 K) to be

$$V_J = 117\,\text{pV}/\sqrt{\text{Hz}}.$$
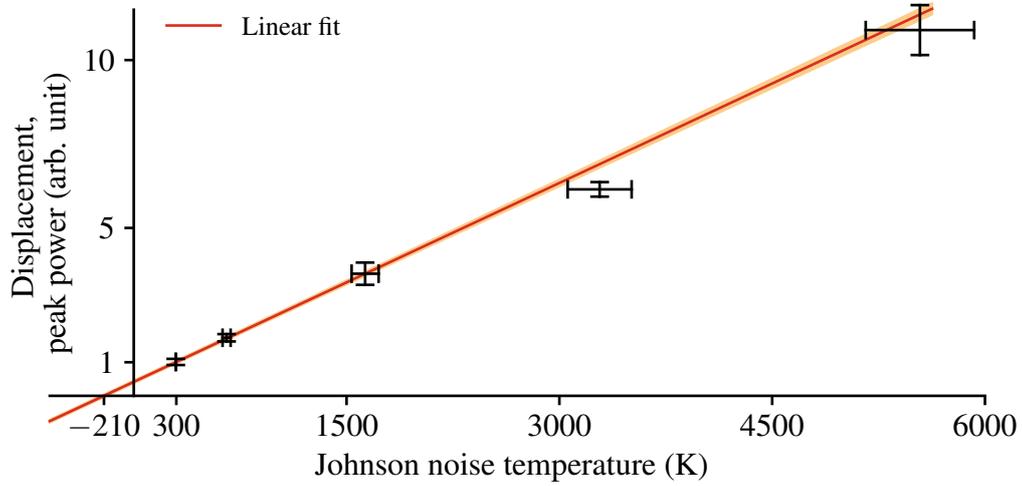
Figure 6.6: Noise-power after transduction, evaluated at the peak of the membrane and plotted against the voltage-drive converted into the temperature that gives an equivalent Johnson noise. The errorbars present one-standard-deviation, calculated from the variance in the voltage-drive and the variance between fits to five spectra. (*Adapted from Simonsen et al. (2019b)*).

Therefore, the intrinsic transducer noise in voltage units must be:

$$V_N = 117 \, \text{pV}/\sqrt{\text{Hz}} \sqrt{\frac{210(16) \, \text{K}}{296.15 \, \text{K}}} = 99(4) \, \text{pV}/\sqrt{\text{Hz}}. \tag{6.4}$$

The circuit model can also convert this number into an equivalent current-noise through the circuit impedance at the signal frequency,

$$|Z_{AC}(\Omega_{LC})| \approx 0.9 \, \Omega.$$

That results in

$$I_N = 113(4) \, \text{pA}/\sqrt{\text{Hz}}$$

.

Finally, the transduction noise also corresponds to a magnetic-field sensitivity, a sensitivity we estimated through the expression

$$B_\perp = L_d \, I_N / A_{\text{coil}} = 8 \, \text{fT}/\sqrt{\text{Hz}}. \tag{6.5}$$

Here, $B_\perp$ represents a uniform magnetic field that is perpendicular to the surface of the coil, and the coil-area A_{coil} was estimated to be $5 \times 10^{-3} \, \text{m}^2$, a number that we got by parametrizing and numerically integrating a picture of the coil.

6.3 Discussion

Theoretical and measured noise temperature agrees nearly within one standard. That means the setup has avoided technical noise issues such as ambient RF noise and sideband-noise from the bias drive, even though the first problem plagued our DC transduction and the latter was a significant challenge for Takeda et al. (2017) and Tominaga et al. (2018), and their experimental setup is very similar to ours.

Importantly, the SNR bandwidth meets the goal of 10 kHz set in section 4.1, unlike our first DC experiments, and the transduction uses AC biasing in a circuit design for NMR/MRI detection. This bandwidth is significantly larger than previous publications on the electro-mechano-optical transduction with a single-mode device (Bagci et al. 2014; Takeda et al. 2017; Tominaga et al. 2018), and it is almost as high as a multi-mode transducer Haghighi et al. 2018.

The corresponding noise-temperature, 210 K, is orders-of-magnitude lower than Haghighi et al. (2018), Takeda et al. (2017), and Tominaga et al. (2018), but unfortunately not as good as our DC transduction in chapter 3 and, nor the state-of-the-art electronic amplifiers discussed in section 3.3. For MRI at 32 MHz specifically, a high-performance, commercial amplifier (Watcom WMA32C) offers a noise-temperature 50 K. Nevertheless, we believe this result is auspicious considering this was our first full characterization using the new devices. Assuming further improvements, Takeda et al. (2017) have projected that the electro-mechano-optical transduction can perform better than standard amplifiers with suitable improvements. For our particular device, improvements would be a more sensitive optical setup and smaller membrane-capacitor gap.

Chapter 7

MRI imaging with the transduction

Disclaimer: parts of this chapter is adapted from a manuscript (Simonsen et al. 2019a) currently submitted for publication and under review. Some sentences and paragraphs may be very similar in language to that manuscript.

At long last, this chapter explains all the considerations needed to implement the transducer in an MRI scanner successfully. The work was done partially in parallel with the characterization in the previous chapter and used a similar setup for the initial tests at the scanner. The main difference was the vacuum chamber; it had to be compatible with the large magnetic field inside the scanner and was therefore made by glass-fiber and pumped by a molecular sieve. This chamber came from our collaborators from the Hypermag group at DTU. It was initially designed to be a cryostat cooled by liquid nitrogen. Once again, Juan Diego Sanchez represented our collaborator, and he he both operated the scanner took care of circuit construction.

The very first test at used the same circuit and scheme as the previous chapter (fig. 6.1b) and ended as well as could be expected for a first attempt: with membrane collapsed onto the electrode, meaning the transducer had been destroyed, with no apparent explanation. Good guesses, sure, but nothing truly known. This collapse happened despite successful test-runs with the setup in a shielded lab at DTU, even though Takeda et al. (2017) and Tominaga et al. (2018) had already demonstrated the transduction in an NMR setup. Our failure prompted more careful analysis, subsequent modifications, and even more failed attempts at operating the transducer in the scanner before, finally, it would work.

Section 7.2 below describes the setup modifications and final procedure that yielded an image using the transducer. First, however, the next section give a

brief crash-course in magnetic-resonance imaging that introduces the technical challenges associated with the scanner. The results in section 7.3 finally present data showing an MRI image recorded through an electro-mechano-optical transducer, with a final summary and outlook in section 7.3.2.

7.1 MRI basics

MRI imaging is a mature technology, the details of which scanning are too complicated to be covered in sufficient detail here. Moreover, several introductory textbooks exist on the topic. This section builds on such a book, Dale et al. (2015, ch. 1 to 6), and merely aims to summarize some fundamental concepts in MRI—phantom, transmit pulse, proton image, NMR signal, and pulsing scheme—that are necessary to understand how we implemented the transduction in a medical MRI scanner.

In standard MRI, a sample goes into a large magnetic field, $\vec{B}$. This sample contains atoms with nuclear spin, and the magnetic field induces an energy-splitting between spin-up and spin-down in the direction of the field. The energy-splitting ΔE depends on both the gyromagnetic ratio γ of the particular nucleus and the amplitude of the main magnetic field. The dependence is:

$$\Delta E = \hbar\,\Omega_{\mathrm{NMR}} = \hbar\,\gamma\,|\vec{B}|, \tag{7.1}$$

where Ω_{NMR} corresponds to the Larmor-frequency for the precession of the spins. This work used a medical scanner (GE MR750) with a 3 T main field and targeted the ^{13}C-isotope. That means the Larmor-frequency was 32 MHz. The sample we imaged was a so-called phantom: a plastic container that holds a liquid of high purity and relatively high content of carbon. Such a phantom gives more consistent results than a living sample, making it easier to compare to other MRI technologies that target the same atoms.

7.1.1 MRI pulse sequence

Usually, an MRI sequence starts by entangling the two energy-levels with an RF pulse—the transmit pulse. It has to have the right orientation and can be very powerful. On a Bloch-sphere, that pulse flips the spin-vector up towards the equator. Each spin then proceeds to precess around the north-to-south axis on the sphere, an axis aligned along the main magnetic field $\vec{B}$. The spins have an associated magnetic moment, and thus generate a magnetic field oscillating

perpendicular to the main field. If this new field penetrates a loop of electrical wire—the detection coil—it induces an electromotive force. The resulting voltage/current constitutes the Nuclear Magnetic Resonance (NMR) signal. Typically, the coil forms a resonant circuit to enhance the induced NMR signal.

Detecting the spin-precession is not sufficient to generate an image. An additional spatial encoding is needed to separate the spins from each other. That space-information comes from additional magnetic-fields that changes the Larmor frequency locally. The gradients affect the spins differently if they are on during or between the transmit pulses; in the first case, a gradient field ensures only specific spins get excited by the transmit pulse. In the second case, the gradient shifts the Larmor-frequency or -phase, depending on the scheme. All the excited spins add up to give the NMR signal. The actual spin-density follows from a Fourier transform of the time-domain NMR-recording, where the magnetic-field gradient defines the Fourier-transform variable (Kumar et al. 1975). In our case, the detection resonance does not move in frequency. Consequently, the last gradient must ensure that the spin-precession ends near the transduction frequency. That is difficult to guarantee in practice, which makes it advantageous to have a large transduction bandwidth.

In a general MRI scan, the spins lose constructive interference over time as the spin-precession dephases, incoherent physical process occusr, and as the spins relax back into the ground-state. This problem is typically addressed through through a variety of pulsing schemes that generate a pulse-echo for later detection, with the exact scheme depending application. We used a pulsing scheme that produced not only an image but also a spectrum of the ^{13}C motion, although using relatively large voxels. For the treatise, the scheme is a black box. It suffices to say that the spin-flip angle was small (10° to 20° on the Bloch sphere), that the spin-coherence was longer than the detection time, and that the gradient fields encoded the position through selective excitation rather than phase-encoding.[1]

After MRI processing, the output is a matrix with spectra corresponding to spin oscillation. The matrix covers a selected slice of space, and each spectrum corresponds to all spins in a particular voxel in that slice.

---

[1]To the best of this authors knowledge.

7.1.2 Calibration

Besides running the scanner to detect carbon atoms, the calibration procedures for the scanner relied on ^{1}H atoms—also called protons. Hence, the scanner comes with a secondary coil more or less built into the machine, the proton coil, that detects hydrogen atoms. This coil had two usages: one, to image the protons reliably and thus locate the phantom plus the settings of the gradients that targets it for detection; and two, to optimize field-homogeneity by minimizing the linewidth of the proton peak by tuning the current in ancillary, shim coils dedicated for field-correction—a process called shimming. Both calibration procedures induce voltage-spikes in electronics inside the scanner than the standard scans. The biggest reason is that the shimming has to switch large currents, but also because the proton-detection typically use a larger pulse-power than the main scan. However, the calibrations are not strictly necessary.

7.2 Methods

While it would have been possible to reconstruct the MRI image from the transducer's output directly, that approach would require timing the signal-sampling to the MRI sequence, not to mention the correct reconstruction algorithm. Since the MRI machine already does that, we chose to feed the transduced signal into the scanner, just like the standard electrical detection scheme would. However, the transduction down-converts the NMR signal to the mechanical resonance frequency, and the MRI scanner expects the signal at the Larmor frequency. More specifically, the transduction shifts the NMR signal down by the AC bias frequency. Therefore, we can recover the original signal by up-converting the transduced output with the AC bias using a mixer. Unfortunately, that up-conversion reduce the SNR, but we deemed it an acceptable trade-off for the added convenience.

The MRI scanner itself stands inside a dedicated room designed to shield against electromagnetic noise. Things inside this rooms risk being pulled into the scanner by the main magnetic field and should, therefore, be as non-magnetic as possible. To avoid issues with noise or magnetic compatibility, we placed all components for the transduction processing outside the shielded room, except for the transducer and the optical circulator. Unfortunately, that meant having to run long cables between the in- and outside.

After the very first failed test, we reduced the transmit pulse power as much as

possible and carefully increased the pulse-power in subsequent measurements. Also, we skipped the calibration procedures to avoid the voltage-spikes induced during these steps. After several iterations of improvements, we were able to operate the scanner with transmit-pulses at maximum power for the particular imaging scheme. Moreover, the transducer could also survive the calibration routines, at least when the bias was off. The next subsections describe the necessary measures that yielded an image, presented for the circuit, optics, and general setup, respectively.

7.2.1 Circuit

While the detection circuit is similar to the previous chapter (fig. 6.1a on page 107, one modification turned out to be essential: to detune the *LC* resonance frequency during the transmit pulse delivered by the MRI scanner. Without sufficient detuning, each pulse induces large voltages in the pick-up coil as the transmit and detection frequencies are identical. Those voltages are large enough to pull the membrane electrostatically down into the second electrode. This detuning is a standard technique in MRI (Edelstein et al. 1986), and there are several ways to implement such circuit detuning. The one that worked for us was to cut the detection coil from before in two and insert a segmenting capacitor with a trap circuit in parallel (Rispoli et al. 2016). Figure 7.1 shows the final circuit diagram along with a photo of the circuit after modifications. The trap activates when a trigger voltage exceeds the forward voltage of a PIN-diode, a trigger delivered by the MRI machine during the transmit pulses.

The new circuit board also changed the bandpass filter at the RF bias frequency, as shown in fig. 7.1. The new filter had a narrower bandpass window compared to before. Though more importantly, the new filter also permitted a DC bias through to the transducer, with the new segmenting capacitor preventing a short of the DC bias through the coil. That was important to reduce spurious noise in the system, as explained below.

Another circuit modification was to put two antiparallel PIN-diodes in parallel to the transducer; these diodes prevent a large voltage across the transducer, thus offering the additional protection of the transducer. Conversely, they also limit the maximum AC bias applied to the transducer. Using crossed diodes like that is a standard way to protect electronic preamplifiers. However, when on their own, they did protect the transducer sufficiently during the transmit pulses.

The above protective measures go beyond what Takeda et al. (2017) and Tom-
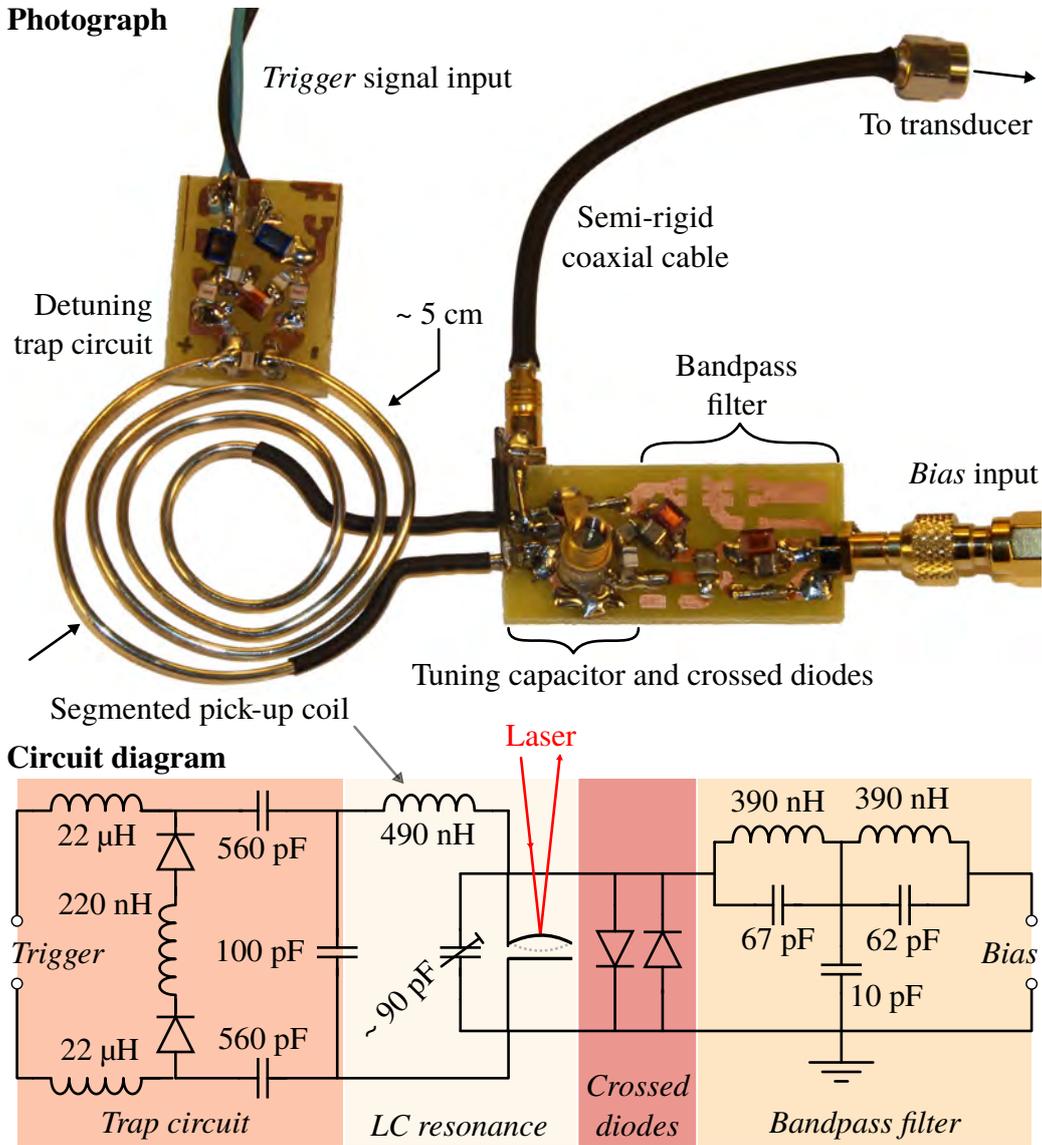
Photograph

Figure 7.1: Photograph (top) and circuit diagram (bottom) of the detection circuit after modifications. The photo has been edited to remove the background for clarity. (*Adapted from Simonsen et al. (2019a).*)

inaga et al. (2018) reported, and also adds to the transduction implementation in the previous chapter. Nonetheless, the whole setup remains very similar. Electrically, the circuit is essentially still an *LC* resonance tuned to the Larmor precession frequency of the ^{13}C nuclear spins and connected in parallel to the transducer. Although this connection looks slightly different as it now goes through a short coaxial cable—see fig. 7.1. Furthermore, the transducer still used the same 8-pin circuit socket and was still shorted by the switch introduced in section 4.3.1 to protect it during transport, mounting, and making electrical connections.

7.2.2   Optical readout

The optical setup also had to be modified before we could use it in the scanner. This alteration started simple—make the fibers longer between laser, transducer, and detector, so the laser and detector can be outside the shielded scanner room—but it ended up being more involved. To begin with, the additional 10 m of fiber went between the circulator and fiber-feedthrough in the previous fiber setup (fig. 5.2 on page 90). Both elements already terminated in a fiber-connection (angle-polished) so the added fiber was just a patch-cable connected to these ends.

However, this added fiber turned out to introduce noise in the optical readout. Over a measurement of a few seconds, the recorded spectra had their background noise-level jump significantly during the measurement. The amplitude and frequency of the jumps depended on the current driving the laser: the jumps could be higher than 10 dB and happen multiple times per second, or they could be as small as 5 dB to 7 dB and happen a little less than once per second. Additionally, both amplitude and frequency instability would drift in time, although it improved overall if the long fiber was polarization-maintaining. The current-dependence suggests mode-hops of the laser. This explanation appears likely considering it was a cheaper diode-laser without a build-in isolator, and that the specified isolation was inadequate for the components in the fiber network. Interestingly, these jumps were not visible without the extra fiber length, no matter the current settings. That could be because some interference has to occur in the fiber-network, interference that converts mode-hops in frequency into noise in amplitude, and that would be affected by the path-length difference.

The solution was to move the circulator closer to the transducer—i.e., into the scanner room, as shown in fig. 7.2—and have two long fibers for getting light to
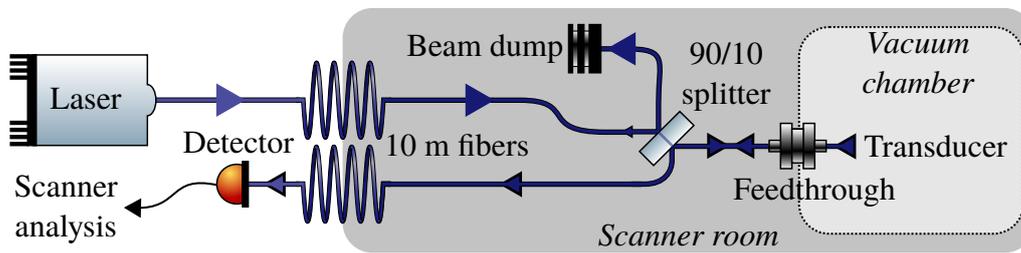
Figure 7.2: The optical network used at the MRI scanner. The setup is partially inside and outside the shielded room. All components are coupled directly to fiber, and everything inside the shielded room is as non-magnetic as possible, per usual.

and from the scanner. Unfortunately, a typical circulator cannot go near the high magnetic field as they usually rely on Faraday rotators with strong, build-in magnets to manifest non-reciprocal beam propagation. As an alternative, we used an optical 90/10 splitter that routed 10 % of light coming from the laser towards the transducer. The transducer's reflection went back into the same splitter, and 90 % of that proceeded towards the detector through a separate, polarization-maintaining fiber. The remaining light from the laser went into an optical terminator, and the remaining reflection went back towards the laser, now attenuated to <1 % of the lasers output power.

Besides these changes, the optical setup remained the same as in fig. 5.2 on page 90. It used the same transducer coupling, fiber-feedthrough, and single-mode fibers, along with the same diode laser that operated at the optical wavelength of 1064 nm.

7.2.3 Mounting in the scanner

Because the transducer needs a vacuum to operate, we placed it inside a custom cryostat, a cryostat comprised of glass-fiber and other (mostly) non-magnetic materials so that it could enter the scanner. Figure 7.3 shows a photo of this cryostat. Although it could be used to cool the transducer and circuit, our setup only employed it as a vacuum chamber that could hold a pressure $<1 \times 10^{-3}$ mbar up to five hours after disconnecting from its pump. The chamber maintained the vacuum-level using two molecular sieves (activated charcoal and sodium aluminum silicate), both cooled to liquid nitrogen temperature (77 K). The typical pressure was in the 1×10^{-5} mbar to 1×10^{-4} mbar range during measurements.

In the first iteration, both the transducer and circuit was inside the vacuum. However, that made it cumbersome and slow to exchange the transducer when
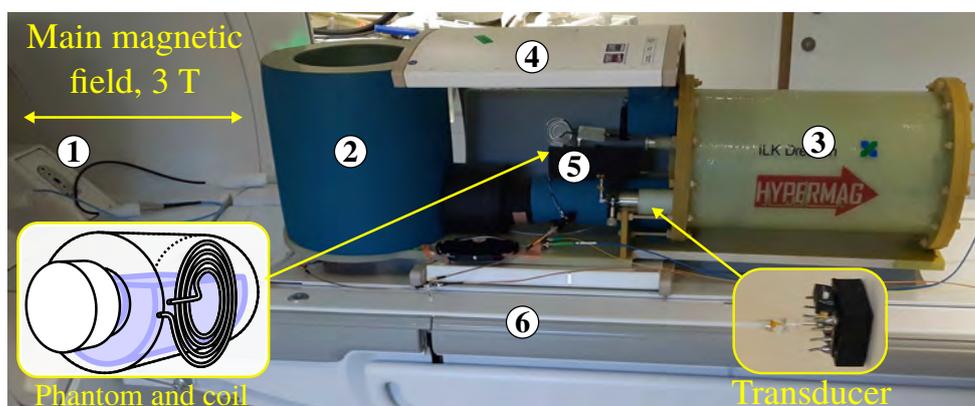
Figure 7.3: Big photo shows the custom cryostat before it enters the scanner bore (1). The liquid nitrogen container (2) and the vacuum chamber (3) goes on either side of the proton coil (4) such that the phantom and coil (5) are centered in the scanner when the patient bed (6) moves in. Left inset depicts the phantom and coil; right inset shows the packaged transducer.

necessary, not to mention mounting the phantom over the coil, a process that often required tuning the *LC* resonance slightly. In the final iteration, only the transducer went into the vacuum, loaded in through a KF25 flange on the back of the chamber. That meant the transducer-circuit connection had to go through a coaxial-cable-feedthrough in the vacuum flange. This connection included a coaxial-cable t-piece, and we used one of its connectors to short the transducer when disconnecting the circuit, a connector that was left open during the actual imaging.

Figure 7.3 shows the placement of the cryostat on the patient bed before we drove it into the center of the MRI scanner bore. The alignment puts the phantom and detection coil roughly near the center of the proton coil, and close to each other. In the end, we used a phantom that was simply a plastic bottle filled with ethylene glycol (purity 99.8 %; natural abundance, 1.1 %, ^{13}C). This particular liquid has an NMR spectrum with three peaks—a triplet—with a splitting of 142 Hz. That unique feature makes it easier to distinguish the actual NMR signal from a potential noise peak, at least compared to an NMR spectrum with only one peak.

Lastly, when mounting the transducer inside the cryostat, we tried to align the membrane perpendicular to the main magnetic field. Without this alignment, the mechanical linewidth would broaden. We attribute this effect to the Lorentz force on the charges in the membrane.

7.2.4 Connecting everything

Figure 7.4 presents a diagram of the signal-processing setup that fed the transduced NMR data back into the scanner. That shows how two long coaxial cables connected the setup inside and outside the scanner room. Inside, we connected the trigger signal directly to the detuning trap on the circuit, while also sending the trigger out through one of the long cables. That same cable simultaneously carried the transduced signal back into the scanner. These two signals are at different frequencies—the trigger and NMR signal being low- and high-frequency, respectively—so we could split them with a bias-tee at each end of the long cable. The second long cable delivered the AC bias to the transducer together with a DC offset. Outside the scanner room, the trigger signal controlled the AC bias such that it turned off during the transmit pulse, although with a slight time-delay (~100 ns) introduced by the long cable. Furthermore, the laser and detector were also outside the room, along with the up-converting mixer introduced above. After mixing and additional amplification, the transduced signal went back into the room, through the first long cable, and into the scanner's input.

As it turned out that, both the biasing- and mixing-tone had to be degenerate in frequency to obtain an image. When we tried separate generators, with an insufficient frequency lock, the image reconstruction failed to give the expected image. To implement that, we split the output of a single generator 50/50, with one part becoming the AC bias and the other becoming the local oscillator to the mixer. Before reaching the transduction circuit, the AC bias passed through these elements, in order: a separate amplification stage, the external filter (see fig. 6.1b
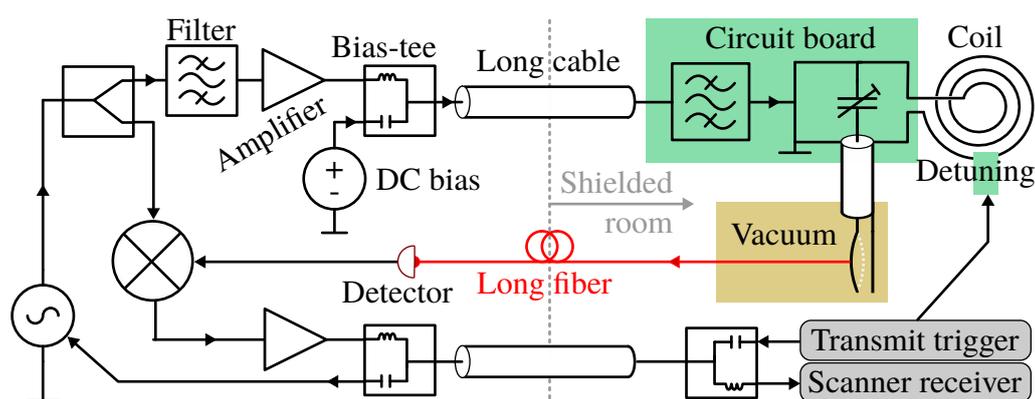
Figure 7.4: Diagram of the electric setup that interfaces the transducer and its output to the scanner hardware. (*Adapted from Simonsen et al. (2019a).*)

on page 107), a bias-tee that added a DC offset, and one of the long cables going to the scanner. In contrast, the other part went directly into the mixer. This scheme worked, despite the image-reconstruction being frequency sensitive and the bias collecting a phase-delay from the long cable compared to the local oscillator.

The plots in fig. 7.6 show the transduction after the up-conversion, but without an NMR signal. Without the AC bias, the spectrum shows the intrinsic membrane peak. With the bias, the peak moves down in frequency and broadens as expected from the electromechanical coupling. This spectrum displayed significant noise in our first experiments but went away with optimization, as also shown in fig. 7.6. As it turned out, this noise came in through the long cables. Exchanging them with semi-rigid coaxial cables, seemingly solved the problem. Probably because they have their outer conductor made out of solid metal, thus shielding the inner conductor from noise far better than the standard, flexible coaxial cables. Although this was an improvement, it was far from sufficient to remove all the noise. The most significant advancement came when we added a DC offset to the bias, an offset tuned to maximize the resonance frequency of the membrane. This DC voltage cancels out the trapped charges on the membrane discovered in section 4.3.2. Because these charges are static, the associated noise cannot have been frequency-shifted. We suspect it either comes from the ambient environment—as the room shielding probably is less efficient at the mechanical frequency—or from the trigger connected to the trap.

7.3   Results and discussion

Figure 7.5 shows the thermally driven motion of the membrane, with all the modification in place, and with the detected response up-converted to the NMR-signal frequency. Note that the actual membrane resonance was at $1.39\,\mathrm{MHz}$. Additionally, the plot includes Lorentzian peaks fitted to the data using eq. (2.42) on page 27. These show that mechanical linewidth broadened with bias from $181\,\mathrm{Hz}$ to $917\,\mathrm{Hz}$. According to equation eq. (2.63) on page 33, that broadening gives the electromechanical cooperativity $C_{\mathrm{em}} \approx 4.0$. Furthermore, they show that the peak area decreased by ≈ 4.8. This peak reduction comes straight from the measurements, ignoring the effect that the bias has on the displacement sensitivity (section 5.1). On the mechanical resonance, the thermal driven response is about four times higher than the optical background. That means this experiment, unfortunately, did not reach the optimal cooperativity. Even if it had, it
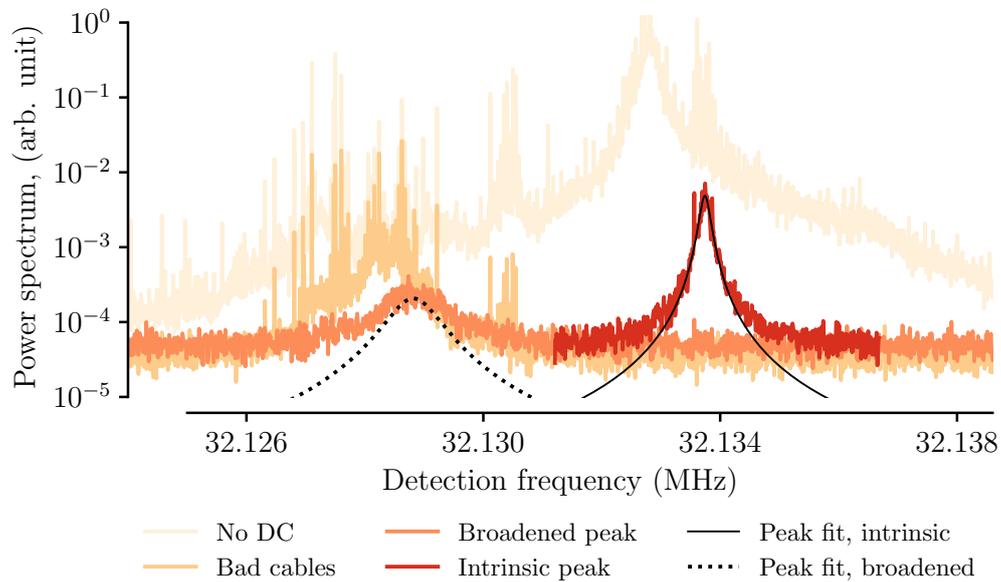
Figure 7.5: The figure shows the mechanical spectra after up-conversion and before it goes to the scanner. They come from different stages in the setup progression, specifically with the added DC bias and better cables. Throughout testing, we used these spectra to evaluate the various improvements.

still would not have worked as well as the setup in chapters 3 and 6. However, it is also the first demonstration with significant changes to the entire setup, so the performance is not unreasonable.

For the relevant electrical and mechanical frequency, the cooperativity should have been higher than $32\,\text{MHz}/1.4\,\text{MHz} \approx 23$ to make the thermal noise of the membrane equal to the Johnson noise in the circuit, according to eq. (2.70) on page 35. That means the mechanical noise dominates these measurements, and its peak-area should consequently reduce by cooperativity. However, it decreased slightly more than that, a discrepancy most likely caused by the effect the bias has on the cavity as described in section 5.1. Nevertheless, the correspondence indicates we managed to reduce all added noise significantly, possibly even eliminating it.

7.3.1   Spectrum

Figure 7.6 shows the NMR spectrum of ethylene glycol obtained with two different measurements: the transduction setup and a typical MRI setup that used a commercial coil with a standard electronic amplifier. Both spectra correspond to a single voxel in the MRI image and illustrate three close peaks—this is characteristic triplet expected from ethylene glycol.

When running the same measurement with the spin-flip pulse turned off, the three NMR peaks disappear in both measurement schemes, leaving only a noise background intrinsic to the setup also shown in fig. 7.6. For the standard electronic amplifier, the noise is flat in a broad, spectral window. In contrast, the transducer's noise fits well to the expected Lorentzian with an added offset. Furthermore, when changing the AC bias frequency, the Lorentzian moved in frequency relative to the NMR peaks, also following our expectations.

All told, these observations confirm that the transduction worked and indeed show the correct NMR spectrum. Note that both noise levels are similar despite the relatively poor performance of this implementation of the transduction.
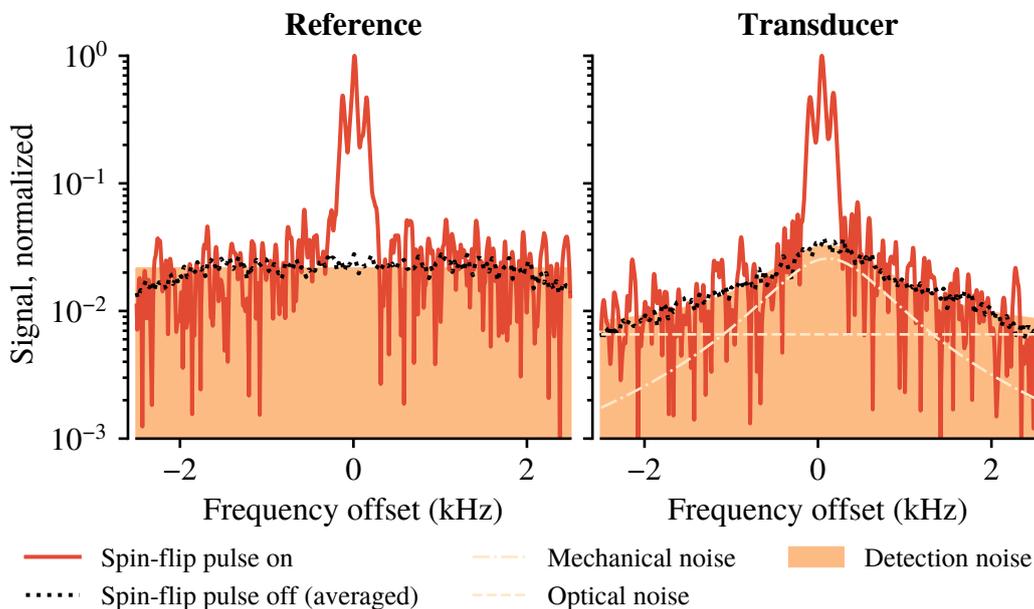
Figure 7.6: NMR spectra for a well-known reference (left) and the transducer (right). Both recorded and processed by the MRI scanner and normalized to the amplitude of the largest peak. (*Adapted from Simonsen et al. (2019a).*)

7.3.2 Image

The MRI scanner reconstructed the images shown in fig. 7.7 for both the commercial setup and the transducer. The images show the spatial density of ^{13}C atoms in a cross-section through the bottle that contains the ethylene glycol, a density which is uniform in reality. However, the detection itself has a spatial dependence. Specifically, the coil's sensitivity depends on the distance between the coil and voxel.[2] That dependence decays with said distance, determined by coil design. Both images used different coils and should, therefore, be expected to produce different images. As is the case in fig. 7.7. With its surface-coil design, the transducer coil has a sensitivity that decays much faster with distance compared to the commercial coil, thus explaining why the transducer-recorded image does not show as much of the phantom. Nevertheless, the measurement still reflects the same shape as the reference, although possibly with the phantom and coil misaligned between each MRI scan—it was difficult to mount everything identically for the two images.

   As far as we know, this data represents the first MRI image obtained through the electro-mechano-optical transduction!

---

[2]There can also be some variation from local perturbations in the magnetic fields and other typical image-degrading effects, but these should be small and, thus, neglectable.
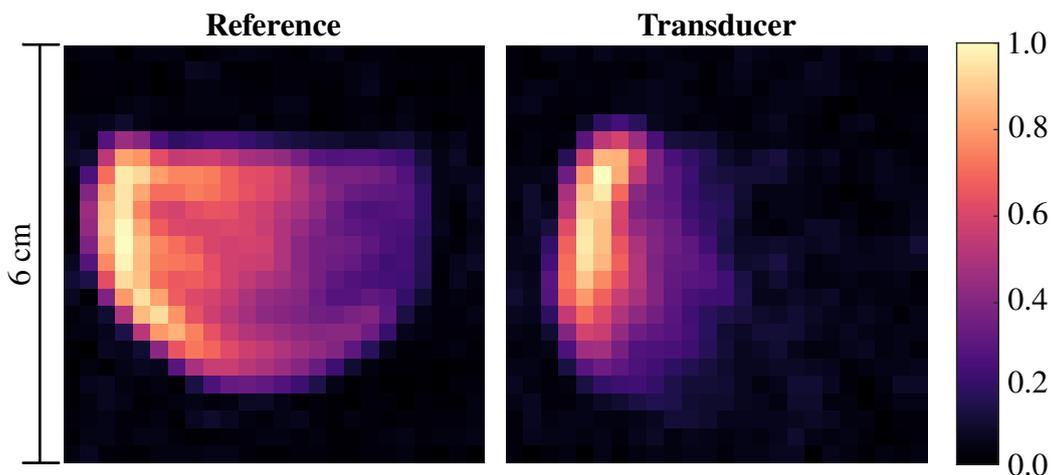
Figure 7.7: MRI image recorded with an established setup (left) and the transduction (right). Both collected and reconstructed by the MRI scanner itself through its standard data-processing routine. (*Adapted from Simonsen et al. (2019a).*)

Summary and outlook

An electro-mechanic-optical transduction scheme can serve as a sensor for faint electronic signals. We verified this through both theoretical analysis and an experimental investigation, an investigation that used either a DC or AC bias for transduction, a Y-factor extrapolation of intrinsic transducer noise-temperature, and ample shielding to reduce noise from the environment. The transduction platform was a capacitor with on plate being a freely suspended membrane, and it could all be integrated on a single chip together with an ultra-short cavity for optical readout. The compact device simplified the transduction setup considerably and enabled us to implement the transduction in a commercial MRI scanner, thus collecting the NMR signal optically and transferring it out of the scanner room via optical fiber. We believe this measurement is the first of its kind.

With a DC bias, we demonstrated an optimal noise-temperature 4 K and saw that it might extend down to 90 mK provided the optical setup gets better. These number´s are competitive with electronic amplifiers that are state-of-the-art and low-noise. However, the transduction has a much narrower bandwidth than conventional electronic solutions and the flip-chip membrane-capacitors yielded inconsistent performance. These inconsistencies that were due to the assembly, prompting an improved design and fabrication procedure. Nevertheless, the electromechanical system behaved as expected from theory.

Improvement on the membrane-capacitor device came by switching material from nitride to aluminum and employing an isotropic dry-etch process to release membranes. The membranes got deposited on top of a sacrificial layer on a mirror, and the thickness of that layer controls the length of an optical formed between mirror and membrane. The resulting optical cavity had a small, unexpected offset in its length but otherwise agreed will with a model for the optical propagation. Furthermore, the cavity could be coupled directly to a fiber with just one focusing lens and the lens-fiber distance being the only free parameter. However, the optical assembly indicated mediocre alignment of the focal point

133

and mode-matching between the fiber and cavity. That plus the unexpected cavity offset leaves room for considerable improvements on the optical detection. Despite these issues, the cavity enabled a sufficient displacement detection for AC biased transduction, and it demonstrated optical backaction from photothermal forces.

When implemented in a circuit designed for MRI and an AC bias, the new device had a noise-temperature of 210 K and an SNR bandwidth of 12 kHz. That is not as good as commercially available low-noise preamplifiers, but it is a significant improvement over similar works on electro-mechano-optical transduction aimed at room-temperature sensing, not to mention that the setup is less complicated and more compact than our first iteration. For the specific circuit, the noise-temperature may be expressed as $99\,\mathrm{pV}/\sqrt{\mathrm{Hz}}$, $113\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$, or $8\,\mathrm{fT}/\sqrt{\mathrm{Hz}}$. That does not mean the transduction had all this noise; each number is merely an equivalent representation of the total noise.

The transduction setup required additional modifications to work in an MRI scanner, most notably it needed protection from the transmit pulse and noise reduction. These precautions degraded the transduction performance, but it was nevertheless possible to see both an NMR spectrum and an MRI image from a phantom with ^{13}C, and the transduced NMR-spectrum showed a signal-to-noise comparable with a commercial system. Even though the transductions lacked in performance, this result still featured up-conversion of the NMR signal into an optical carrier and detection of this carrier outside the scanner room. Such remote sensing circumvents the technical issues with amplifiers and cables inside the MRI machine, issues that otherwise challenge MRI.

Future work on electro-mechano-optical transduction could address the technical issues described in this thesis. For the transduction circuit in MRI, that means increasing the AC bias amplitude and reducing the parallel capacitance. For MRI, other avenues for investigation are delivering the bias wirelessly like Menke et al. (2017)—this could be done by through the transmit pulse—and constructing an array of detection coils. For the membrane-capacitor design, potential advancement could be a lighter membrane, eliminating the trapped charges, and shortening the distance between membrane and electrode—a reduction by roughly a factor of three should be readily available with the existing design. Coupling to multiple mechanical modes like Haghighi et al. (2018) is an exciting approach for boosting the transduction bandwidth. For the optical cavity, possible improvement includes increased reflectivity of mirrors, for instance by moving to another wavelength of the laser, and better coupling to the cavity

or at least better control of the focal point on the laser. Boosting the input-mirror reflectivity is straightforward with the current fabrication, but that also make tolerance more stringent and therefore require better control of the cavity length. Another hindrance for transduction in MRI is the need for external vacuum. It might be possible to attain sufficient vacuum through micro-fabrication (Rushton et al. 2014), but that has not been considered in this work. So vacuum packaging remains a critical challenge to be solved before electro-mechano-optical transduction becomes a viable alternative to conventional electronics in MRI.

Bibliography

Abraham, E. R. I. and E. A. Cornell (1998). "Teflon feedthrough for coupling optical fibers into ultrahigh vacuum systems". In: *Applied Optics* 37.10, pp. 1762–1763.

Agarwal, G. S. and S. Huang (2010). "Electromagnetically induced transparency in mechanical effects of light". In: *Physical Review A - Atomic, Molecular, and Optical Physics* 81.4, pp. 4–7.

Aspelmeyer, M., T. J. Kippenberg, and F. Marquardt (2014). "Cavity optomechanics". In: *Reviews of Modern Physics* 86.4, pp. 1391–1452.

Bagci, T. (2014). "Opto-electromechanical devices for low-noise detection of radio waves". PhD thesis. University of Copenhagen, p. 159.

Bagci, T, A Simonsen, S Schmid, L. G. Villanueva, E Zeuthen, J Appel, J. M. Taylor, A Sørensen, K Usami, A Schliesser, and E. S. Polzik (2014). "Optical detection of radio waves through a nanomechanical transducer." In: *Nature* 507.7490, pp. 81–85.

Barg, A., Y. Tsaturyan, E. Belhage, W. H. P. Nielsen, C. B. Møller, and A. Schliesser (2018). "Measuring and imaging nanomechanical motion with laser light". In: *Applied Physics B* 123.1, pp. 71–85.

Braginskii, V. B. and A. B. Manukin (1967). "Ponderomotive effects of electromagnetic radiation". In: *Soviet Physics JETP* 25.4, p. 3.

Braginskii, V. B., A. B. Manukin, and M. Y. Tikhonov (1970). "Investigation of dissipative ponderomotive effects of electromagnetic radiation". In: *Soviet Physics JETP* 31.5, pp. 829–830.

Brittain and J. Brittain (1990). "Thevenin's theorem". In: *IEEE Spectrum* 27.3, p. 42.

Brooker, G (2008). *Modern Classical Optics*. 1st ed. Oxford University Press.

Carlson, G. T. and B. L. Illman (1994). "The circular disk parallel plate capacitor". In: *American Journal of Physics* 62.12, p. 1099.

Chowdhury, S, M Ahmadi, and W. C. Miller (2005). "A closed-form model for the pull-in voltage of electrostatically actuated cantilever beams". In: *J. Micromech. Microeng.* 15, pp. 756–763.

Cicak, K., M. Allman, J. Strong, K. Osborn, and R. Simmonds (2009). "Vacuum-gap capacitors for low-loss superconducting resonant circuits". In: *IEEE Transactions on Applied Superconductivity* 19.3, pp. 948–952.

Clerk, A. A., M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf (2010). "Introduction to quantum noise, measurement, and amplification". In: *Reviews of Modern Physics* 82.2, pp. 1155–1208.

Dale, B. M., M. A. Brown, R. C. Semelka, and M. A. Preceded by (work): Brown (2015). *MRI Basic Principles and Applications*. Chichester, UK: John Wiley & Sons, Ltd.

De Liberato, S., N. Lambert, and F. Nori (2011). "Quantum noise in photothermal cooling". In: *Physical Review A - Atomic, Molecular, and Optical Physics* 83.3, pp. 1–7.

Deng, C., M. Otto, and A. Lupascu (2013). "An analysis method for transmission measurements of superconducting resonators with applications to quantum-regime dielectric-loss measurements". In: *Journal of Applied Physics* 114.5, p. 054504.

DeVore, J. R. (1951). "Refractive Indices of Rutile and Sphalerite". In: *Journal of the Optical Society of America* 41.6, p. 416.

Dobrindt, J. M., I Wilson-Rae, and T. J. Kippenberg (2008). "Parametric Normal-Mode Splitting in Cavity Optomechanics". In: *Phys. Rev. Lett.* 101, p. 263602.

Dougherty, W. M., K. J. Bruland, J. L. Garbini, and J. A. Sidles (1996). "Detection of AC magnetic signals by parametric mode coupling in a mechanical oscillator". In: *Measurement Science and Technology* 7.12, pp. 1733–1739.

Edelstein, W., C. Hardy, and O. Mueller (1986). "Electronic decoupling of surface-coil receivers for NMR imaging and spectroscopy". In: *Journal of Magnetic Resonance* 67.1, pp. 156–161.

Filograno, M. L., M. Pisco, A. Catalano, E. Forte, M. Aiello, A. Soricelli, D. Davino, C. Visone, A. Cutolo, and A. Cusano (2016). "Triaxial fiber optic magnetic field sensor for MRI applications". In: *Sixth European Workshop on Optical Fibre Sensors* 9916.May 2016, 99160S.

Foreman-Mackey, D., D. W. Hogg, D. Lang, and J. Goodman (2013). "emcee: The MCMC Hammer". In: *Publications of the Astronomical Society of the Pacific* 125.925, pp. 306–312.

Frickey, D. A. (1994). "Conversion between S, Z, Y, h, ABCD and T parameters which are valid for complex source and load impedances". In: *IEEE Transactions on microwave theory and techniques* 42.2, pp. 205–211.

Gardiner, C. W. and P. Zoller (2004). *Quantum Noise*. Third. Berlin: Springer, p. 449.

Gerry, C. and P. L. Knight (2005). *Introductory Quantum Optics*. Cambridge: Cambridge University Press, p. 316.

Gillespie, D. T. (1996). "The mathematics of Brownian motion and Johnson noise". In: *American Journal of Physics* 64.3, pp. 225–240.

Gröblacher, S, K Hammerer, M. R. Vanner, and M Aspelmeyer (2009). "Observation of strong coupling between a micromechanical resonator and an optical cavity field". In: *Nature* 460, pp. 724–727.

Haghighi, I. M., N. Malossi, R. Natali, G. Di Giuseppe, D. Vitali, I. Moaddel Haghighi, N. Malossi, R. Natali, G. Di Giuseppe, and D. Vitali (2018). "Sensitivity-bandwidth limit in a multimode optoelectromechanical transducer". In: *Physical Review Applied* 9.3, p. 034031.

Haus, H. A. (1984). *Waves and fields in optoelectronics*. Ed. by N. Holonyak Jr. New Jersey: Prentice-Hall, p. 415.

Higginbotham, A. P., P. S. Burns, M. D. Urmey, R. W. Peterson, N. S. Kampel, B. M. Brubaker, G. Smith, K. W. Lehnert, and C. A. Regal (2018). "Harnessing electro-optic correlations in an efficient mechanical converter". In: *Nature Physics* 14.10, pp. 1038–1042.

Horowitz, P and W Hill (1989). *The Art of Electronics*. 2nd ed. Cambridge University Press.

Hyun, S., W. L. Brown, and R. P. Vinci (2003). "Thickness and temperature dependence of stress relaxation in nanoscale aluminum films". In: *Applied Physics Letters* 83.21, pp. 4411–4413.

Jöckel, A., M. T. Rakher, M. Korppi, S. Camerer, D. Hunger, M. Mader, and P. Treutlein (2011). "Spectroscopy of mechanical dissipation in micro-mechanical membranes". In: *Applied Physics Letters* 99.14, p. 143109.

Kathiravan, S. and J. Kanakaraj (2013). "A review on potential issues and challenges in MR imaging." In: *The Scientific World Journal* 2013, p. 783715.

Kemiktarak, U., M. Metcalfe, M. Durand, and J. Lawall (2012). "Mechanically compliant grating reflectors for optomechanics". In: *Applied Physics Letters* 100.6, p. 061124.

Khalil, M. S., M. J. A. Stoutimore, F. C. Wellstood, and K. D. Osborn (2012). "An analysis method for asymmetric resonator transmission applied to superconducting devices". In: *Journal of Applied Physics* 111.5, p. 054510.

Kim, H. S., J. M. Kim, Y. S. Bang, E. S. Song, C. H. Ji, and Y. K. Kim (2012). "Fabrication of a vertical sidewall using double-sided anisotropic etching of <100> oriented silicon". In: *Journal of Micromechanics and Microengineering* 22.9, p. 095014.

Kumar, A., D. Welti, and R. R. Ernst (1975). "NMR Fourier zeugmatography". In: *Journal of Magnetic Resonance (1969)* 18.1, pp. 69–83.

Landau, L. D. and E. M. Lifshitz (1986). *Course of theoretical physics, theory of elasticity volume 7*. Third. Butterworth-Heinemann.

Leong, K., J. Mazierska, and J. Krupka (1997). "Measurements of unloaded Q-factor of transmission mode dielectric resonators". In: *IEEE MTT-S International Microwave Symposium Digest*. Vol. 3. IEEE, pp. 1639–1642.

Malitson, I. H. (1965). "Interspecimen Comparison of the Refractive Index of Fused Silica*,†". In: *Journal of the Optical Society of America* 55.10, p. 1205.

Mao, P. and J. Han (2005). "Fabrication and characterization of 20 nm planar nanofluidic channels by glass-glass and glass-silicon bonding". In: *Lab on a Chip* 5.8, pp. 837–844.

McCulloch, M. A., J. Grahn, S. J. Melhuish, P.-A. Nilsson, L. Piccirillo, J. Schleeh, and N. Wadefalk (2017). "Dependence of noise temperature on physical temperature for cryogenic low-noise amplifiers". In: *Journal of Astronomical Telescopes, Instruments, and Systems* 3.1, p. 014003.

McPeak, K. M., S. V. Jayanti, S. J. P. Kress, S. Meyer, S. Iotti, A. Rossinelli, and D. J. Norris (2015). "Plasmonic films can easily be better: Rules and recipes". In: *ACS Photonics* 2.3, pp. 326–333.

McRae, T. G., K. H. Lee, G. I. Harris, J. Knittel, and W. P. Bowen (2010). "Cavity optoelectromechanical system combining strong electrical actuation with ultrasensitive transduction". In: *Physical Review A* 82.2, p. 023825.

Memis, O. G., Y. Eryaman, O. Aytur, and E. Atalar (2008). "Miniaturized Fiber-Optic Transmission System for MRI Signals". In: *Magnetic Resonance in Medicind* 59, pp. 165–173.

Menke, T., P. S. Burns, A. P. Higginbotham, N. S. Kampel, R. W. Peterson, K. Cicak, R. W. Simmonds, C. A. Regal, and K. W. Lehnert (2017). "Reconfigurable re-entrant cavity for wireless coupling to an electro-optomechanical device". In: *Review of Scientific Instruments* 88.9, p. 094701.

Metzger, C., I. Favero, A. Ortlieb, and K. Karrai (2008a). "Optical self cooling of a deformable Fabry-Perot cavity in the classical limit". In: *Physical Review B - Condensed Matter and Materials Physics* 78.3, pp. 1–12.

Metzger, C., M. Ludwig, C. Neuenhahn, A. Ortlieb, I. Favero, K. Karrai, and F. Marquardt (2008b). "Self-Induced oscillations in an optomechanical system driven by bolometric backaction". In: *Physical Review Letters* 101.13, pp. 26–29.

Metzger, C. H. and K. Karrai (2004). "Cavity cooling of a microlever". In: *Nature* 432.7020, pp. 1002–1005.

Møller, C. B., R. A. Thomas, G. Vasilakis, E. Zeuthen, Y. Tsaturyan, M. Balabas, K. Jensen, A. Schliesser, K. Hammerer, and E. S. Polzik (2017). "Quantum back-Action-evading measurement of motion in a negative mass reference frame". In: *Nature* 547.7662, pp. 191–195.

Nielsen, W. H. P., Y. Tsaturyan, C. B. Møller, E. S. Polzik, and A. Schliesser (2017). "Multimode optomechanical system in the quantum regime". In: *Proceedings of the National Academy of Sciences of the United States of America* 114.1, pp. 62–66.

Petersan, P. J. and S. M. Anlage (1998). "Measurement of resonant frequency and quality factor of microwave resonators: Comparison of methods". In: *Journal of Applied Physics* 84.6, pp. 3392–3402.

Pinard, M. and A. Dantan (2008). "Quantum limits of photothermal and radiation pressure cooling of a movable mirror". In: *New Journal of Physics* 10.10, p. 095012.

Pinard, M, Y Hadjar, and A Heidmann (1999). "Effective mass in quantum effects of radiation pressure". In: *European Physical Journal D* 7, pp. 107–116.

Polzik, E. S., A. Schliesser, S. Schmid, A. S. Sørensen, J. M. Taylor, K. Usami, T. Bagci, A. Simonsen, L. G. Villanueva, E. Zeuthen, J Appel, E. S. Polzik, A. Schliesser, S. Schmid, A. S. Sørensen, J. M. Taylor, K. Usami, T. Bagci, A. Simonsen, L. G. Villanueva, E. Zeuthen, and J Appel (2014). *Optical detector and amplifier for RF-detection having a position dependent capacitor with a displaceable membrane*.

Pozar, D. M. (2012). *Microwave Engineering*. 4th ed. John Wiley and Sons, Inc.

Restrepo, J., J. Gabelli, C. Ciuti, and I. Favero (2011). "Classical and quantum theory of photothermal cavity cooling of a mechanical oscillator". In: *Comptes Rendus Physique* 12.9-10, pp. 860–870.

Riley, K. F., M. P. Hobson, and S. J. Bence (2006). *Mathematical methods for physics and engineering*. Third. Cambridge: Cambridge University Press.

Rispoli, J. V., I. E. Dimitrov, S. Cheshkov, C. Malloy, S. M. Wright, and M. P. McDougall (2016). "Trap Design and Construction for High-Power Multinuclear Magnetic Resonance Experiments." In: *Concepts in magnetic resonance. Part B, Magnetic resonance engineering* 46B.4, pp. 162–168.

Rugar, D., C. S. Yannoni, and J. A. Sidles (1992). "Mechanical detection of magnetic resonance". In: *Nature* 360.6404, pp. 563–566.

Rushton, J. A., M. Aldous, and M. D. Himsworth (2014). "Contributed Review: The feasibility of a fully miniaturized magneto-optical trap for portable ultracold quantum technology". In: *Review of Scientific Instruments* 85.12, p. 121501.

Safavi-Naeini, A. H., D. Van Thourhout, R. Baets, and R. Van Laer (2019). "Controlling phonons and photons at the wavelength scale: integrated photonics meets integrated phononics". In: *Optica* 6.2, p. 213.

Sahoo, H. K., L. Ottaviano, Y. Zheng, O. Hansen, and K. Yvind (2018). "Low temperature bonding of heterogeneous materials using Al2O3 as an intermediate layer". In: *Journal of Vacuum Science & Technology B, Nanotechnology and Microelectronics: Materials, Processing, Measurement, and Phenomena* 36.1, p. 011202.

Saulson, P. (1990). "Thermal noise in mechanical experiments". In: *Physical Review D* 42.8, pp. 2437–2445.

Schmid, S, T Bagci, E Zeuthen, J. M. Taylor, P. K. Herring, M. C. Cassidy, C. M. Marcus, L. G. Villanueva, B Amato, A Boisen, Y. C. Shin, J Kong, A. S. Sœrensen, K Usami, and E. S. Polzik (2013). "Graphene on silicon nitride for optoelectromechanical micromembrane resonators". In: *arxiv*, p. 22.

Schmid, S., T. Bagci, E. Zeuthen, J. M. Taylor, P. K. Herring, M. C. Cassidy, C. M. Marcus, L. Guillermo Villanueva, B. Amato, A. Boisen, Y. C. Shin, J. Kong, A. S. Sørensen, K. Usami, and E. S. Polzik (2014). "Single-layer graphene on silicon nitride micromembrane resonators". In: *Journal of Applied Physics* 115.5, p. 054513.

Sidles, J. A. and D. Rugar (1993). "Signal-to-noise ratios in inductive and mechanical detection of magnetic resonance". In: *Physical Review Letters* 70.22, pp. 3506–3509.

Simonsen, A. (2014). "Opto-electro-mechanics with nanomembranes". PhD thesis. University of Copenhagen, p. 154.

Simonsen, A., J. D. Sanchez, S. A. Saarinen, J. H. Ardenkjær-Larsen, A. Schliesser, and E. S. Polzik (2019a). "Magnetic resonance imaging with optical preamplification and detection". In: *arxiv*, pp. 1–6.

Simonsen, A., S. A. Saarinen, J. D. Sanchez, J. H. Ardenkjær-Larsen, A. Schliesser, and E. S. Polzik (2019b). "Sensitive optomechanical transduction of electric and magnetic signals to the optical domain". In: *Optics Express* 27.13, pp. 18561–18578.

Sobol, W. T. (2012). "Recent advances in MRI technology: Implications for image quality and patient safety". In: *Saudi Journal of Ophthalmology* 26.4, pp. 393–399.

Solgaard, O., A. A. Godil, R. T. Howe, L. P. Lee, Y. A. Peter, and H. Zappe (2014). "Optical MEMS: From micromirrors to complex systems". In: *Journal of Microelectromechanical Systems* 23.3, pp. 517–538.

Sosale, G., S. Prabhakar, L. G. Frechette, and S. Vengallatore (2011). "A microcantilever platform for measuring internal friction in thin films using thermoelastic damping for calibration". In: *Journal of Microelectromechanical S.* 20.3, pp. 764–773.

Southworth, D. R., R. A. Barton, S. S. Verbridge, B. Ilic, A. D. Fefferman, H. G. Craighead, and J. M. Parpia (2009). "Stress and silicon nitride: A crack in the universal dissipation of glasses". In: *Physical Review Letters* 102.22, pp. 1–4.

Sridaran, S and S. A. Bhave (2011). "Electrostatic actuation of silicon optomechanical resonators". In: *Optics Express* 19.10, p. 7.

Stedman, R. and G. Nilsson (1966). "Dispersion Relations for Phonons at 80 and 300 ºK". In: *Physical Review* 145.2, pp. 492–500.

Taffoni, F., D. Formica, P. Saccomandi, G. Di Pino, and E. Schena (2013). "Optical fiber-based MR-compatible sensors for medical applications: an overview." In: *Sensors (Basel)* 13.10, pp. 14105–14120.

Takeda, K., K. Nagasaka, A. Noguchi, R. Yamazaki, Y. Nakamura, E. Iwase, J. M. Taylor, and K. Usami (2017). "Electro-mechano-optical detection of nuclear magnetic resonance". In: *Optica* 5.2, pp. 152–158.

Taylor, J. M., A. S. Sørensen, C. M. Marcus, and E. S. Polzik (2011). "Laser cooling and optical detection of excitations in a LC electrical circuit". In: *Physical Review Letters* 107.27, p. 273601.

Thompson, J. D., B. M. Zwickl, A. M. Jayich, F. Marquardt, S. M. Girvin, and J. G. E. Harris (2008). "Strong dispersive coupling of a high-finesse cavity to a micromechanical membrane." In: *Nature* 452.7183, pp. 72–5.

Tominaga, Y., K. Nagasaka, K. Usami, and K. Takeda (2018). "Studies on NMR-signal up-conversion from radio-frequency to optical regimes using a lightweight nanomembrane transducer". In: *Journal of Magnetic Resonance* 298, pp. 6–15.

Tsaturyan, Y., A. Barg, E. S. Polzik, and A. Schliesser (2017). "Ultracoherent nanomechanical resonators via soft clamping and dissipation dilution". In: *Nature Nanotechnology* 12.8, pp. 776–783.

Tsaturyan, Y., A. Barg, A. Simonsen, L. G. Villanueva, S. Schmid, A. Schliesser, and E. S. Polzik (2014). "Demonstration of suppressed phonon tunneling losses in phononic bandgap shielded membrane resonators for high-Q optomechanics." In: *Optics express* 22.6, pp. 6810–21.

Voorthuyzen, J. A. and P Bergveld (1984). "The influence of tensile forces on the deflection of circular diaphragms in pressure sensors". In: *Sensors and Actuators* 6, pp. 201 –213.

Weis, S., R. Rivière, S. Deléglise, E. Gavartin, O. Arcizet, A. Schliesser, and T. J. Kippenberg (2010). "Optomechanically induced transparency." In: *Science (New York, N.Y.)* 330.6010, pp. 1520–3.

Wiggins, G. C., J. R. Polimeni, A. Potthast, M. Schmitt, V. Alagappan, and L. L. Wald (2009). "96-Channel receive-only head coil for 3 Tesla: design optimization and evaluation." In: *Magnetic resonance in medicine* 62.3, pp. 754–62.

Williams, K. R., K. Gupta, and M. Wasilik (2003). "Etch Rates for Micromachining Processing - Part II". In: *Journal of Microelectromechanical Systems* 12.6, pp. 1689–1699.

Wilson, D., C. Regal, S. Papp, and H. Kimble (2009). "Cavity Optomechanics with Stoichiometric SiN Films". In: *Physical Review Letters* 103.20, p. 207204.

Wilson-Rae, I. (2008). "Intrinsic dissipation in nanomechanical resonators due to phonon tunneling". In: *Physical Review B* 77.24, p. 245418.

Winger, M, T. D. Blasius, T. P. Mayer Alegre, A. H. Safavi-Naeini, S Meenehan, J Cohen, S Stobbe, and O Painter (2011). "A chip-scale integrated cavity-electro-optomechanics platform". In: *Optics Express* 19.25, p. 17.

Wolf, K. (1979). *Integral transforms in science and engineering*. Boston MA: Springer US, p. 489.

Ylivaara, O. M. E., X. Liu, L. Kilpi, J. Lyytinen, D. Schneider, M. Laitinen, J. Julin, S. Ali, S. Sintonen, M. Berdova, E. Haimi, T. Sajavaara, H. Ronkainen, H. Lipsanen, J. Koskinen, S. P. Hannula, and R. L. Puurunen (2014). "Alu-

minum oxide from trimethylaluminum and water by atomic layer deposition: The temperature dependence of residual stress, elastic modulus, hardness and adhesion". In: *Thin Solid Films* 552, pp. 124–135.

Yu, P.-L., T. P. Purdy, and C. A. Regal (2012). "Control of material damping in high-Q membrane microresonators". In: *Physical Review Letters* 108.8, p. 083603.

Zalalutdinov, M., A. Zehnder, A. Olkhovets, S. Turner, L. Sekaric, B. Ilic, D. Czaplewski, J. M. Parpia, and H. G. Craighead (2001). "Autoparametric optical drive for micromechanical oscillators". In: *Applied Physics Letters* 79.5, pp. 695–697.

Zeuthen, E. (2015). "Electro-optomechanical transduction & quantum hard-sphere model for dissipative Rydberg-EIT media". PhD thesis. Copenhagen, p. 158.

Zeuthen, E., A. Schliesser, J. M. Taylor, and A. S. Sørensen (2017). "Electro-optomechanical equivalent circuits for quantum transduction". In: *Physical Review Applied* 10.1, p. 044036.

Zhang, W.-M., H. Yan, Z.-K. Peng, and G. Meng (2014). "Electrostatic pull-in instability in MEMS/NEMS: A review". In: *Sensors and Actuators A: Physical* 214, pp. 187–218.

Zhu, H., M. Holl, T. Ray, S. Bhushan, and D. R. Meldrum (2009). "Characterization of deep wet etching of fused silica glass for single cell and optical sensor deposition". In: *Journal of Micromechanics and Microengineering* 19.6.