A FULLY-IMPLANTABLE LOW-NOISE EMI-RESISTANT PIEZOELECTRIC-POLYMER MICROPHONE AND AMPLIFIER FOR THE MIDDLE EAR

by

AARON YEISER

B.S. Electrical Engineering and Computer Science
Massachusetts Institute of Technology, 2021

SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF ENGINEERING IN ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

Author .....
Department of Electrical Engineering and Computer Science
May 13, 2022

Certified by .....
Jeffrey H. Lang
Professor of Electrical Engineering
Thesis Supervisor

Accepted by .....
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

A Fully-Implantable Low-Noise EMI-Resistant Piezoelectric-Polymer Microphone and Amplifier for the Middle Ear

by

Aaron Yeiser

Submitted to the Department of Electrical Engineering and Computer Science
on May 13, 2022, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

We present a fully implantable piezoelectric microphone designed to operate with a cochlear implant. This thesis details the design, fabrication, and potential surgical implantation scheme for a fully differential and shielded cantilever made from polyvinylidene difluoride (PVDF)—a common piezoelectric polymer, as well as a low noise differential charge amplifier designed for small capacitance sensors. The amplifier and sensor combination has a noise floor of 385 e^- (0.062 fC) over its bandwidth of 100 Hz to 20 kHz, equivalent to 0.015 nm of displacement. When implanted, we achieve a pressure sensitivity of 80–100 fC/Pa referenced to ear canal pressure below 2 kHz and 8–10 fC/Pa above 4 kHz. We expect this sensitivity at high frequency to substantially improve when measured relative to free-field sound pressure, as the horn-like outer ear and ear canal provide up to 20 dB pressure gain above 1 kHz. Our design also provides significant EMI protection—we measured a sensitivity to external electric potentials of only 0.6 fF compared to over 200 fF for an unshielded 4 mm-diameter sphere. We believe this microphone design is competitive with commercial electret microphones used for cochlear implants, especially since our design is fully implantable and interfaces with the existing middle ear structures.

Thesis Supervisor: Jeffrey H. Lang
Title: Professor of Electrical Engineering

Acknowledgments

I would like to thank Professor Jeffrey Lang, Professor Heidi Nakajima, Professor Lisa Olson, and Professor John Kymissis for their leadership on this project and excellent guidance.

I would also like to thank Kurt Broderick for his help developing the sensor fabrication process and Lukas Graf for his work on the surgical support for the microphone.

Many thanks to John Zhang, Annesya Bannerjee, and Benjamin Cary for their foundational work and help running experiments.

Thanks to the people of MITERS for good company and esoteric technical expertise.

And of course, I am forever grateful to my friends and family for their love and support.

This project is funded by the NIH through grant R01 DC016874.

Contents

List of Figures

List of Tables

Chapter 1

Prior Work

1.1 Introduction

Cochlear implants are one of the most successful implants designed to restore sensory function, if not the most successful to date. A modern cochlear implant, shown in Figure 1-1, comprises an external portion containing microphones and sound-processing hardware that powers and communicates with an implanted electrode array inside the cochlea. The implant converts an audio signal into a series of electrical impulses in the cochlear implant electrodes in a manner that mimics the natural stimulation from the hair cells within the cochlea. Low frequencies map to electrodes near the apex of the cochlea and high frequencies to electrodes near the base. Thus the cochlear implant can restore hearing to individuals with damage to the middle ear or to the hair cells within the cochlea.

One of the major drawbacks of current cochlear implants is their external microphone. Human directional hearing relies on the shape of the outer ear—the frequency response of the outer ear is highly dependent on the azimuth and elevation of the incoming sound wave [17]. As cochlear implant microphones are located on the side of the head, these directional cues are not available to cochlear implant users. The outer ear and ear canal also act like a horn above 1 kHz, boosting sound pressure level at the ear drum.

External microphones pose additional problems for cochlear implant users. They

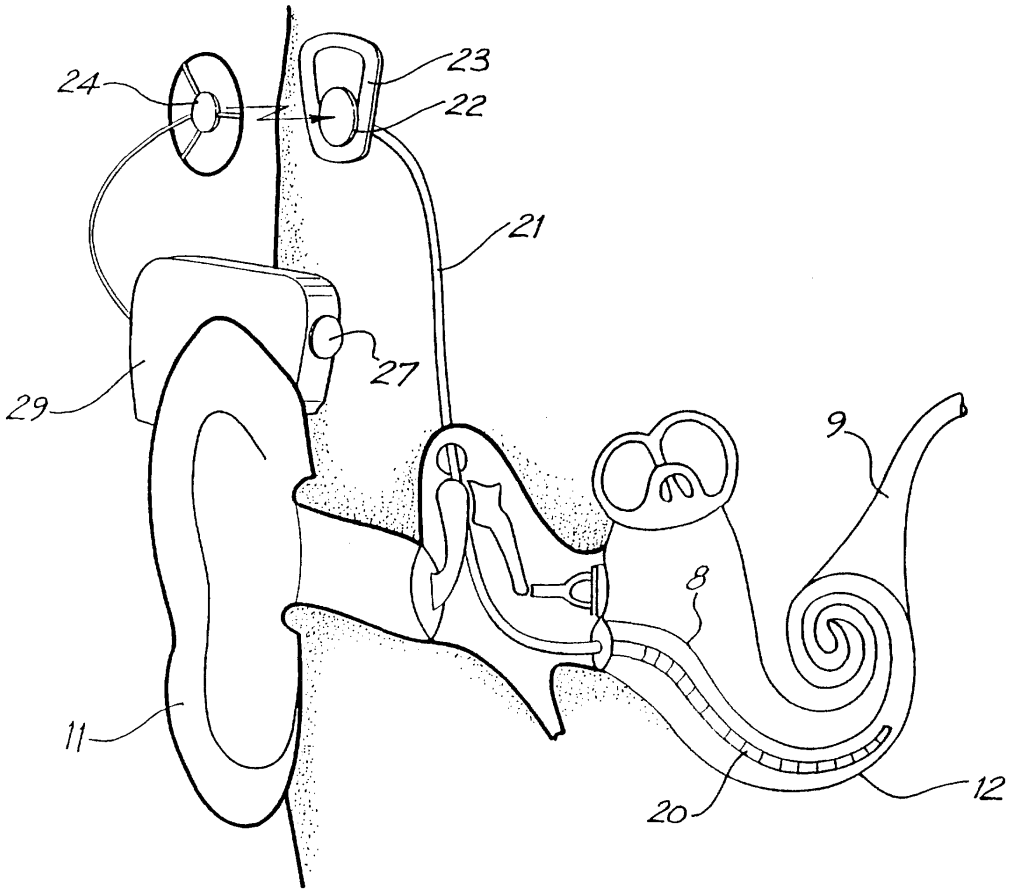

Figure 1-1: Figure 1 from US Patent No. 2003/0171787, showing a typical cochlear implant's microphone (27), sound processing electronics (29), wireless inductive transponder (22, 23, 24), and electrode array (20) inside the cochlea (12) near the basiliar membrane containing the auditory nerve (8).

are very visible, restrict the user's ability to wear hats or helmets, are susceptible to wind noise, and cannot be worn at night. There is active research to develop an implantable cochlear implant microphone that dates back at least two decades, and dozens of papers detailing different approaches to microphone design have been discussed.

This thesis explores the construction of a low-noise EMI-resistant piezoelectric microphone and amplifier, as well as some of the surgical considerations for installing said microphone. This microphone design takes the form of a differential-mode triangular cantilever and is heretofore formally denoted as the "cantilever-mic." The cantilever-mic addresses many of the issues of signal strength, EMI pickup, and fabrication

and implantation difficulties raised by previous designs. Like any good engineering project, the decisions that make the cantilever-mic a viable design build heavily on the features and antifeatures of rejected designs over the course of years of iteration and tinkering.

1.2 A brief technical review

Most modern hearing devices use external microphones, and for the past two decades, extensive research has taken place to try to develop implantable microphones. An excellent summary of current approaches to this problem can be found in [3]. Aside from subcutaneous microphones, middle ear devices usually fall into the categories of microphones, force transducers, and accelerometers. These device topologies primarily use piezoelectric or capacitive sensing, although piezoresistivity and optical sensing can be used as well. The location of device implantation varies fairly widely as well.

There are also devices designed to be implanted in the cochlea. The device described in [16] uses a MEMS microphone coupled to the fluid pressure in the scala tympani.

1.3 Why PVDF?

Our lab has focused on building implantable microphones that take advantage of the middle and outer ear structure. We build these devices out of polyvinylidene difluoride (PVDF)—a common piezoelectric polymer. PVDF has a few key advantages over typical piezoelectric materials, which are typically hard crystalline ceramics. The hardness of most piezoelectric materials presents problems when interacting with the auditory system, which tends to have moving parts with lower mechanical impedances. PVDF has a Young's modulus two orders of magnitude smaller than piezoelectric materials like lead zirconate-titanate (PZT), allowing for the much simpler fabrication of such structures. PZT and similar very efficient piezoelectric materials like potassium-

sodium niobate (KNN) also have a very high electrical permittivity. A PZT device would therefore have a much higher capacitance than a similar PVDF device. Generally speaking, op amps and FETs with very low noise figure tend to have a very high input noise impedance, and to achieve a good noise impedance match at audio frequencies, the sensor's capacitance must be very low. To achieve this low capacitance, PVDF seems to be the material of choice for implantable piezoelectric sensors.

The next few sections illustrate some of the PVDF–based microphones our lab has built. The three designs listed below all worked in the sense that when successfully implanted, they produced an electrical signal in response to an acoustic stimulus. The conclusion of this thesis is that these designs are an excellent starting point for more research, but they are not mature enough to stand on their own without substantial modification. A somewhat detailed description of these devices is warranted here to illustrate the (many) pitfalls of developing a usable implantable microphone.

1.3.1 Intracochlear-mic: lessons in impedance matching

This design, developed by John Zhang building on work done in [15] and [8], consists of a strip of PVDF embedded in silicone. The intracochlear microphone (abbreviated as the coch-mic) is implanted into the scala tympani through the round window, and the PVDF produces a voltage proportional to the fluid pressure. While this design is reasonably simple to implant and could be potentially integrated with a cochlear implant, it has an acoustic impedance several orders of magnitude higher than that of the inner ear. The extreme mismatch means that the vast majority of available acoustic energy in the scala tympani simply bypasses the device. This inefficiency is evidenced in the measured input equivalent noise floor of 50 to 60 dB SPL—about 20 dB worse than a hearing aid microphone [15].

The coch-mic also has serious problems with electromagnetic interference. In its current iteration, the coch-mic is single–ended and unshielded, thus making it an exceptional antenna for stray electric fields. This issue of pickup was severe enough that great care in ground loops and large amounts of shielding were required during testing to avoid electrical pickup from driving actuator, which was otherwise completely

swamping the desired signal. Inside the body, this device would have to contend with electrical noise from external sources, such as 60 Hz electrical noise and harmonics and switched–mode power supplies, as well as internal noise from random neural impulses and from the cochlear implant itself.

In my opinion, the combined issues of low sensitivity and high potential for electromagnetic pickup make the present coch-mic design impractical without substantial modification to address these problems. Attempts to refine the coch-mic, including a focus on acoustic impedance matching and a layered PVDF design, formed the foundation for the current cantilever microphone design.

1.3.2 Drum microphone: Why linearity matters

The drum microphone, which was the focus of Benjamin Cary's masters thesis [5], consists of a PVDF drum that rests under the umbo—the tip of the malleus that protrudes towards the middle-ear cavity from the center of the ear drum. With a small amount of preload between the umbo and the drum, movement of the umbo changes the tension in the drum, producing a charge across the drum's membrane.

This design has higher sensitivity than the coch-mic, but it has a few unique drawbacks. Firstly, the charge generated by the device varies roughly quadratically with displacement, and so does the mechanical impedance. This nonlinearity means that the device's mechanical impedance and sensitivity are dependent on the static offset of the drum. Although this nonlinearity does not cause substantial audio distortion, the static offset makes repeatable implantation challenging. In Ben's thesis, it is shown that changing the static offset by as little as 6 µm can double the sensitivity and presumably the mechanical impedance as well. In practice, the tolerance for placing the drum-microphone is 100–200 µm, since the drum-mic is much stiffer than the umbo connected to the ossicular chain. Secondly, the drum microphone rests on the promontory—a bony protrusion of variable proportions somewhat underneath the ear drum. If the shape of the promontory is wrong (which it often is), the drum microphone can rock back and forth and behave extremely inconsistently. Also, the Jacobsen's nerve runs over the promontory and is often dehiscant without

bone covering, further precluding implantation in this location.

The drum microphone also does not address the issue of electromagnetic pickup. Although it is significantly more sensitive than the coch-mic and it is not submerged in a conductive liquid, it still has no shielding and is still a single–ended device. Reimagining the drum design with a focus on linearity and EMI shielding led directly to the cantilever-mic design described by this thesis.

1.3.3 Diving board microphone: implantability issues

Another microphone design took the form of a relatively long cantilever mounted to the skull and resting on the umbo, similar to the drum microphone design. While the diving board behaves linearly as its mechanical impedance and sensitivity are mostly independent of offset, the diving board microphone proved difficult to implant. The anchor point for the cantilever had to be located relatively far from the umbo, making the stiffness of the diving board unacceptably low, especially at high frequencies. These factors conspired to make consistently implanting the diving board microphone exceptionally difficult.

The diving board functions by mounting a piece of PVDF to a non–piezoelectric backing material. When the device bends, a net tension or compression force is imparted to the PVDF, generating a net charge. This construction technique solves the problem of static offset, as it will readily generate charge with zero preload and tends to behave linearly. However, the mechanical impedance of the ear drum and ossicular chain is still moderately sensitive to static offset, and we struggled to control this factor.

The diving board was shelved and unpublished for a few years due to these implantation issues, but these difficulties seemed to originate from an excessively long cantilever length. A much shorter and thinner cantilever with a better solution for implantation could potentially perform quite well, with sensitivity comparable to the drum microphone but with none of the nonlinearity.

1.4 Lessons learned

These three designs demonstrated the feasibility of PVDF-based piezoelectric microphones. We decided that the drum's location under the umbo was best suited to intercepting incoming acoustic energy, and the diving board microphone's bending action held potential advantages over the drum-mic's membrane stretching action and the coch-mic's bulk compression action. A bending-mode device could be very linear and utilize PVDF efficiently, reducing parasitic capacitance.

None of the three designs addressed electrical shielding and EMI reduction in a satisfactory way. Even if electrical shielding of the sensor itself could be accomplished, there was no good way to attach the sensor to the amplifier. John's introduction of hot-bar bonding flex wires to the PVDF helped improve the reliability of the electrical connection but it did not provide any EMI shielding.

It was also clear that an improved low noise charge amplifier was possible. Improving the noise floor of the amplifier could improve the signal to noise ratio and hence system performance "for free" without modifying the sensor design.

Chapter 2

Sensor Design and Fabrication

2.1 The differential cantilever design

Around July or August 2021 we settled on a rounded triangular bimorph cantilever 3–4 mm wide and 3–4 mm long. The sensor detects motion of the umbo—the tip of the malleus that forms the point of the conical ear drum. Since umbo motion is large and mostly unidirectional in most humans, this seemed like a good location for an implantable microphone. Our sensor's particular shape was chosen to give high torsional stiffness, relatively uniform stress distribution in the PVDF, and a good mechanical impedance match to the ear drum and ossicular chain.

The sensor construction is illustrated in Figure 2-1, and pictures of the sensor assembly are shown in Figures 2-2–2-5. By building the sensor around a flexible printed circuit board (flex-PCB) substrate, we streamline the process of shielding the sensor and attaching cables without introducing significant parasitic capacitance. This ground shield completely enveloping the sensor combined with its symmetrical construction and differential output give extremely low sensitivity to external electromagnetic interference. We measured a differential capacitance to infinity of 0.6 fF—for comparison, a 4 mm-diameter sphere has a capacitance to infinity of over 200 fF.

The ground shield serves double duty as a ground electrode for the PVDF. On the inside of the sensor, the positive and negative sense electrodes are sputter coated

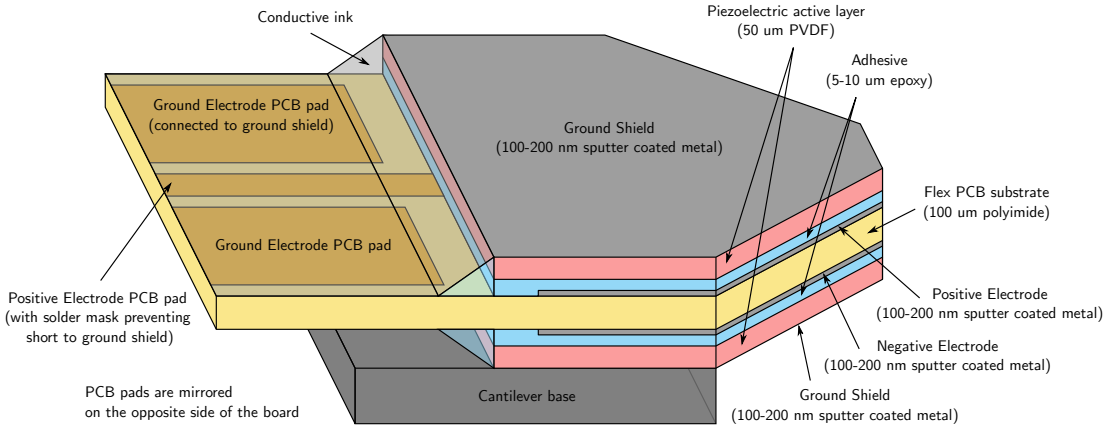

Figure 2-1: The cantilever-mic stackup, not to scale. The sensor is based on a flex-PCB substrate with sputtered aluminum electrodes. PVDF is attached with a thin layer of epoxy—this layer is thin enough to permit capacitive coupling. A sputter-coated ground shield encloses the whole sensor.

onto the flex-PCB substrate and directly connected to the flex-PCB copper pads. Although there is no direct electrical connection between the charge sense electrodes and the PVDF, most of the charge from the PVDF capacitively couples across the thin epoxy layer. This approach yields high mechanical strength and a reliable electrical connection to solderable flex-PCB pads.

50 µm poled PVDF was purchased from PiezoPVDF (formerly PolyK).

2.2 Fabrication process

Developing the fabrication process for these sensors was a months-long trial of gradual improvements and abject failure, described in detail in Appendix A. While the process described in this section could benefit from further refinement, it is mature enough to consistently produce sensors that survive the stress of implantation.

2.2.1 Electrode patterning

Figure 2-2 shows the progression from factory flex-PCB blank to flex-PCB with patterned electrodes. The flex-PCB, purchased from OSH Park with 4 mil polyimide and 1 oz copper, is the mechanical and electrical backbone for the cantilever sensor. It pro-

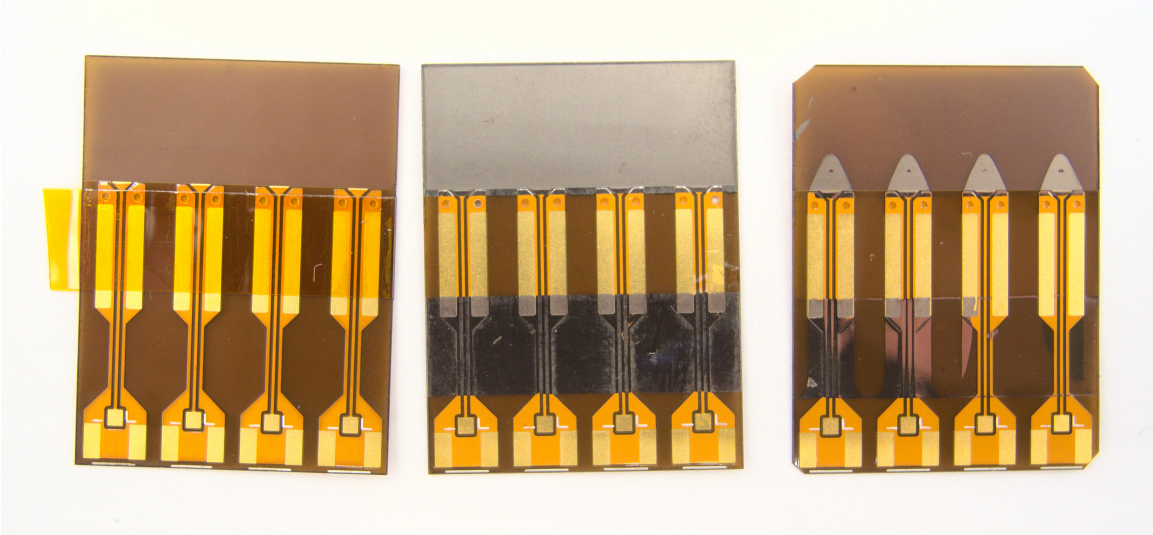

Figure 2-2: Left: a flex-PCB (20×26 mm) straight from the factory. Gold areas are ENIG gold-plated exposed copper, smooth orange is solder mask, and matte brown is base Kapton with no solder mask. Center: 200 nm aluminum or chromium electrode is sputtered onto either side of the flex-PCB blank. Right: A photolithography process is used to dissolve away the unwanted metal, leaving behind the shape of the charge-sense electrode.

vides the bridge between delicate sputter-coated electrodes and mechanically robust connectors, all while protecting the signal path from EMI. Our sensor design calls for an electrode on the flex-PCB directly beneath the PVDF layer. Although a copper flex-PCB pad would be a quick solution to this problem, the copper has too much mechanical stiffness, and we had to sputter coat own electrodes. Using MIT.nano's AJA ATC Orion sputter coater, we deposited approximately 200 nm chromium or aluminum, with chromium providing better step coverage between the flex-PCB trace

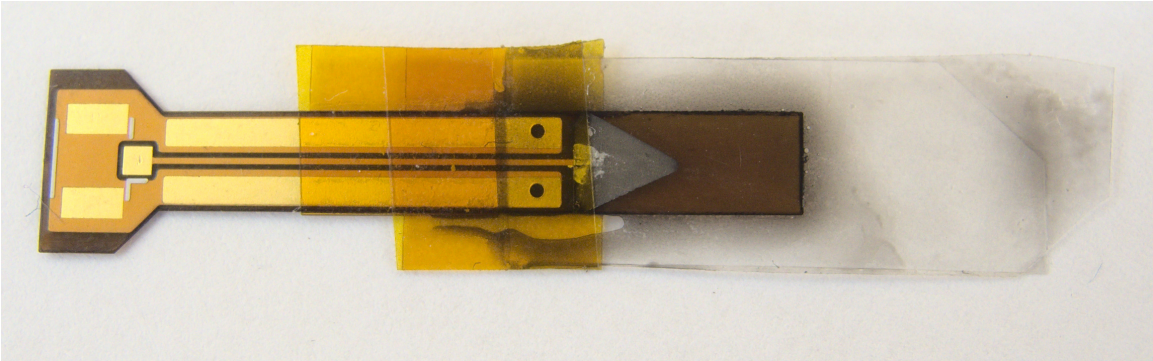

Figure 2-3: The sensor is masked off with Kapton tape and PVDF is glued on with epoxy. The masking process controls the location of the epoxy. The PVDF layers have opposite poling. We use a doctor blade to squeeze as much epoxy out from between the PVDF and the electrode as possible.

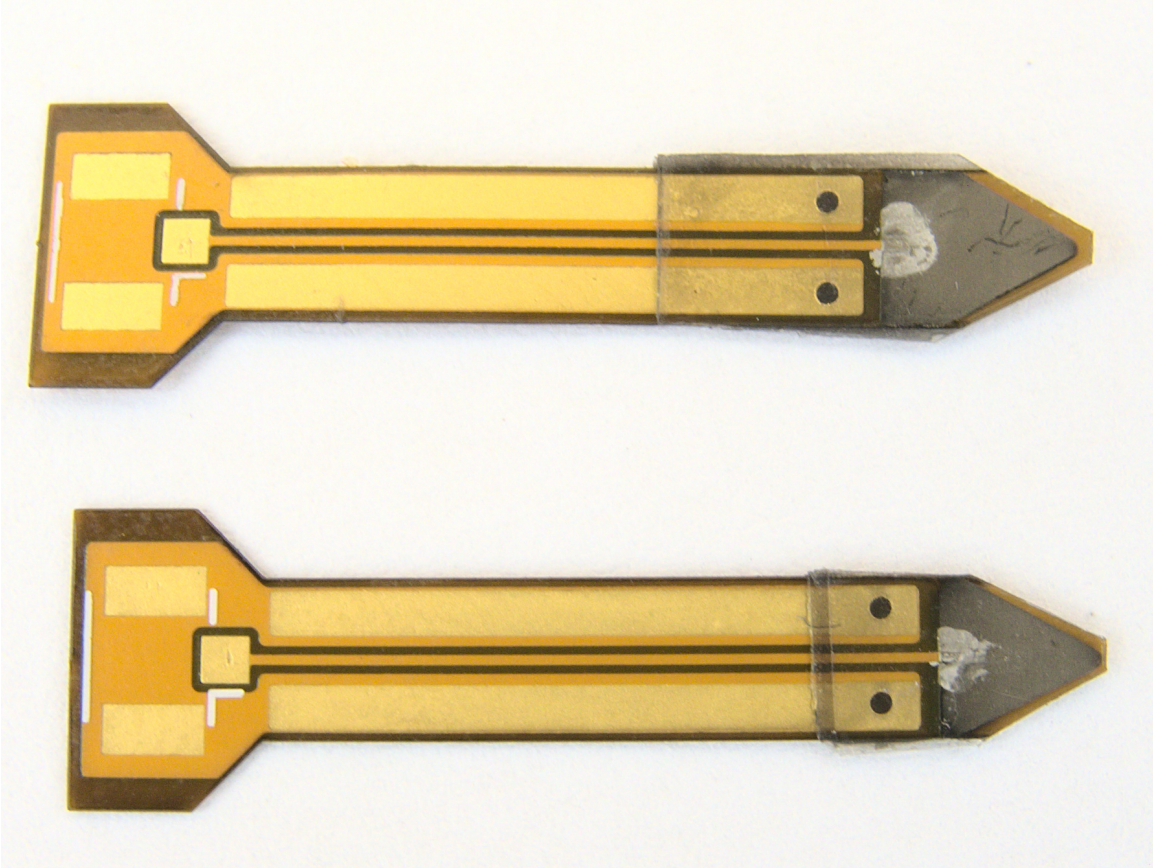

Figure 2-4: The sensor is trimmed down to shape with scissors. A small offset of electrode-free polyimide is left between the electrode and the edge of the sensor to provide waterproofing.

and the electrode layer and aluminum providing better electrical conductivity. We found that a one minute oxygen substrate etch was helpful for improving layer adhesion.

Waterproofing is achieved by ensuring an offset between the electrode and the edge of the sensor. The sensors therefore require the electrodes to be patterned before the PVDF is glued over them. At first we achieved this by masking areas of the flex-PCB with Kapton tape before sputtering, and then peeling off the tape. Unfortunately this procedure is time-consuming and imprecise. We therefore settled on a contact photolithography process.

Our photolithography process is as follows: The flex-PCB is spin-coated on both sides with AZ3312 positive photoresist. This is achieved by spin-coating one side of the flex-PCB for 30 seconds at 3000 RPM, baking it for approximately two minutes

Figure 2-5: 200 nm of aluminum or chromium is sputter-coated onto the surface of the PVDF. Conductive ink can be used to ensure electrical contact between the electrode and the flex-PCB trace. A U.FL connector is soldered onto the flex-PCB pads on the left.

at 110 °C, then spin-coating the other side and baking again. The photoresist-covered flex-PCB is placed into a contact lithography mask shown in Figure 2-6 and flood exposed for 30 seconds on each side. The UV-exposed photoresist degrades and can be dissolved in a tetramethyl ammonium hydroxide (TMAH) solution. The TMAH also dissolves aluminum, so no additional etching is needed for an aluminum electrode layer. For a chromium electrode layer, an acid etch in chromium etchant (nitric acid and ceric ammonium nitrate) is required. Finally the unexposed photoresist is dissolved in acetone and the electrode patterning process is complete.

2.2.2 PVDF layering

After significant effort spent on conductive adhesives and adhesive tapes, we settled on capacitively coupling our electrodes to the PVDF through a very thin non-conductive glue layer. This strategy came with its own challenges—the 50 µm PVDF has a unit capacitance of 2.2 pF/mm^2. For the capacitive coupling to be effective, the glue layer must have a substantially higher unit capacitance than the PVDF. Assuming the adhesive had a relative permittivity of at least 3.0, a glue layer less than 5 µm thick would have a sufficiently high unit capacitance of at least 5.3 pC/mm^2.

We tried a few different methods of applying very thin layers of adhesive before we found something that worked. Spin-coating and then curing SU-8 had poor gap

Figure 2-6: A 3D-printed jig to hold the flex-PCB blanks in place for photolithography flood exposure.

coverage and poor adhesion. We briefly experimented with spray adhesives but were unable to achieve consistent layer height or surface adhesion. Finally, we settled on using a Devcon Plastic Steel epoxy and repeatedly squeezing the two surfaces together with a doctor blade to push out as much excess glue as possible. Kapton tape was used to mask off areas of the sensor that needed to remain glue-free. This tape is shown in Figure 2-3, as well as a piece of tape used to tack down the left edge of the PVDF to the flex-PCB before gluing. This piece of tape acted as a hinge—the PVDF was lifted up, uncured epoxy globbed in, and then the doctor blade was used to squeeze epoxy out from the base of the cantilever to the tip. Figure 2-3 does not show Kapton over the U.FL pads, which is necessary since epoxy tends to get everywhere during this process.

We found that sanding the bottom of the PVDF with 3000-grit sandpaper was necessary to improve peel strength. Glue thickness of under 2.5 µm was measured with a micrometer; this was within the desired thickness range by a substantial margin.

2.2.3 Ground shield

After trimming the sensor to shape as shown in Figure 2-4, a final ground shield was sputtered over each side of the sensor. This ground shield forms the ground electrode

for the PVDF, and it also protects the charge sense electrodes from electromagnetic interference. By extending the ground shield up all the way to the U.FL connector, it effectively protects the entire signal path from EMI.

2.3 Pitfalls and potential process improvements

The process as described above is still highly developmental and in need of refinement. Quite possibly the biggest problem with the current process is its reliance on sputter coating step coverage. Covering 90° steps with a sputter-coated electrode is unreliable—the inner corners develop stress concentrations and the layer thickness of the step is concerningly thin. We experimented with conductive ink pens to electrically "reinforce" these steps. It is possible that a conductive epoxy would provide a stronger and more permanent solution. Eliminating sharp steps with filler material or design changes is another potential solution. If step coverage is not a concern, aluminum is more conductive and requires less processing (and nasty chemicals) than chromium.

The photolithography process in general was fairly experimental. Our process relied on evenly spin-coating both sides of the flex-PCB with positive photoresist, but the photoresist would often drip onto the back of the flex-PCB, causing unevenness. The photoresist also tended to accumulate at the step between the solder mask-free Kapton of the flex-PCB and the copper electrodes.

The 3D-printed jigs were a quick and dirty solution for a photolithography mask. While they were effective, they are unsuitable for more complicated patterns. A sufficiently high resolution contact lithography mask can be made by printing the design onto an overhead projector transparency with a decent quality toner printer, although this mask would have to be hand-aligned. Alternatively, a maskless contact lithography setup can be built for cheap using a UV lamp and a high DPI, high contrast LCD screen, similar to budget resin 3D printers.

There are several good reasons for replacing the pre-purchased flex-PCB with a pattern etched into a sputter-coated electrode on a Kapton blank. OSH Park was the

only fab that delivered adhesiveless flex PCBs for a reasonable price, and they only have one stackup available and a three week lead time. Replacing the pre-fabricated flex-PCB could help mitigate issues with sputter-coating step coverage on the first layer as well as problems with photoresist thickness in the lithography step. It also creates some new problems. Soldering to sputter-coated electrodes is extremely difficult and we would likely have to reinforce solder pads with an electroplated copper layer. We would also need our own insulating coating somewhat like the flex-PCB's solder mask to prevent short-circuiting the ground shield to the charge sense electrodes.

Chapter 3

Transducer Optimization

3.1 Introduction

Many different designs of PVDF piezoelectric transducers have been investigated to make an implantable microphone for cochlear implants. These devices have two main categories—umbo microphones and intracochlear microphones. Umbo microphones measure deflection of the umbo by placing some flexible element in contact with the umbo and detecting the resulting signal, where intracochlear microphones detect fluid pressure variations in the scala tympani.

We will briefly explore why we chose an an umbo-microphone design over an intracochlear microphone and then detail the design process for optimizing the shape and dimensions of the cantilever-mic for maximum mechanical performance. We also optimize the electrical interaction between the piezoelectric sensor and amplifier.

3.1.1 Piezoelectric mechanics

We can define the constitutive relations for a piezoelectric material:

$$\text{Strain tensor} \quad \mathbf{S} = \mathbf{s}\mathbf{T} + \mathfrak{d}^T\mathbf{E} \quad , \tag{3.1}$$

$$\text{Electric displacement field} \quad \mathbf{D} = \mathfrak{d}\mathbf{T} + \epsilon\mathbf{E} \quad , \tag{3.2}$$

with stress tensor $\mathbf{T}$, electric field $\mathbf{E}$, compliance tensor $\mathbf{s}$, electrical permittivity ϵ, and piezoelectric tensor $\mathfrak{d}$. By choosing an appropriate spatial reference frame, most of the terms of $\mathfrak{d}$ cancel to zero, with d_{33} and d_{31} most significant for PVDF. As PVDF is often sold in thin sheets, d_{33} represents the coupling between electric charge across the sheet and compression of the sheet across its thickness. d_{31} quantifies coupling between electric charge across the sheet and stretching along its length. For the purpose of building piezoelectric cantilevers, we mostly care about d_{31}.

We can also define the piezoelectric coupling coefficient k_{31}^2, which represents the strength of coupling between elastic energy and electrical energy in a piezoelectric material. For PVDF we can compute

$$k_{31}^2 = \frac{d_{31}^2}{s_{11}^E \epsilon_{33}} \approx \frac{2.7\,\text{GPa} \cdot (30\,\text{pC/N})^2}{12.5 \cdot 8.85 \times 10^{-12}\,\text{F/m}} = 0.022, \tag{3.3}$$

meaning that roughly 2.2 % of the elastic strain energy can be converted into electrical energy by a PVDF sensor.

3.2 Umbo vs intracochlear microphones

The performance of a piezoelectric microphone is closely related to the amount of acoustic energy that microphone can convert to electrical energy across the range of relevant frequencies. The ratio of electrical energy out to elastic strain energy in is bounded above by the piezoelectric coupling coefficient—about 2.2 % for PVDF. This available acoustic energy is roughly determined by the acoustic impedance of the structure connected to the microphone and the pressure at that location given a specific sound pressure level in the ear canal. For umbo microphones, we can assume that the ear drum moves pistonically at lower frequencies, giving an available acoustic power of

$$L_{umbo} = \frac{P_{EC}^2 A}{|Z|}, \tag{3.4}$$

where P_{EC}^2 is the RMS pressure in the ear canal, A is the area of the ear drum (roughly $50\,\text{mm}^2$), and Z is the acoustic impedance of the ossicular chain, measured

in $\mathrm{Pa} \cdot \mathrm{s/m}$. Umbo velocity is typically around $30\,\mu\mathrm{m}/(\mathrm{s} \cdot \mathrm{Pa})$ at 100 Hz and approximately $300\,\mu\mathrm{m}/(\mathrm{s} \cdot \mathrm{Pa})$ from 1 kHz to 4 kHz [14], giving

$$\frac{L_{umbo}}{P_{EC}^2} = 1.5 \text{ to } 15\,\mathrm{nW/Pa^2}. \tag{3.5}$$

Intracochlear microphone performance can be determined in a similar manner, with

$$L_{ST} = \frac{P_{ST}^2}{|Z|}, \tag{3.6}$$

P_{ST} being the RMS pressure in the scala tympani and Z being the parallel acoustic impedances across the round window and scala media, measured in $\mathrm{Pa} \cdot \mathrm{s/m^3}$ The pressure gain from the ear canal to the scala tympani is roughly 2 across a wide range of frequencies. The acoustic impedance is bounded above at 20 to $40\,\mathrm{GPa} \cdot \mathrm{s/m^3}$ by the differential impedance between the scala vestibuli and scala tympani, and bounded below at roughly $5\,\mathrm{GPa} \cdot \mathrm{s/m^3}$ by the acoustic impedance of the round window at roughly 1 kHz [11]. Combining all of these factors gives

$$\frac{L_{ST}}{P_{EC}^2} \approx 0.1 \text{ to } 0.8\,\mathrm{nW/Pa^2}, \tag{3.7}$$

although this equation only reflects energy available from hydrostatic pressure and ignores wave phenomena that become significant at high frequency. Purely on the basis of available acoustic energy, we should expect umbo microphones to outperform intracochlear microphones operating solely on hydrostatic pressure.

3.2.1 Acoustic impedance matching

The story does not end with available acoustic energy. This acoustic energy is useless unless a substantial fraction of the acoustic energy passes through the device, which requires a rough acoustic impedance match. As we will see later in this chapter, tuning the cantilever-mic's acoustic impedance is fairly straightforward, allowing for an impedance match to the umbo and attached ossicles by design. The drum-mic's mechanical impedance varies substantially with static offset and the drum is some-

what stiffer than would be ideal. The coch-mic's acoustic impedance to hydrostatic pressure is

$$Z = \frac{K}{j\omega V} \tag{3.8}$$

with K as the bulk modulus and V the volume, assuming the structure is solid and not subject to wave phenomena. With PVDF's bulk modulus of 3 GPa, a 1 mm × 15 mm × 50 µm intracochlear hydrophone has an acoustic impedance of

$$|Z| \approx 6.4 \times 10^{14}\,\mathrm{Pa \cdot s/m^3}, \quad f = 1\,\mathrm{kHz}, \tag{3.9}$$

approximately five orders of magnitude higher than the input impedance to the scala tympani at this frequency [11]. Since the coch-mic must fit through the round window, creative solutions to improve the acoustic impedance match to hydrostatic pressure are challenging.

3.3 Piezoelectric microphone optimization

There are several multiplicative factors that determine piezoelectric microphone efficiency. The microphone must be placed where a high amount of acoustic energy is available, and a rough acoustic impedance match is required to translate this acoustic energy into elastic strain energy in the piezoelectric material. The distribution and direction of this strain energy is important too—ideally strain is uniform through the PVDF and aligned with d_{31}, with the conversion of strain energy to electrical energy is bounded by k_{31}^2. To improve signal to noise ratio, the amplifier's noise figure should be as low as possible, and the noise impedance of the amplifier should be matched to the electrical impedance of the sensor.

3.3.1 Maximizing acoustic energy

Getting good performance out of a microphone generally requires as much acoustic energy flowing through the microphone as is feasible, given the design constraints. This means placing the microphone in an area that is as loud as possible and ensuring

that the microphone is appropriately coupled to its environment. The "loudness" of the environment can be approximated using the pressure and acoustic impedance of relevant structures, given that those structures are significantly smaller than one quarter wavelength. As stated above, tracking umbo motion is superior to measuring pressure in the scala tympani both in terms of available acoustic energy and ease of impedance matching.

3.3.2 Elastic strain efficiency

The direction and distribution of elastic strain is very important for a piezoelectric device. As much strain as possible should be aligned with the primary poling axis—d_{31} is much larger than d_{32}, so any strain perpendicular to the primary poling axis is effectively wasted.

Similarly important is the stress distribution. Assuming each piece of PVDF is flat and uniform with thickness h and area A, the charge accumulated across the sheet can be written as

$$Q = \frac{1}{h} \int d_{31} T_{11}(x, y, z)\, d\Omega = A T_{\text{avg}}, \tag{3.10}$$

where T_{avg} is the average T_{11} over the sheet. We also assume T_{11} is the only nonzero component of the stress tensor. The electrical energy stored in the sheet is

$$U_{\text{elec}} = \frac{Q^2}{2C} = \frac{1}{2\epsilon A h} \left(\int d_{31} T_{11}(x, y, z)\, d\Omega \right)^2 = \frac{d_{31}^2 T_{\text{avg}}^2 A h}{2\epsilon}. \tag{3.11}$$

Meanwhile, the mechanical energy stored in the device is

$$U_{\text{mech}} = \frac{1}{2C_{11}^E} \int T_{11}^2(x, y, z)\, d\Omega. \tag{3.12}$$

Chebyshev's inequality states that

$$\int_\Omega 1\, du \int_\Omega u^2\, du \geq \left(\int_\Omega u\, du \right)^2, \tag{3.13}$$

with equality when u is uniform. Therefore the ratio of electrical to mechanical energy is maximized when the stress is uniformly distributed through the piezoelectric material.

3.3.3 Noise figure and impedance matching

The previous subsections have covered heuristics for optimizing sensor performance purely by increasing electrical energy out, but this figure cannot be interpreted without a noise floor. For any given amplifier we can specify input–referred noise voltage v_n and current i_n. with correlation β. The noise voltage and current also induce

$$v_\mathrm{n,ind} = \frac{1}{j\omega C} i_\mathrm{n}, \qquad i_\mathrm{n,ind} = j\omega C v_\mathrm{n}, \tag{3.14}$$

where C is the sensor capacitance. We are interested in the overall noise power (including reactive power)—with $i_\mathrm{n,tot} = i_\mathrm{n} + i_\mathrm{n,ind}$, this gives us

$$\overline{P_\mathrm{n}} = \left\| \frac{1}{j\omega C} \overline{i_\mathrm{n,tot}^2} \right\| = \frac{1}{\omega C} \overline{i_\mathrm{n}^2} + \omega C \overline{v_\mathrm{n}^2} + 2\mathfrak{Im}(\beta) i_\mathrm{n,rms} v_\mathrm{n,rms}. \tag{3.15}$$

This noise power is minimized when a noise impedance match occurs:

$$\frac{1}{\omega C_\mathrm{ideal}} = \frac{v_\mathrm{n,rms}}{i_\mathrm{n,rms}} = Z_\mathrm{noise}. \tag{3.16}$$

In practice i_n, v_n, and β are frequency-dependent, and $(\omega C)^{-1}$ should roughly match Z_noise over the bandwidth of the device. In practice there is a range of acceptable sensor capacitances for any given amplifier and bandwidth. Choosing a larger sensor capacitance tends to result in more high-frequency and less low-frequency noise than a smaller-capacitance sensor. The shape of the noise floor is tunable to the sensitivity of the microphone.

Fortunately, the FETs (and FET-based op amps) that offer the absolute best noise figure at audio frequencies have a range of C_ideal overlapping with practical implantable PVDF piezoelectric sensors.

3.4 Optimization in practice: Mechanical modelling

Since the cantilever is thin relative to its length and has a variable cross-section, we can use the Euler-Bernoulli beam bending equation

$$M = -EI(x)\frac{d^2w}{dx^2} \tag{3.17}$$

where M is the bending moment, w the beam deflection, E the Young's modulus of the beam, and $I(x)$ the position-dependent second moment of area of the beam's cross-section. For a rectangular cross section with width b and thickness h, the second moment of area is

$$I = \frac{bh^3}{12}. \tag{3.18}$$

Our beam is a cantilever with a fixed base and a point load F at length L. We can represent these boundary conditions as

$$w(0) = 0, \qquad w'(0) = 0, \qquad M(x) = F(x - L). \tag{3.19}$$

It should be evident that for a cantilever of uniform thickness, the curvature is directly proportional to the stress in the piezoelectric layer. Going by the conclusion of Section 3.3.2 that stress should be evenly distributed through the piezoelectric material for maximum efficiency, the width of the cantilever should taper inwards from the base to the tip to match the bending moment, resulting in relatively uniform curvature over the whole device. In theory efficiency is maximized when the load is applied exactly at the tip of the triangle. The sharp tip generates concerning stress concentrations and is difficult to surgically place with precision, so a trapezoidal cantilever is a reasonable compromise.

3.4.1 Acoustic impedance matching

For a perfectly triangular cantilever loaded at the tip with length L and base B we have

$$b(x) = B\left(1 - \frac{x}{L}\right), \qquad I = \frac{Bh^3(L-x)}{12L}, \qquad M = -F(L-x). \tag{3.20}$$

We can now compute the deflection at the tip:

$$\frac{d^2w}{dx^2} = -\frac{M}{EI} = \frac{12FL(L-x)}{EBh^3(L-x)} = \frac{12FL}{EBh^3}, \qquad w = \iint_0^L \frac{12FL}{EBh^3}\,dx = \frac{6FL^3}{EBh^3}. \tag{3.21}$$

Therefore the spring constant k and mechanical impedance Z are

$$k = \frac{F}{w} = \frac{EBh^3}{6L^3}, \qquad Z(j\omega) = \frac{k}{j\omega} = \frac{EBh^3}{6L^3 j\omega}. \tag{3.22}$$

This formula is a reasonably good approximation for the significantly more complicated formula for a trapezoidal cantilever. Note that significant changes in mechanical impedance can be made by slightly adjusting the length and thickness of the cantilever, as the impedance is proportional to the cube of the thickness to length ratio.

A typical human eardrum behaves mostly spring-like until its first resonance at approximately 1 kHz. Umbo displacement in this regime is roughly 50 nm/Pa. Since the eardrum motion is more or less pistonic at these low frequencies, we can crudely approximate mechanical impedance as

$$Z_{\text{umbo}} \approx \frac{PA_{\text{drum}}}{v_{\text{umbo}}}, \qquad k_{\text{umbo}} \approx j\omega Z_{\text{umbo}} = \frac{PA_{\text{drum}}}{d_{\text{umbo}}}. \tag{3.23}$$

Using $A_{\text{drum}} \approx 50\,\text{mm}^2$, we get an eardrum spring constant of roughly 1 N/mm. At higher frequencies above 1 kHz, the eardrum tends to have multiple modes of vibration and the pistonic motion assumption breaks down, making the umbo more compliant than umbo velocity data would suggest. Therefore, an approximate low-frequency impedance match between cantilever and eardrum is acceptable. Substituting $k =$

1 N/mm into Equation 3.22 along with $E = 2.7\,\text{GPa}$ and $B = 3.5\,\text{mm}$ gives

$$\frac{L}{h} = \sqrt[3]{\frac{EB}{6k}} \approx 12. \tag{3.24}$$

In our case, $h = 200\,\mu\text{m}$ giving optimal $L = 2.4\,\text{mm}$. Our cantilevers are 3 mm long, and the stiffness can be increased by moving the contact point towards the base.

3.4.2 Cantilever stackup fine tuning

The ratio of piezoelectric thickness to base thickness is also relevant. Using (3.11) we see that electrical energy stored in a flat piezoelectric sheet is proportional to the volume of material and average stress. We can imagine a single-ended cantilever of total thickness h, created by stacking a PVDF layer of thickness h_1 on top of a Kapton layer of thickness $h - h_1$. With the assumption that Kapton and PVDF have similar Young's modulus,

$$T_{\text{avg}} = E\frac{d^2w}{dx^2}\left(\frac{h}{2} - \frac{h_1}{2}\right) \propto h - h_1, \qquad U_{\text{elec}} = \frac{d_{31}^2 T_{\text{avg}}^2 A h_1}{2\epsilon} \propto h_1(h - h_1)^2. \tag{3.25}$$

Keeping cantilever thickness (and hence mechanical impedance) constant, we see that the optimal PVDF thickness is $h/3$.

We can easily double the sensitivity of the sensor by putting another PVDF layer on the other side of the sensor. In the case of the differential cantilever there is the additional parasitic capacitance of the positive and negative electrodes through the Kapton. Assuming charge symmetry, the negative half of the cantilever is equivalent to a ground at the plane midway through. This gives a capacitance of

$$C = \frac{A\epsilon_{\text{PVDF}}}{h_1} + \frac{A\epsilon_{\text{Kapton}}}{h/2 - h_1} \propto \frac{\alpha}{h_1} + \frac{1}{h/2 - h_1}, \text{ where } \alpha = \frac{\epsilon_{\text{PVDF}}}{\epsilon_{\text{Kapton}}}. \tag{3.26}$$

With $\epsilon_{\text{PVDF}} = 12.5\epsilon_0$ and $\epsilon_{\text{Kapton}} = 3.2\epsilon_0$, we have $\alpha \approx 4$. Now we can write

$$U_{\text{elec}} = \frac{Q^2}{2C} \propto \frac{T_{\text{avg}}^2}{C} \propto (h - h_1)^2 \left(\frac{\alpha}{h_1} + \frac{1}{h/2 - h_1}\right)^{-1}. \tag{3.27}$$

Note that as $\alpha \to \infty$ this approaches the single-ended case in (3.25). For $\alpha \approx 4$ this maximum occurs when $h_1 \approx 0.24h$. Our cantilever design uses 50 µm PVDF over a 100 µm Kapton substrate, giving a total stackup height of 200 µm. Our stackup is very nearly optimal by this metric.

3.5 Summary

The current cantilever-mic design makes use of most of these rules of thumb for optimization. The cantilever is linear and predictable—it therefore has a controllable mechanical impedance that is tuned to match that of the umbo with attached ossicles by adjusting the sensor's dimensions. The placement of the PVDF on either side of the cantilever ensures that most of the device's elastic strain is stored in the PVDF parallel to d_{31}. By choosing a triangular cantilever and the ratio of the PVDF layer thickness to the Kapton center layer, we achieve relatively uniform stress in the PVDF, resulting in more efficient conversion of elastic energy to electrical energy. We also minimize unnecessary PVDF area and use shorter cables to minimize parasitic capacitance. Finally, the noise impedance of the charge amplifier is a rough impedance match to the electrical impedance of the cantilever-mic at audio frequencies.

Chapter 4

Charge Amplifier Design

4.1 Design criteria

Piezoelectric microphones behave as a charge source in parallel to a capacitor (or alternatively, a voltage source in series with a capacitor). Our particular piezoelectric sensors have a very small capacitance—on the order of 10 to 50 pF. Below 1 kHz the amplifier's noise current is often the dominant contribution to the noise floor of low capacitance sensors. Sensitive amplifiers on the market tend to be designed for larger capacitance sensors—they have higher or unspecified current noise, and they tend to use larger feedback capacitors. Therefore designing a sensitive amplifier specifically designed for low capacitance sensors was of paramount importance. The amplifier topology that makes the most sense for small piezoelectric sensors is the charge amplifier. It has similar noise specifications to a voltage amplifier, but with the advantage that its gain is unaffected by parasitic capacitance. The basic charge amplifier topology is similar to a transimpedance amplifier but with a feedback capacitor instead of a high value feedback resistor. The op amp servos the terminals to have the same voltage, meaning that all of the charge on capacitor C_{piezo} is transported to C_f, giving a midband charge to voltage gain of $1/C_f$. We can design a differential charge amplifier by connecting two matched charge amplifier input stages to a difference amplifier, as shown in Figure 4-1.

The noise floor of the charge amplifier was found to be comparable to that of

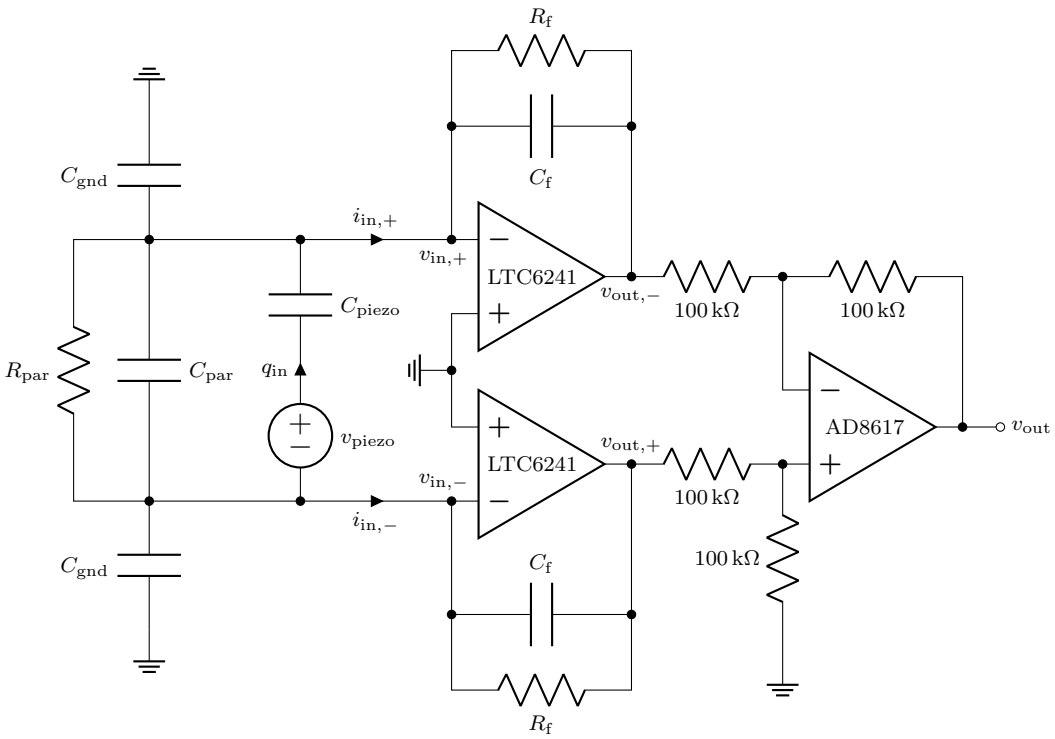

Figure 4-1: The differential charge amplifier topology. The differential-mode output of the piezoelectric sensor is modeled as the voltage source v_piezo in series with the capacitor C_piezo, together with a parasitic leakage capacitor C_par and resistor R_par. This amplifier has a gain of $v_\mathrm{out}/q_\mathrm{in} = 2/(C_\mathrm{f}+1/(j\omega R_\mathrm{f}))$, which is independent of parasitics. Our implementation uses $C_\mathrm{f} = 1\,\mathrm{pF}$ and $R_\mathrm{f} = 10\,\mathrm{G\Omega}$.

a voltage amplifier (implemented with a non–inverting amplifier). Since the gain stability of the charge amplifier confers a significant advantage over the voltage amp, we chose to pursue the charge amp topology.

4.2 Noise floor: A theoretical approach

There are five significant sources of noise in this circuit: Johnson noise from R_f and R_par, voltage noise and current noise from the first stages, and voltage noise from the second stage difference amplifier. For Johnson noise, it is easiest to treat the noise as a parallel current source. The input-referred current noise induced by the two resistors is therefore

$$\overline{i^2}_{\mathrm{in,Johnson}} = \frac{4k_BT}{2R_\mathrm{f} \parallel R_\mathrm{par}} \quad . \tag{4.1}$$

Note that common-mode noise is ignored here because the difference amplifier will reject it.

The noise voltage at $v_{\text{in},+}$ and $v_{\text{in},-}$ effectively produces a noise current $i_{\text{in,op1,v}}$ determined by the impedance of the sensor and the feedback network. While the noise voltage is not completely frequency-independent, flicker noise for the LTC6240 is negligible above 100 Hz, so v_{op1} can be modeled as a white noise source. Defining

$$Z_{\text{comm}} = R_{\text{f}} \parallel \frac{1}{j\omega(C_{\text{f}} + C_{\text{gnd}})} \tag{4.2}$$

and

$$Z_{\text{diff}} = R_{\text{par}} \parallel \frac{1}{j\omega(C_{\text{par}} + C_{\text{piezo}})} \quad , \tag{4.3}$$

the induced noise currents are

$$i_{\text{in,op1,v+}} = \frac{v_{\text{op1,}-} - v_{\text{op1,}+}}{Z_{\text{diff}}} - \frac{v_{\text{op1,}+}}{Z_{\text{comm}}}, \tag{4.4}$$

with $i_{\text{in,op1,v}-}$ defined equivalently. Next, define the differential input current $i_{\text{in}} = \frac{1}{2}(i_{\text{in,}+} - i_{\text{in,}-})$ so that

$$i_{\text{in,op1,v}} = (v_{\text{op1,}-} - v_{\text{op1,}+}) \cdot \left(\frac{1}{Z_{\text{diff}}} + \frac{1}{2Z_{\text{comm}}} \right) \quad . \tag{4.5}$$

Finally, because the two op amps are independent noise sources, the differential induced noise current has magnitude

$$\overline{i^2}_{\text{in,op1,v}} = 2\overline{v^2}_{\text{op1}} \cdot \left| \frac{1}{2Z_{\text{comm}}} + \frac{1}{Z_{\text{diff}}} \right|^2 \quad . \tag{4.6}$$

Current noise from the LTC6240 can be modeled as

$$\overline{i^2}_{\text{op1,total}} = \overline{i^2}_{\text{op1}} + \omega^2 \overline{q^2}_{\text{op1}} \tag{4.7}$$

where i_{op1} represents the white noise component and q_{op1} represents the frequency dependent component. We assume q_{op1} is more or less independent from v_{op1} [2]; the

correlation coefficient between op amp current noise and voltage noise is unspecified in the datasheet [2]. When two op amps are used for a differential charge amplifier, and only differential current noise is considered, the resulting noise becomes

$$\overline{i^2}_{\text{in,op1,i}} = \frac{\overline{i^2}_{\text{op1}} + \omega^2 \overline{q^2}_{\text{op1}}}{2} \quad . \tag{4.8}$$

The charge gain G of the op amp is given by

$$G = \frac{v_{\text{out}}}{q_{\text{in}}} = 2j\omega R_{\text{f}} \parallel C_{\text{f}}^{-1} = \frac{2}{C_{\text{f}} + (j\omega R_{\text{f}})^{-1}} \quad . \tag{4.9}$$

Therefore, the input-referred charge noise contribution from the second stage may be expressed as

$$\overline{q^2}_{\text{in,op2}} = \frac{\overline{v^2}_{\text{op2}}}{|G|^2} = \frac{\overline{v^2}_{\text{op2}}}{4} \left(C_{\text{f}}^2 + \frac{1}{\omega^2 R_{\text{f}}^2} \right) \quad . \tag{4.10}$$

Dividing $\overline{i^2}_{\text{in}}$ by ω^2 gives noise charge power density $\overline{q^2}_{\text{in}}$. Combining (4.1) through (4.10), expanding Z_{diff} and Z_{comm}, and collecting terms gives the equivalent noise charge (ENC) power density of

$$\begin{aligned} \overline{q^2}_{\text{in}} = {} & 2\overline{v^2}_{\text{op1}} C_{\text{tot}}^2 + \frac{1}{2}\overline{q^2}_{\text{op1}} + \frac{1}{2}\overline{v^2}_{\text{op2}} C_{\text{f}}^2 \\ & + \frac{1}{\omega^2} \left(\frac{4k_B T}{2R_{\text{f}} \parallel R_{\text{par}}} + \frac{\overline{i^2}_{\text{op1}}}{2} + \frac{2\overline{v^2}_{\text{op1}}}{(2R_{\text{f}} \parallel R_{\text{par}})^2} + \frac{\overline{v^2}_{\text{op2}}}{(2R_{\text{f}})^2} \right) \quad , \end{aligned} \tag{4.11}$$

where

$$C_{\text{tot}} = C_{\text{piezo}} + C_{\text{par}} + \frac{1}{2} C_{\text{f}} + \frac{1}{2} C_{\text{gnd}} \quad .$$

4.3 Practical component selection

Several design guidelines can be extracted from (4.11). Parasitic leakage conductance and capacitance are universally bad from a noise perspective, and should be minimized for any given sensor design. Minimizing parasitic capacitance is especially important, as the $2\overline{v_{\text{op1}}^2} C_{\text{tot}}^2$ term of (4.11) is a significant part of the amplifier noise floor. Furthermore, the ratio of C_{piezo} to C_{f} is effectively the voltage gain of the first

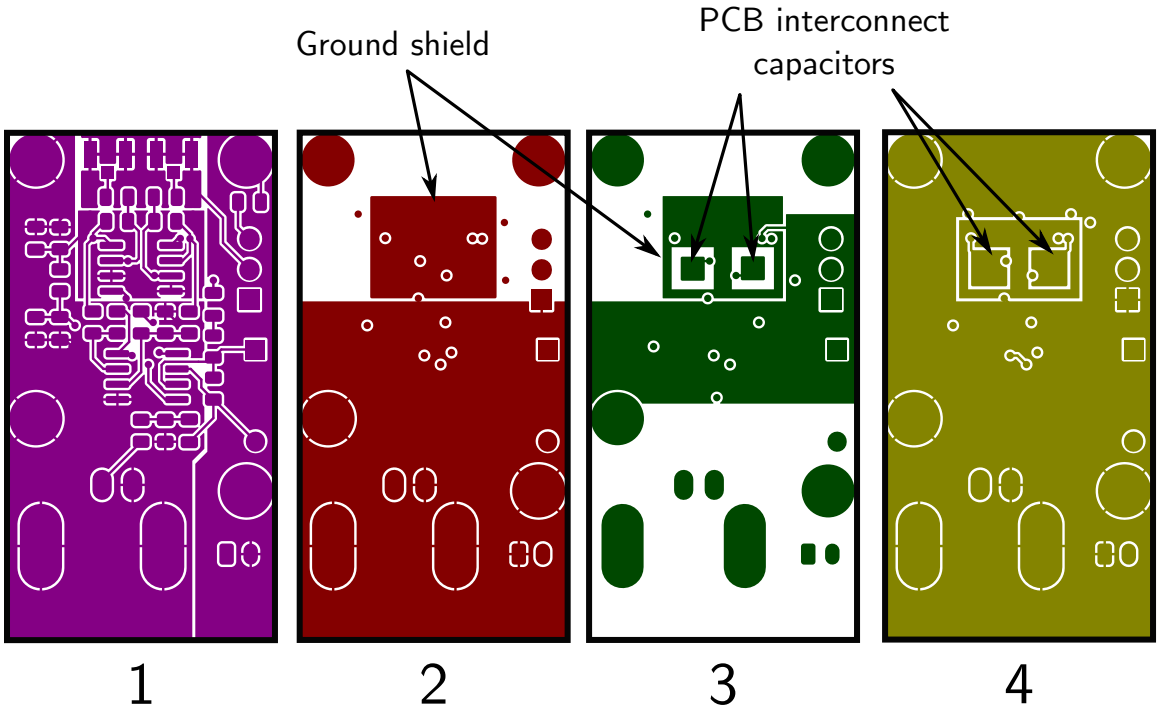

Figure 4-2: The four layer PCB design, showing 1 pF feedback capacitors on layers 3 and 4.

stage; C_f should be several times smaller than C_piezo to minimize the second-stage contribution to the noise floor. Working prototypes have been built with C_f up to 10 pF, but those with $C_\mathrm{f} = 1\,\mathrm{pF}$ work best. Since the differential charge amplifier requires good matching between the two input stages to achieve a good common-mode rejection ratio, PCB capacitors are used to implement C_f. By using a four layer PCB, and building the capacitors between the bottom two layers, each capacitor can be implemented in a 3 mm × 3 mm area with good matching and shielding as shown in Figure 4-2.

The value of R_f requires more care. Ideally R_f should be made as large as possible, but increased R_f gives worse bias stability. It has been observed that increasing R_f beyond 10 GΩ does not yield significant performance benefits.

One downside of increasing R_f is increased sensitivity to DC leakage current, which can come both from the op amp inputs and from the outside environment. The design presented in this paper can tolerate approximately 100 to 200 pA of DC leakage current, limited by the rail voltages and R_f. This requires the amplifier board and sensor to be designed very carefully to avoid leakage paths. One possible

Property	**LT1792**	**LTC6240**	**LTC6081**	**LTC6078**
$\bar{v}_{\text{in}}$ (nV/$\sqrt{\text{Hz}}$)	4.2	7	13	18
$\bar{i}_{\text{in}}$ (fA/$\sqrt{\text{Hz}}$)	10	0.56	0.5	0.56
C_{in}	27 pF	3.5 pF	3 pF	10 pF
Power	76 mW	6.7 mW	1.2 mW	0.2 mW

Table 4.1: For small capacitance devices, the LTC6240 offers the best performance, at the expense of power consumption compared to the LTC6081 and LTC6078.

solution to the low–frequency response of these charge amps is presented in [1]. This design still depends on a large feedback resistor R_{C}, which behaves similarly to R_{f} both in terms of noise floor and leakage current cancelling. A solution especially suited to integrated circuit design could be to replace R_{C} with a push–pull stage with low W/L MOSFETs. In practice, we found that this additional complexity was unnecessary; minimizing exposed copper of sensitive signals and careful control of surface contamination for the sensor and amplifier were sufficient to reduce leakage to tolerable levels.

The centerpiece of the amplifier is the low-noise op amp used for the first stage, as this sets the absolute lower bound on the noise floor. Choosing this op amp based on (4.11) requires balancing $\overline{v_{\text{in}}^2}$ and $\overline{i_{\text{in}}^2}$ over the desired frequency range and sensor capacitance. This requirement rules out op amps with bipolar or JFET input stages because these op amps typically have unacceptably high current noise. Op amps with CMOS input stages have voltage noise several times higher than top-of-the-line JFET or bipolar op amps, but with substantially lower current noise. Of these, the LTC6240 appears to offer the best combination of voltage noise and current noise, with the LTC6081 and LTC6078 providing respectable performance with lower power consumption. Previous use of the LT1792, which has significantly worse current noise than the LTC6240, caused the current noise to dominate the sensor noise floor at low frequencies. The LTC6240 also can run off of a 3.7 V lithium–ion cell, significantly reducing power requirements. Table 4.1 gives a more detailed comparison of the available op amps.

The second stage charge amp requirements are far more relaxed. The AD8617 has a noise floor of roughly 25 nV/$\sqrt{\text{Hz}}$ and the 100 kΩ feedback resistors contribute

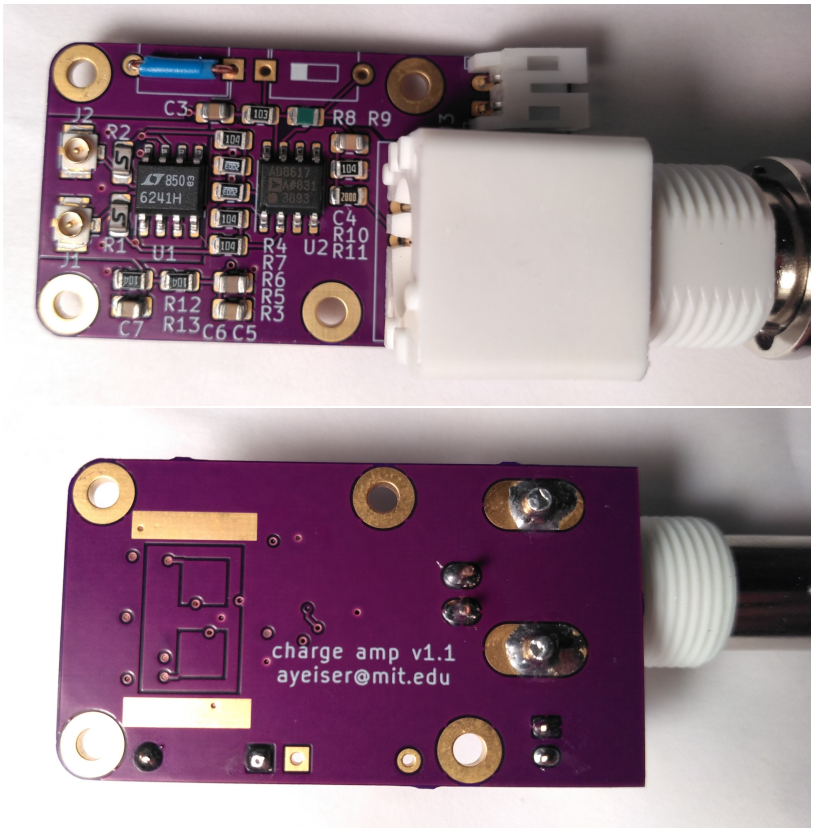

Figure 4-3: The assembled differential amplifier board. The terminals of C1 and C2 connected to `OUT+` and `OUT-` are visible on the back of the board.

$41\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$, for a total of $48\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$. Using (4.10) gives an input-referred white noise contribution of only $0.15\,\mathrm{e}^-/\sqrt{\mathrm{Hz}}$, which is insignificant compared to the noise floor of the complete amplifier.

Figure 4-3 shows the assembled differential charge amplifier. The present implementation features a switchable 20 dB gain stage and two single-order high-pass filters for DC signal rejection and level shifting. These provide an overall gain of $20\,\mathrm{V/pC}$ with a $-3\,\mathrm{dB}$ bandwidth of 170 Hz to 50 kHz.

4.4 Measurement and results

The amplifier noise floor was measured by open circuiting the amplifier inputs to emulate a 0 pF load, and recording several seconds of output noise. Over the target bandwidth of 200 Hz to 20 kHz, with a midband gain of $20\,\mathrm{V/pC}$, the amplifier

achieved an equivalent noise charge (ENC) of 2.96×10^{-17} C, or 185 e^-. Extrapolating from (4.11), the noise floor is expected to rise by approximately 8.7 e^- for every additional picofarad of load capacitance once this load capacitance dominates the noise floor. This approximation holds up fairly well; the measured noise floor with a 25 pF sensor attached is approximately 385 e^-, not too far off from the rough figure of 403 e^-. Figure 4-4 shows the measured noise spectral density.

The amplifier frequency response was measured by feeding a chirp signal into each input through a 10 pF capacitor. A very low signal of 10 mV_{pp} was required to avoid saturating the amplifier output. The amplifier exhibited a midband gain of 19.1 V/pC, within 5 % of the design gain of 20 V/pC, with a -3 dB bandwidth of 170 Hz to 50 kHz, as shown in Figure 4-5. The common-mode rejection ratio was measured to be approximately 30 to 40 dB, as illustrated in Figure 4-6.

4.5 Comparison to existing solutions

The charge amplifier reported here performs well compared to both charge amplifiers in the literature and on the market. Comparing these amplifiers can be difficult though. Noise specifications are usually given as an RMS value over a bandwidth, or a spectral noise density at specific frequencies, often as output-referred voltage noise. To compare amplifiers with different gain, one must compare ENC over similar frequency ranges or noise charge density at comparable frequencies. Ideally, such a comparison should be made at different input capacitances, but this information is often not available.

For the purposes of this thesis, the metric of "goodness" of a charge amplifier is determined solely by its noise performance at audio frequencies with a load of 10 pF or less. Systems with poor noise performance at audio frequencies and small-capacitance loads are not necessarily bad products, but rather likely optimized against different metrics, such as different target bandwidth, input capacitance, configurability, and power consumption, for example. Nevertheless it appears that no charge amplifier on the market, differential or otherwise, approaches the 185 e^- ENC over the audio

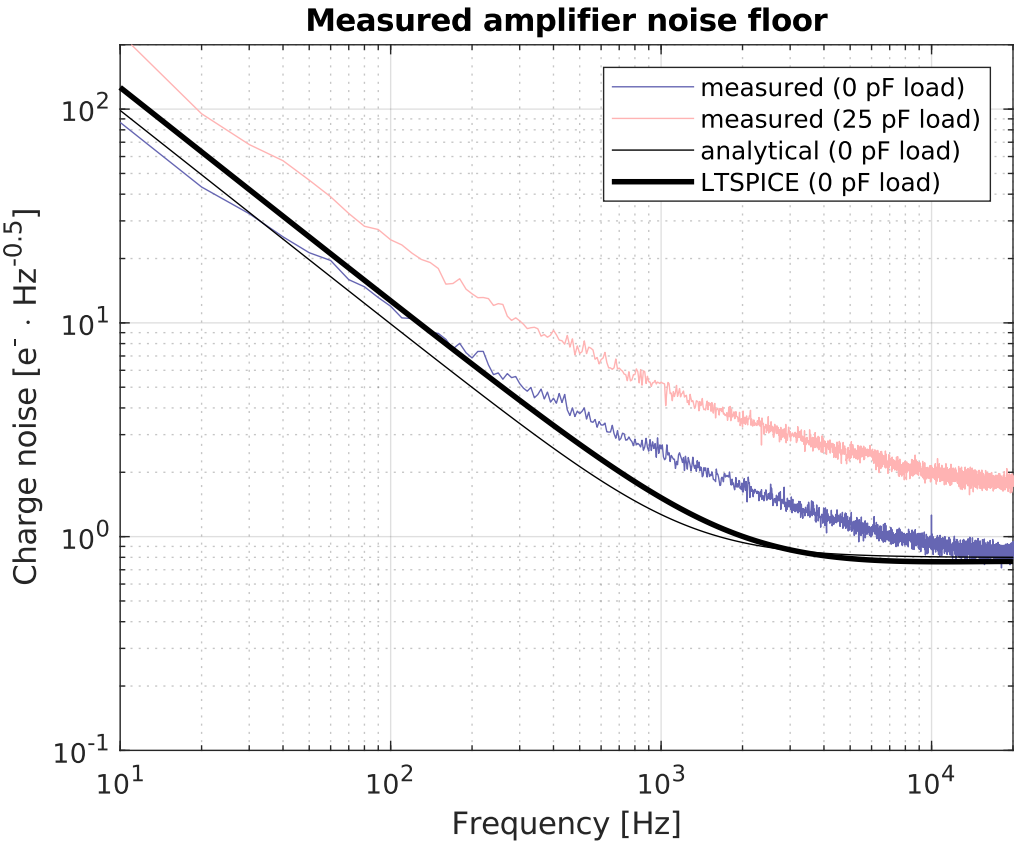

Figure 4-4: The output voltage noise of the amplifier under no load on its high gain setting (20 V/pC). An ENC of roughly 185 e^- was achieved over a bandwidth of 200 Hz to 20 kHz. With a 25 pF sensor, the ENC was roughly 385 e^- over the same bandwidth.

frequency range of the amplifier presented here.

The best alternative differential charge amplifier appears to be the CEC 1-328. This amplifier's datasheet specifies a bandwidth of 10 kHz and an ENC of 3100 e^- [6]. The Femto HQA-15M-10T appears to be the best single-ended charge amplifier, with an ENC of 2200 e^- over its 250 Hz to 15 MHz bandwidth specified in its datasheet [10]. However, at least 630 e^- of the noise between 250 Hz and 20 kHz comes from the 1-GΩ bleed resistor incorporated in the input stage. All other amplifiers reviewed including the charge amplifiers from PCB Piezotronics, B&K, Kistler, and Measurement Specialties (now TE Connectivity) had an ENC well over 10 000 e^- over their respective bandwidths.

The charge amplifiers described in the literature are fairly diverse in their construction and application. Many of these devices, such as those described in [12]

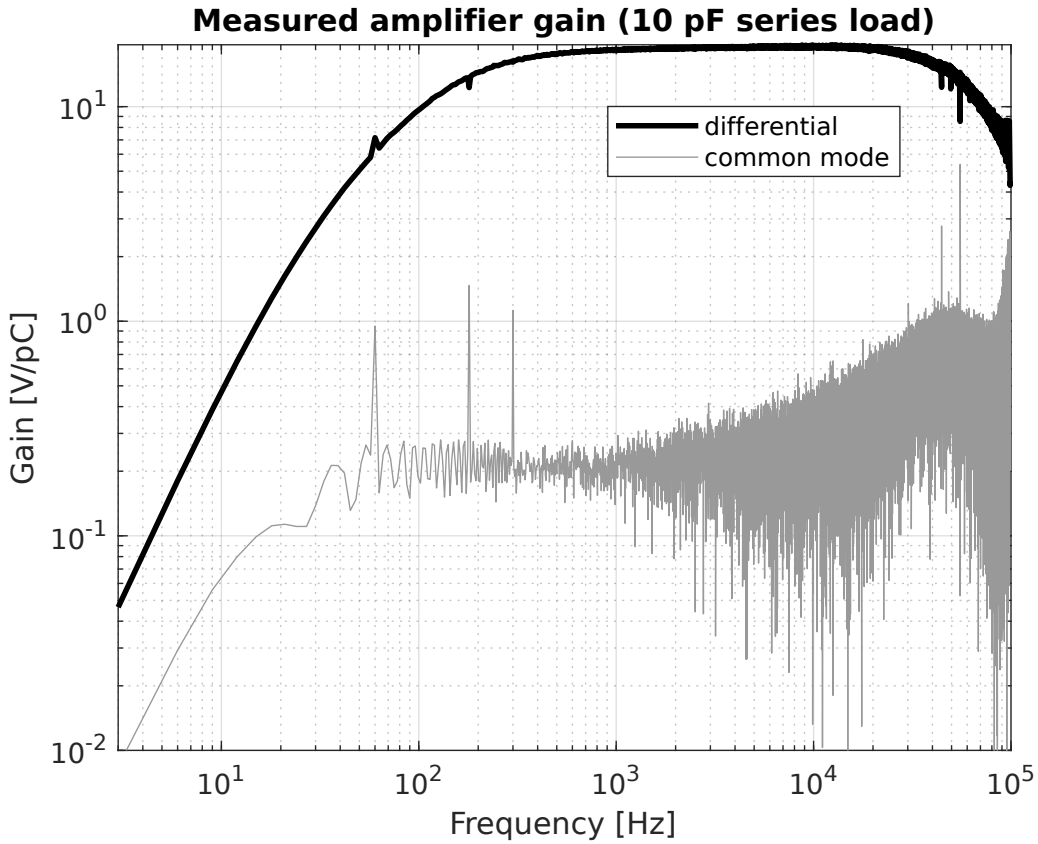

Figure 4-5: Gain was measured by connecting a 10 pF capacitor in series with each input. We achieved a charge gain of 1.91×10^{13} V/C, within 5 % of the design gain.

or [4] are designed for very-high-speed applications like ultrasonic sensing or particle detectors, applications where frequencies below 10 kHz can be safely ignored. Those amplifiers are less affected by the ω^{-2} term in (4.11) than are audio-frequency charge amplifiers, allowing for lower-value feedback resistors and input stages with higher current noise. Some designs also feature reset switches instead of feedback resistors. While this strategy is effective at reducing noise for accelerometers as described in [7], it is unsuitable for devices that require continuous audio output.

Some of the best audio-frequency charge amplifiers described in the literature are designed as dust particle sensors. A recent paper [13] describes a differential charge amplifier implemented in custom silicon. The ENC of this amplifier was only 114 e^- over its 7 Hz to 300 kHz bandwidth with a 5.4 pF input capacitance. This exceptional noise figure was implemented with custom silicon, using 0.1 pF feedback

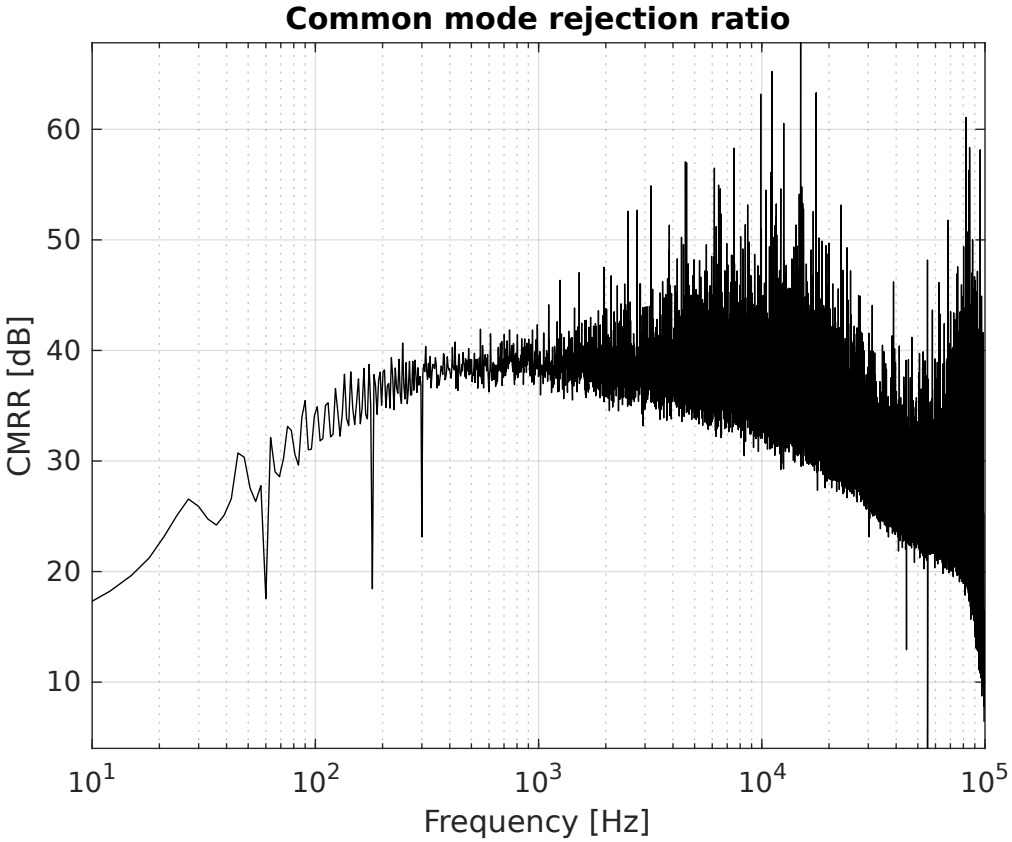

Figure 4-6: A common-mode rejection ratio of 30 to 40 dB is a reasonable expectation for this amplifier. The spikes at 60, 180, and 300 Hz appear to be due to poor shielding and long cable runs to test equipment.

capacitors and 7.6-TΩ pseudoresistors. While this design is extremely impressive, it is not realizable without custom silicon. A more conservative single-ended charge amplifier using off the shelf components was designed for the Electrostatic Lunar Dust Analyzer [9]. This amplifier used the LTC6240 with a 30-GΩ feedback resistor and 5-pF feedback capacitor, and had a measured ENC of 356 e^- over a 10-kHz bandwidth with an unknown input capacitance [18].

4.6 Summary

Our differential charge amplifier design outperforms commercially available charge amplifiers by a substantial margin for low-capacitance sensors at audio frequencies. Our design also performs comparably to charge amplifiers in literature using off-

the-shelf hardware, only falling behind highly sensitive charge amplifiers on custom silicon. We believe our design is quite competitive for these design criteria—our charge amplifier is an affordable high-performance solution to sensing signals from small piezoelectric devices.

Chapter 5

Implantation

5.1 Cantilever platform

Although a micromanipulator attached to an optical air table is effective for holding the cantilever-mic in place for testing on cadaveric temporal bones, this arrangement is clearly insufficient for implantation in a living patient. As such, we have been prototyping the surgical hardware required to hold the cantilever-mic in place. This hardware must be mechanically stiff but also low-profile and easily adjustable by a surgeon to overcome variations in the anatomy.

5.1.1 Middle ear surgical constraints

Implantation of a cochlear implant requires access to the middle ear cavity through a surgical facial recess opening. A mastoidectomy is performed—the mastoid bone with ear cells behind the ear are drilled away. This allows access to the middle ear cavity through the facial recess with is bordered anteriorally by the chorda tympani and posteriorally by the facial nerve. Great care must be taken to avoid damaging the facial nerve at the border of the facial recess, which controls the ipsilateral muscles of the face. The chorda tympani is responsible for taste perception on the ipsilateral anterior two thirds of the tongue—it is preserved when possible although it is frequently sacrificed when necessary. The cochlear implant electrode array is installed

through the round window of the cochlea.

The facial recess is typically about 2 to 3 mm high anterior-posterior. The length superior-inferiorly varies depending on bordering nerves, but 6 mm is a reasonable estimate. The umbo, which is the tip of the manubrium of the malleus, is typically approximately 8 mm past the opening of the facial recess in the medial direction. Any supporting mechanism for the cantilever microphone must be thin enough to fit through the facial recess and provide clearance for the cochlear implant electrode array's passage through the facial recess.

To ensure a stable and mechanically stiff anchor point, the supporting mechanism must be anchored to solid bone close to the facial recess. These constraints rule out the superior and posterior walls of the mastoid cavity, as these surfaces are in proximity to important structure like the brain or the sigmoid sinus venous system. Therefore we consider the anterior wall of the mastoid cavity, which is solid bone a few millimeters thick with only the bony ear canal behind. Sinking bone screws into this region is under investigation and seems to be feasible.

5.1.2 Supporting system for the cantilever

Our current approach firmly support the cantilever-mic surgically implant a supporting system. This supporting system consists of a "platform" that is screwed into the anterior wall of the mastoid cavity. This platform contains a lockable ball joint that interfaces with an arm of the supporting system that extends through the facial recess into the middle ear and holds the tip of the cantilever directly under the umbo. Figure 5-1 shows platform screwed into the ear with the arm extending into the facial recess. The joint between the platform and the arm is located within the mastoid cavity lateral to the facial recess, allowing for adjustment of the arm, demonstrated by an endoscope picture in Figure 5-2.

We built some older prototypes out of bent 175 µm 304 stainless steel sheet, but this approach was limited by the geometries achievable with sheet metal. These designs used loop on the arm to slide over a hook on the platform, allowing small adjustments of the arm in and out, side to side, as well as up and down by bending

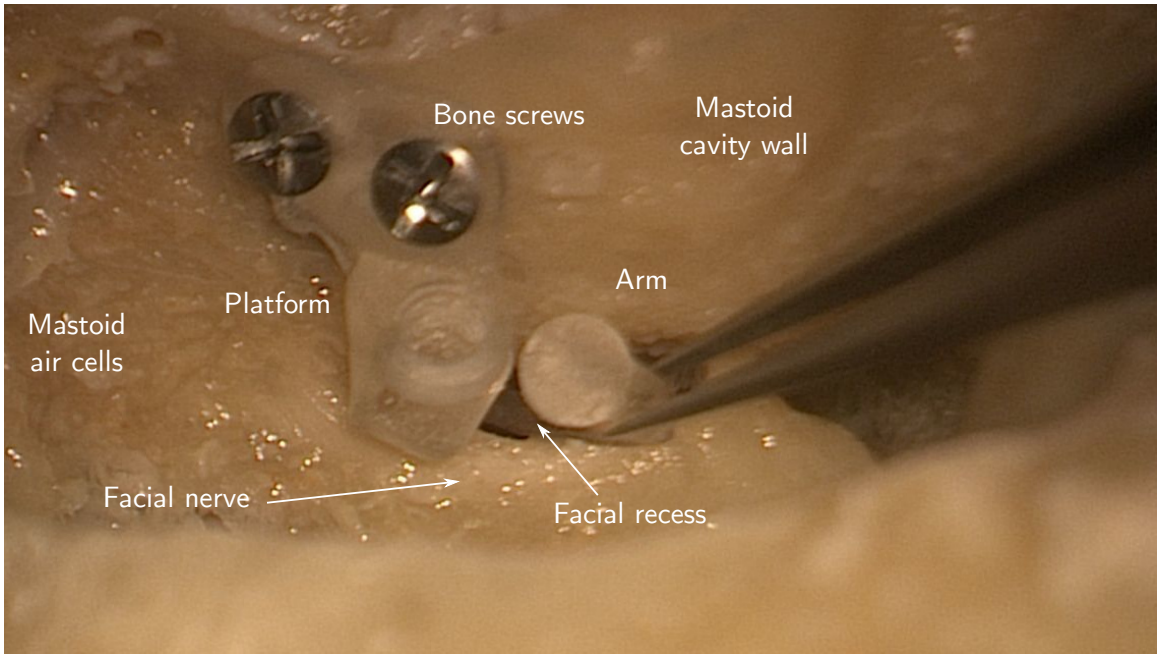

Figure 5-1: The supporting system is screwed into the anterior wall of the mastoid cavity. Here the arm is held with forceps, about to be inserted into the socket on the support. The arm extends through the facial recess into the middle ear cavity.

the hook. Unfortunately the sheet metal bends near the hook and the loop were too flexible, so force on the tip of the sensor bent the platform more than the sensor. The hook and loop system was not particularly stable either—there was no good way to lock the platform and the arm together other than with glue.

With sheet metal fabrication out the window, we decided to 3D-print our design from titanium and replace the hook and loop with a ball and socket, with the socket located on the platform where the hook previously was. Several of these designs are shown in Figure 5-4. So far we have been developing the design using a resin printer. The main difficulty is that the dimensions, angles, and curvature of the mastoid cavity, facial recess, and middle ear vary significantly between people, so we have been testing our designs in several different preserved specimens.

The current iteration of our design is shown in Figure 5-3. A curved platform with ball joint socket is screwed into the anterior mastoid cavity wall such that the bottom corner of the platform near the socket makes contact approximately 1 mm lateral to the opening of the facial recess. Titanium bone screws are fed through the two holes at the top of the platform. Tightening the screws slightly flattens the platform

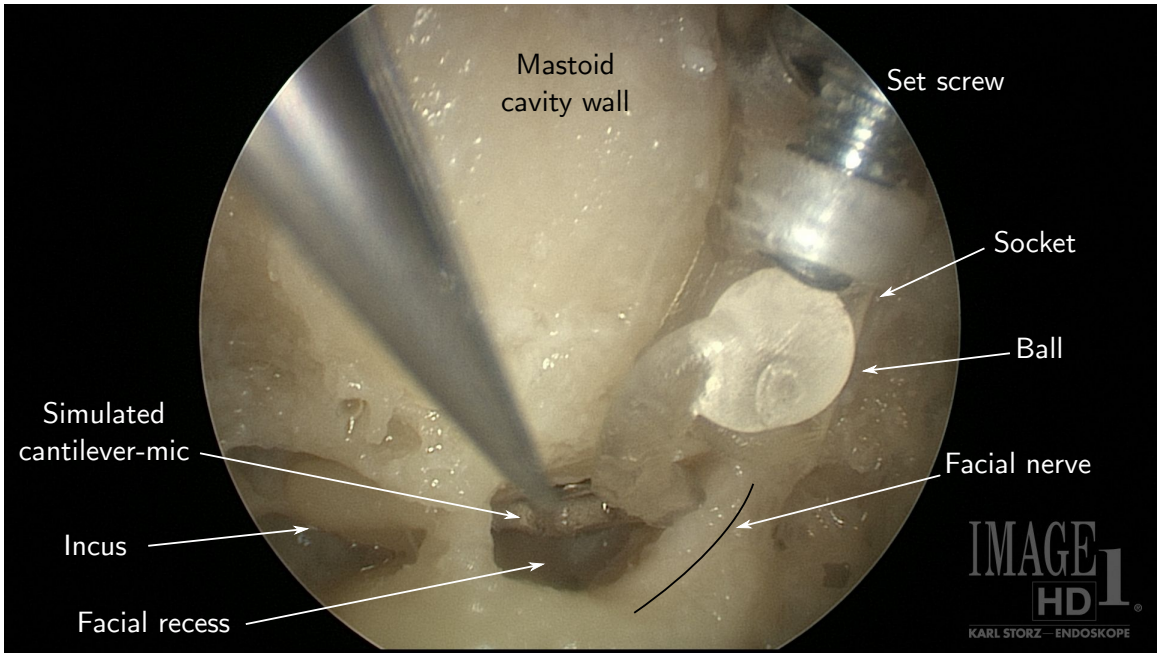

Figure 5-2: The ball joint, taken through an endoscope.

against the anterior mastoid cavity wall, providing extra stiffness. The 2.5 mm ball is locked into place with an M2 set screw (not shown in the figure) when the arm is in the correct position holding the sensor under the umbo.

This ball joint design does require some minor redesign of the sensor in order to work. The inert "tail" of the current cantilever microphone does not fit through the facial recess and gets in the way of the rigid arm. By chopping down the size of the tail to 2 mm or so and replacing the U.FL connectors with thin a shielded twisted pair cable, we should be able to fit cantilever sensor onto the arm with no further modifications.

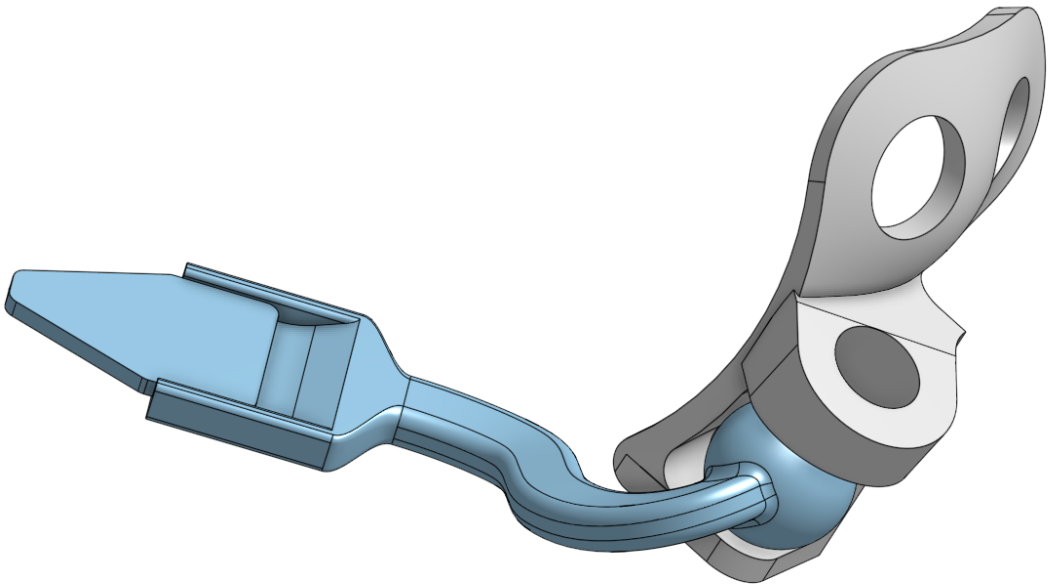

Figure 5-3: The ball and socket design (for a right ear). The cantilever-mic is glued into the shovel-like piece at the end of the arm (blue) of the supporting system. The platform section of the supporting system is drawn in gray and interfaces with the arm at the ball joint.

Figure 5-4: A list of prototype platform designs with initial sheet metal designs on the left and improved ball joint designs to the right.

Chapter 6

Measurement Techniques

6.1 Experimental goals

Our experiments are designed to produce unambiguous, replicable measures of sensor and amplifier performance. Where possible, we try to present data in a manner agnostic to the particular details of our experimental setup. For instance, all noise and sensitivity data are specified relative to the charge output of the microphone rather than the voltage output of the amplifier, allowing a meaningful comparison between similar amplifiers with different gain and transfer function. One of the key advantages of a charge amplifier here is the consistent transfer function from input charge to output voltage irrespective of parasitic resistance and capacitance.

The gold standard of microphone sensitivity statistics is equivalent input noise (EIN), averaged over 1/3-octave intervals. EIN can be calculated by dividing the electrical noise floor by the sensitivity. Two microphones with completely different mechanisms of operation can be compared side by side with EIN, although EIN is often not directly specified in the literature. Unfortunately, EIN discards critical information about the sensitivity transfer function and the shape of the noise floor, and so for comparing PVDF piezoelectric microphones, we also provide the sensitivity and noise floor independently.

6.2 Instrumentation

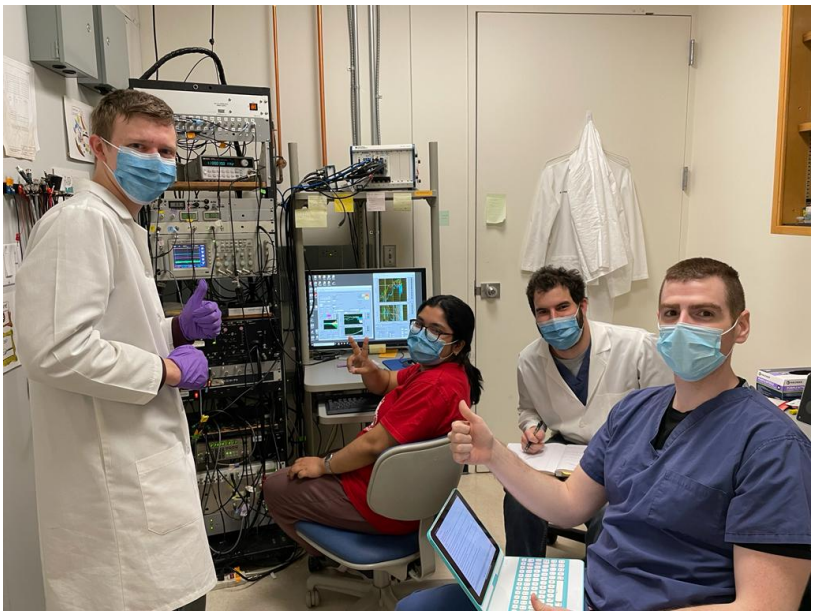

Figure 6-1: From left to right: the author, Annesya, Lukas, and Chris. Behind me an equipment rack connected to the soundproof room is visible, featuring a signal generator, oscilloscope, audio power amplifier, laser doppler vibrometer, and tunable analog antialiasing filters. Annesya is operating an National Instruments PXI running custom measurement software.

Most measurements took place at Massachusetts Eye and Ear Institute (MEEI), where we have a substantial collection of acoustic measurement equipment shown in the background of Figure 6-1, as well as a soundproofed electrically quiet room. The experiments were carried out in this chamber on an air table to reduce the effect of external acoustic noise or vibration. Our electronic equipment rack allows us to plug different pieces of hardware together, with data collection taking place on a National Instruments PXI running custom software. This software can output various stimuli, record up to four analog inputs, and compute the frequency response of each input channel relative to the output stimulus. Many of our experiments required cadaveric human temporal bones. MEEI has the facilities to process, store, and modify these bones for experimental use.

```
function [fft_avg, f] = noise_avg_fft(noise, fsamp, fbin)
% Computes smoothed fft of noise with sampling frequency fsamp
% Final bin size is fbin

n_s = floor(fsamp/fbin);              % samples per segment
fbin = fsamp / n_s;                   % adjust bin size
n_avgs = floor(length(noise) / n_s);  % number of segments

% Chop noise into n_avgs segments of length n_s, compute fft
fft_avg = fft(reshape(noise(1:n_s*n_avgs), n_s, n_avgs), [], 1);

% RMS average of fft's, normalized to V/Hz^0.5
fft_avg = sqrt(sum(abs(fft_avg).^2, 2).*2.*fbin./n_avgs)./fsamp;
fft_avg = fft_avg(2:(floor(length(fft_avg)/2)+1));
f = fbin * (1:length(fft_avg)).';
end
```

Figure 6-2: Code for computing smoothed noise floor based on a recording of noise.

6.3 Charge amplifier performance

6.3.1 Noise floor

Noise floor was measured by measuring ten seconds of amplifier output with no external stimulus. This interval was broken up into one hundred individual 100 ms segments, an FFT was computed for each segment, and then the FFT was normalized to units of $V/\sqrt{Hz}$. The resulting noise voltage could then be divided by the known frequency response of the charge amplifier, giving input-referred noise charge. MATLAB code for computing this noise floor is listed here in Figure 6-2.

Since the noise floor depends on the capacitance and parasitic conductance of the attached sensor, care was taken to measure noise floor with the amplifier open-circuited and with the sensor on the lab bench and implanted. A comparison between the amplifier's predicted noise floor, measured open-circuit noise floor, and measured noise floor with sensor is present in Chapter 4. Data comparing the noise floor with sensor on the lab bench and implanted is shown in Chapter 7.

6.3.2 Gain and CMRR

Charge amplifier gain can be measured by connecting a voltage source to the amplifier through a series a 10 pF capacitor. This allows a known amount of charge to be injected into the amplifier. To measure our amplifier's frequency response at 20 V/pC, the maximum usable stimulus amplitude was 10 mV.

The 10 pF series capacitor was used to measure the response from the positive and negative differential inputs. Subtracting and adding these frequency responses gives the differential and common mode gain, respectively. The common mode rejection ratio (CMRR) is the ratio between the measured common mode and differential mode gain. The measured gain and CMRR are reported in Chapter 4.

6.3.3 EMI sensitivity

Electromagnetic interference is not particularly easy to quantify. Nevertheless we can make a few simplifying assumptions. Since our sensor is an audio-frequency device we can assume all EMI comes from quasistatic sources–either purely electric field or purely magnetic field. Magnetic interference is effectively zero for high impedance piezoelectric sensors, and our sensor has negligible flux-collecting loops, so we will pursue a model of interference from external electric fields. The general effect of electric-field interference from distant sources is to create a relatively flat potential relative to sensor ground. Therefore we can define the EMI capacitance as the capacitance between the sensor's charge electrodes and an "electrode at infinity."

Experimentally determining the EMI capacitance can be done by testing the charge response of the sensor to a flat potential. The sensor is placed in the middle of a rough sphere of aluminum foil a few inches in diameter, with no contact between the sensor and the aluminum foil. The foil ball is insulated from ground and connected to a voltage source. The EMI capacitance is the charge output of the sensor divided by the voltage input to the foil electrode. The measured EMI capacitance is reported in Chapter 7 in comparison to our group's previous sensors.

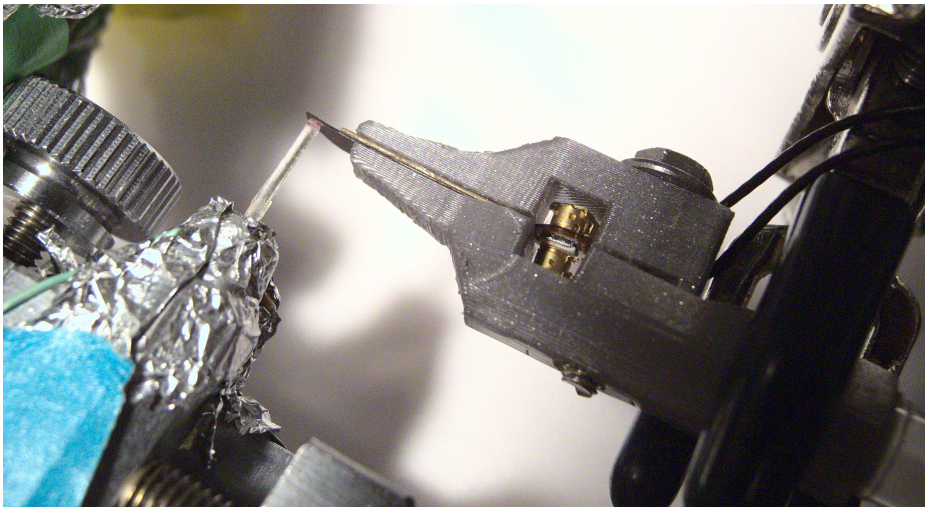

Figure 6-3: Testing displacement sensitivity with a thin glass rod glued to a piezo stack

6.4 Displacement sensitivity

Before implanting sensors in temporal bones, we had to test them and make sure they are working as desired. Our testing procedure was fairly simple here—a piezo stack was connected to a voltage source and calibrated with a laser doppler vibrometer (LDV). A simple voltage-displacement relation could be established: for our particular piezo stack, we measured 16.87 nm/V over a wide frequency range. A thin glass rod rounded at one end was glued to the piezo stack to eliminate electrical contact between the stack and the cantilever microphone. Stimulating the piezo stack with a log chirp and measuring the output of the cantilever through our charge amplifier gave the cantilever-mic's displacement sensitivity.

The displacement sensitivity serves the dual purpose of testing the sensor's actual performance as compared to our theoretical models as well as providing a baseline for implanted measurements.

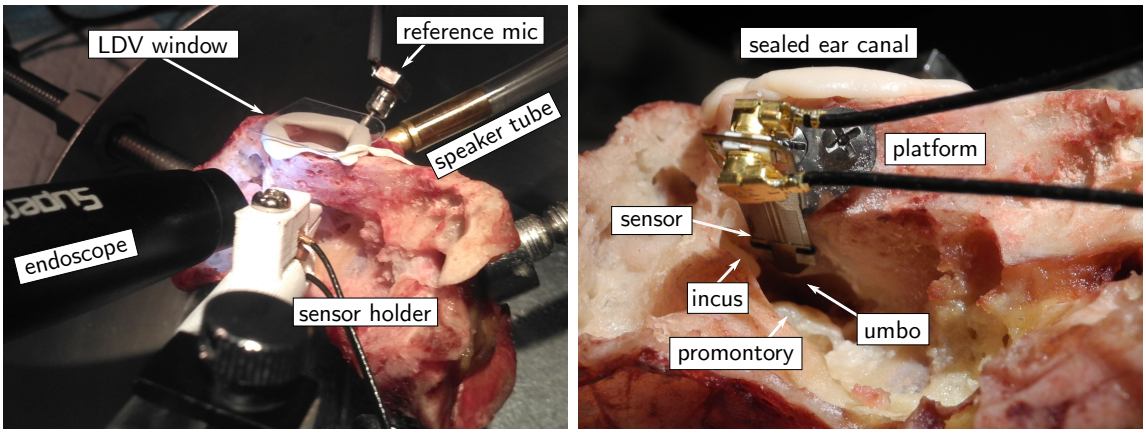

Figure 6-4: Left: a cantilever experiment with cantilever held in a micromanipulator. Right: a rigid platform prototype showing the cantilever tip under the umbo.

6.5 Temporal bone measurements

6.5.1 Ear canal pressure

We drilled out the cadaveric temporal bones through the mastoid air cells to expose the middle ear cavity, similar to the mastoidectomy for a cochlear implant. The cantilever microphone was rigidly held in place under the umbo using either a micromanipulator or a custom steel platform, depending on the experiment. With the micromanipulator we could perform a parameter sweep of sensor static offset and insertion depth and measure the effect of each parameter on sensitivity.

For these experiments we sealed the ear canal and excited sound pressure through a tube connecting the ear canal to a Beyerdynamic speaker. Ear canal pressure was monitored with Knowles EK3103 microphone with a thin probe tube 1 to 2 mm from the umbo. The frequency response of the reference mic with probe tube installed was calibrated periodically using a Larson Davis reference microphone with a known flat frequency response. The Larson Davis microphone was also calibrated against a standard 10 Pa 250 Hz reference source.

With the Knowles probe-tube microphone calibrated we can directly measure the pressure sensitivity of the cantilever-mic by dividing the amplifier output by the probe-mic pressure frequency response. We also take laser doppler vibrometry (LDV) measurements of umbo velocity by placing a retroreflective dot on the umbo as a

target for the LDV laser. The LDV allows us to measure the cantilever's displacement sensitivity while implanted, as well as the umbo velocity vs pressure. We report the sensitivity referenced to ear canal pressure and to umbo displacement in Chapter 7; we also compare implanted displacement sensitivity to measurements taken with the piezo stack. The eardrum's compliance is reported as the ratio umbo velocity to ear canal pressure—Chapter 7 compares compliance with and without the cantilever-mic implanted.

6.6 Location of results

In general, results pertaining to the amplifier alone like gain, CMRR, and open-circuit noise floor are found in Chapter 4, whereas Chapter 7 provides a comprehensive study of microphone-and-amplifier system. In particular we display the effect of altering the contact point with the umbo on cantilever-mic's pressure sensitivity, displacement sensitivity, and its effect on umbo velocity. We compare the cantilever-mic's pressure sensitivity, EMI sensitivity, and noise floor to that of our lab-group's previous sensors and we compare EIN to a commercial hearing aid microphone.

Chapter 7

Results

7.1 Introduction

In this section we present the results of the experiments described in Chapter 6 pertaining to the cantilever sensor with a focus on implanted performance. We evaluate how varying the conditions of implantation affects sensor performance and compare the cantilever-mic's performance to our lab's previous designs and to a commercial hearing aid microphone when applicable.

7.2 Electrical noise

7.2.1 Noise floor

The differential amplifier with a sensor connected has a measured noise floor without stimulus of roughly 385 e^- over the bandwidth 100 Hz to 20 kHz. Figure 7-1 illustrates the measured curves. The "dry sensor" curve was a noise floor measurement taken with the cantilever-mic in ideal laboratory conditions. "Sensor off umbo" and "sensor on umbo" were both taken with the cantilever-mic inside the middle ear cavity. Note that the "sensor off umbo" noise floor is nearly identical to the dry sensor. The "sensor on umbo" noise floor is significantly higher at low frequencies and has some additional noise around 1 to 2 kHz. The increased low-frequency noise of the "sensor

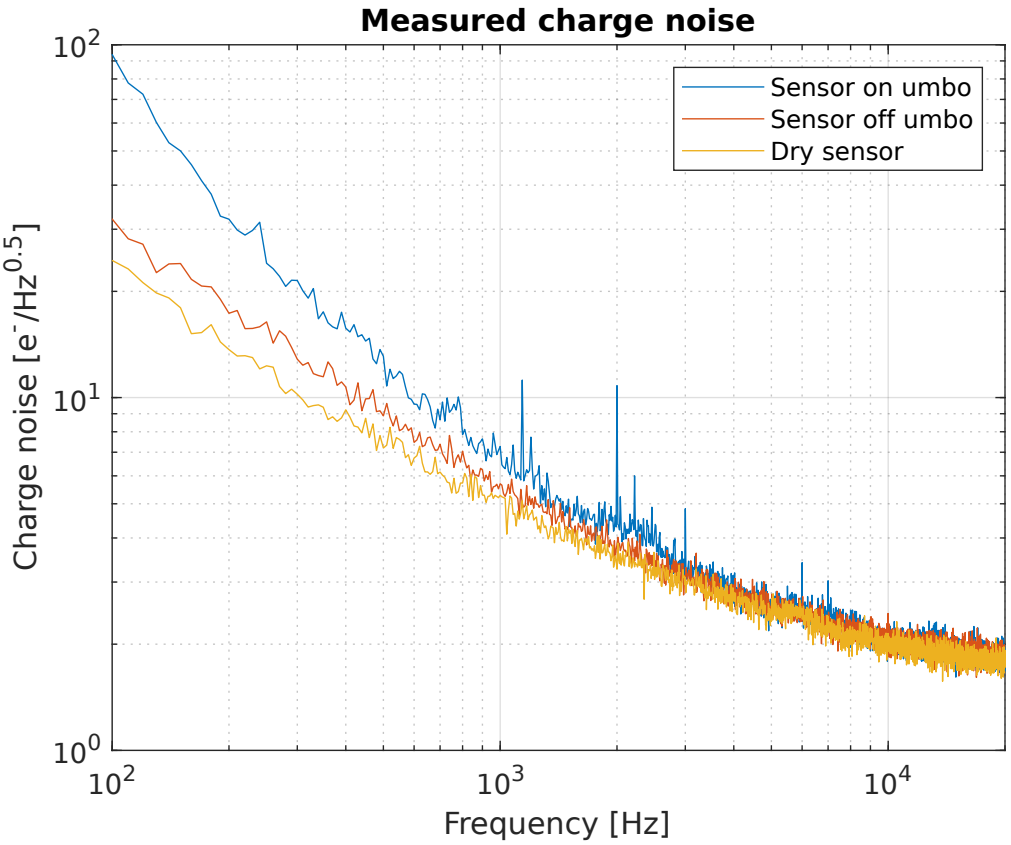

Figure 7-1: Noise floor with sensor connected to the amplifier. The blue curve shows the sensor contacting the umbo, the red curve shows the sensor a fraction of a millimeter below the umbo, and the yellow curve shows a dry sensor outside of a temporal bone.

on umbo" curve is likely due to structural vibration in the building coupling through the relatively long lever arm of the micromanipulator, moving the cantilever-tip relative to the umbo. The spikes between 1 to 2 kHz are likely acoustic background noise coupling through the tympanic membrane. Though the experiments were conducted in a sound-proofed chamber on an optical air table, the chamber cannot isolate all low-frequency noise, especially building structural vibration through the floor.

This noise floor is significantly lower than our previous charge amplifiers, which had a noise floor of approximately 1250 e^- with no load over a comparable bandwidth. Our amplifier also outperformed any commercially available preamplifiers over this bandwidth and stacked up well against custom low noise amplifiers described in literature, as described in Section 4.5.

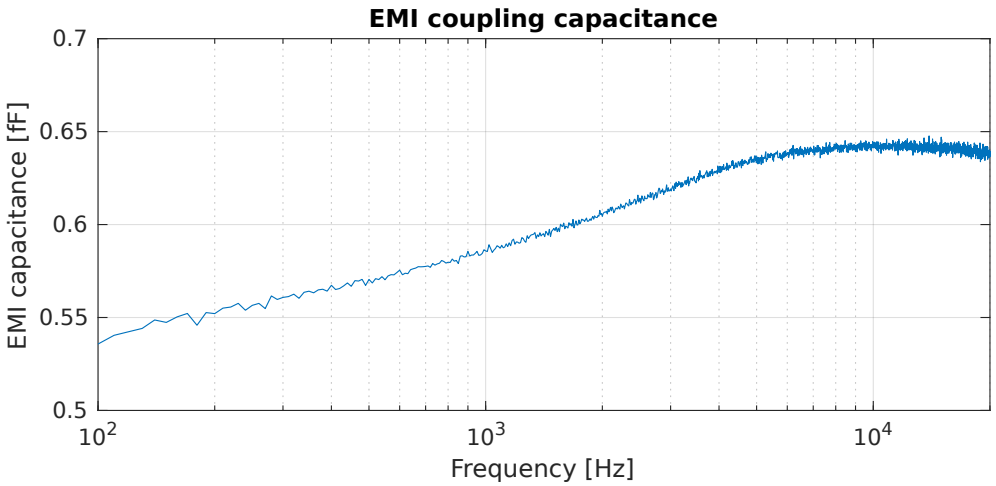

Figure 7-2: Measured EMI coupling capacitance for the cantilever. For comparison, a 4 mm-diameter sphere has an EMI coupling capacitance of over 200 fF.

7.2.2 Electromagnetic interference

The cantilever sensor's fully differential design and integrated ground shield led to a dramatic reduction in EMI, with this EMI coupling capacitance demonstrated in Figure 7-2. We measured 54 dB reduction in EMI compared to the cantilever's unshielded single-ended predecessors. This was readily empirically apparent—our old microphone designs would output a strong signal at mains frequency and harmonics when the lights were turned on in the room, whereas the shielded cantilever was immune to this issue. Similarly, care had to be taken to ground the Beyer speaker with the drum and coch-mics, whereas the cantilever-mic picked up very little coupling either way.

7.3 Cantilever sensitivity

As described in Section 6.5, we tested the cantilever-mic's response to ear canal pressure in a temporal bone while holding the sensor in a micromanipulator. This allowed us to make fine adjustments to the insertion depth and the static offset, changing the contact point between the sensor and the umbo. Moving the contact point affects the stiffness of the cantilever. Pressing the cantilever harder against the umbo causes the eardrum and ossicular chain to deform, which can affect the overall

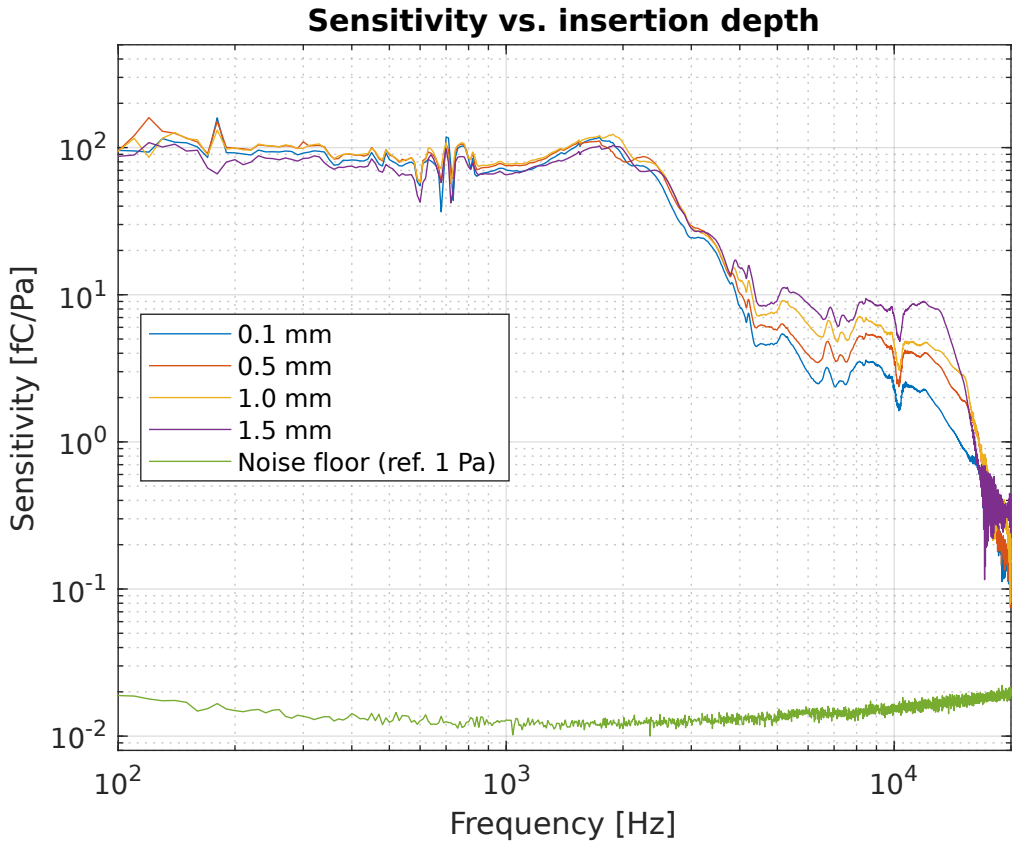

Figure 7-3: Pressure sensitivity (referenced to ear canal pressure) vs. insertion depth. Noise floor is averaged over a 1/3-octave bandwidth and referenced relative to 1 Pa (94 dB SPL) ear canal pressure. Insertion depth is referenced to the tip of the cantilever.

sensitivity. Overall we wanted to determine the effect that the insertion depth and static offset had on overall performance both to explore the possibility of adjusting these values for improved sensitivity and evaluating the tolerance of the cantilever to placement errors.

7.3.1 Insertion depth

Increasing the insertion depth of the cantilever-mic moves the contact point cantilever with the umbo closer to the cantilever base. Moving the cantilever up towards the eardrum increases both the stiffness and displacement sensitivity of the cantilever. At low frequencies where the eardrum is less stiff, increasing the cantilever stiffness seems to damp umbo motion rather than contributing to higher sensitivity. At high frequencies where the eardrum and attached ossicles are mass-dominated, umbo

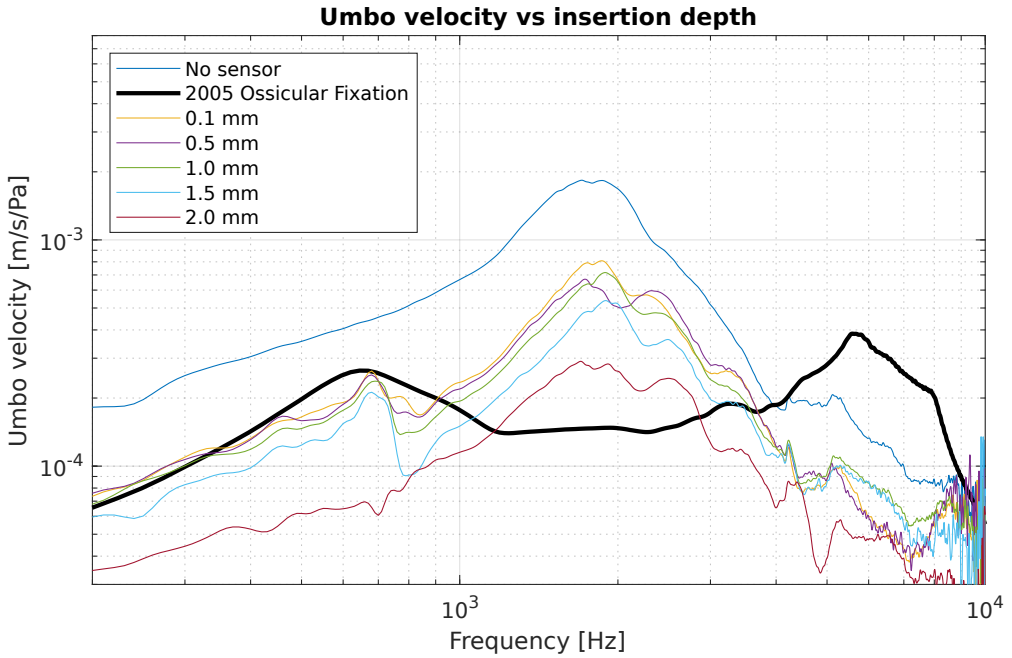

Figure 7-4: Increasing insertion depth can be seen to reduce umbo velocity at low frequencies.

motion is less affected and the higher displacement sensitivity boosts overall performance. These relations are shown in Figures 7-3, 7-4, and 7-5. In these figures, noise floor is measured by recording several seconds of amplifier output with no stimulus, performing a Fourier transform, and then normalizing to 1/3-octave bin size.

One interesting sidenote is that the measured displacement sensitivity in the ear is lower than when measured with the piezo stack. This is possibly due to undesired flexibility in the fixture holding the cantilever in place, or due to imperfect coupling between the umbo and the cantilever tip. More investigation is required to pinpoint the source of this loss of performance.

Overall we found that insertion depth has a substantial effect on performance. However this effect is small enough that placement error of $\pm 0.5\,\mathrm{mm}$ does not result in substantial loss of performance.

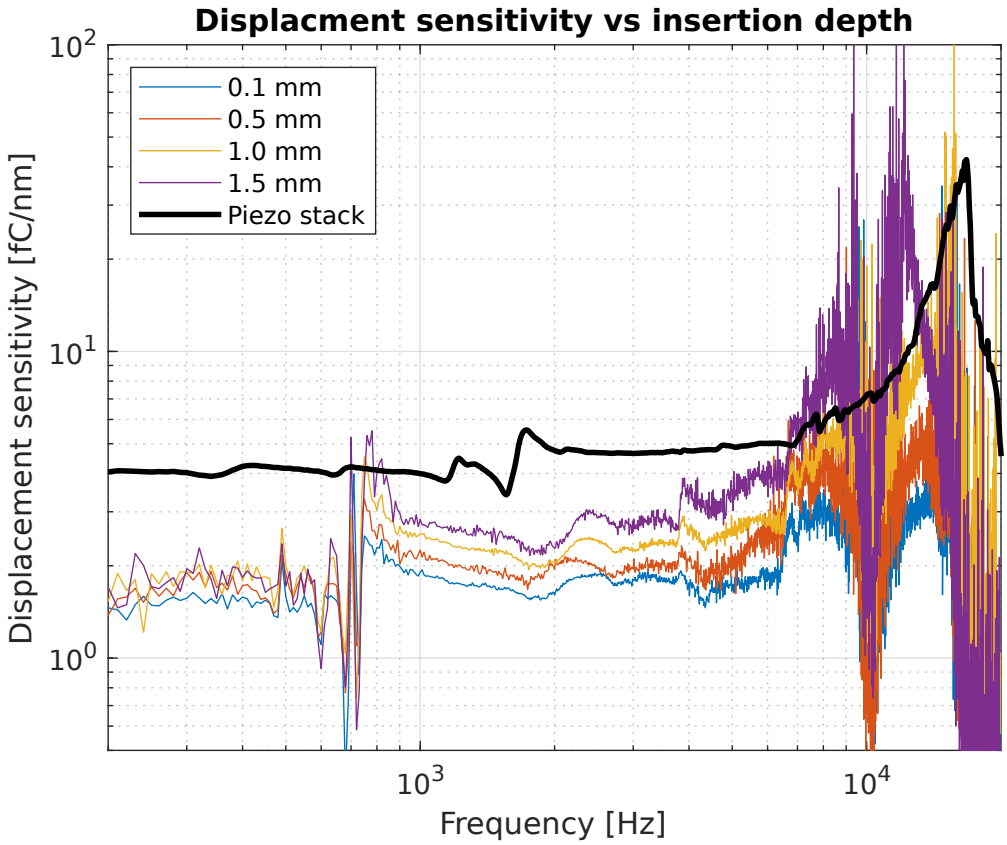

Figure 7-5: Insertion insertion depth increases both stiffness and displacement sensitivity

7.3.2 Static offset

We wanted to measure the impact of static offset on cantilever performance. For the drum-mic this was a significant concern as the sensitivity and mechanical stiffness varied hugely with static offset. Ben Cary's thesis demonstrates that a drum deformation of as little as 6 µm can change the sensitivity by several decibels [5]. Since the drum-mic is several times stiffer than the eardrum, this translates to a static offset of several tens of microns. In contrast, the cantilever-mic showed very little change in either sensitivity or mechanical stiffness in response to static deformation. Figure 7-6 shows an almost constant pressure sensitivity up to an offset of around 300 µm and Figure 7-7 demonstrates that static offset has virtually no effect on cantilever displacement sensitivity whatsoever.

It appears as though most of the variation in sensitivity caused by static offset is actually caused by deformation of the ear drum and ossicular chain in response to

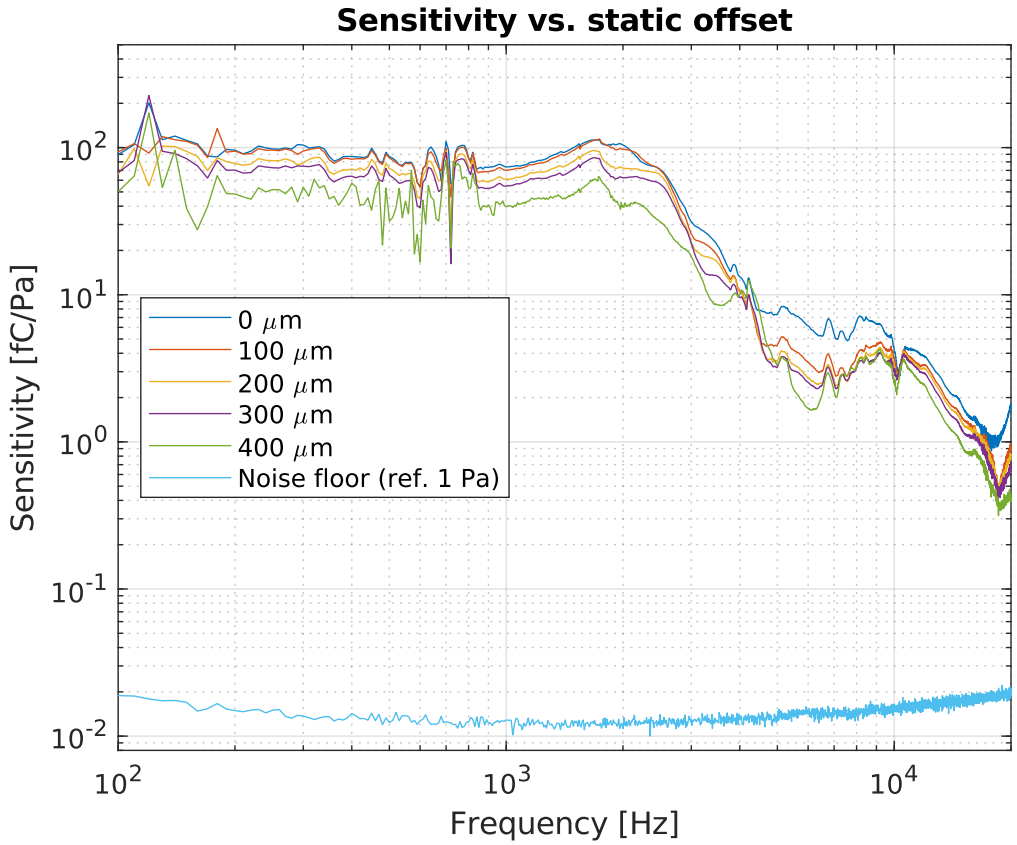

Figure 7-6: Pressure sensitivity vs. static offset. It appears as though up to around 300 µm of static offset has little impact on sensitivity, reducing the precision needed to place the cantilever-mic. Static offset is referenced to marginal umbo contact.

excessive static offset. Therefore there is a large window of acceptable static offset, reducing the precision required to effectively implant the cantilever-mic.

7.4 Comparison to other microphones

The cantilever-mic and differential charge amplifier has higher sensitivity, lower noise floor, and lower EMI than either the drum-mic or the coch-mic. The high sensitivity is attributable to the consistent and controllable mechanical impedance match with the eardrum and concentrating most of the device's elastic strain in the PVDF parallel to d_{31}. The charge sense electrodes in contact with the PVDF are the minimum size necessary and are sealed inside the cantilever—the electrodes' small size reduces parasitic capacitance, and sealing the electrodes away from moisture greatly reduces

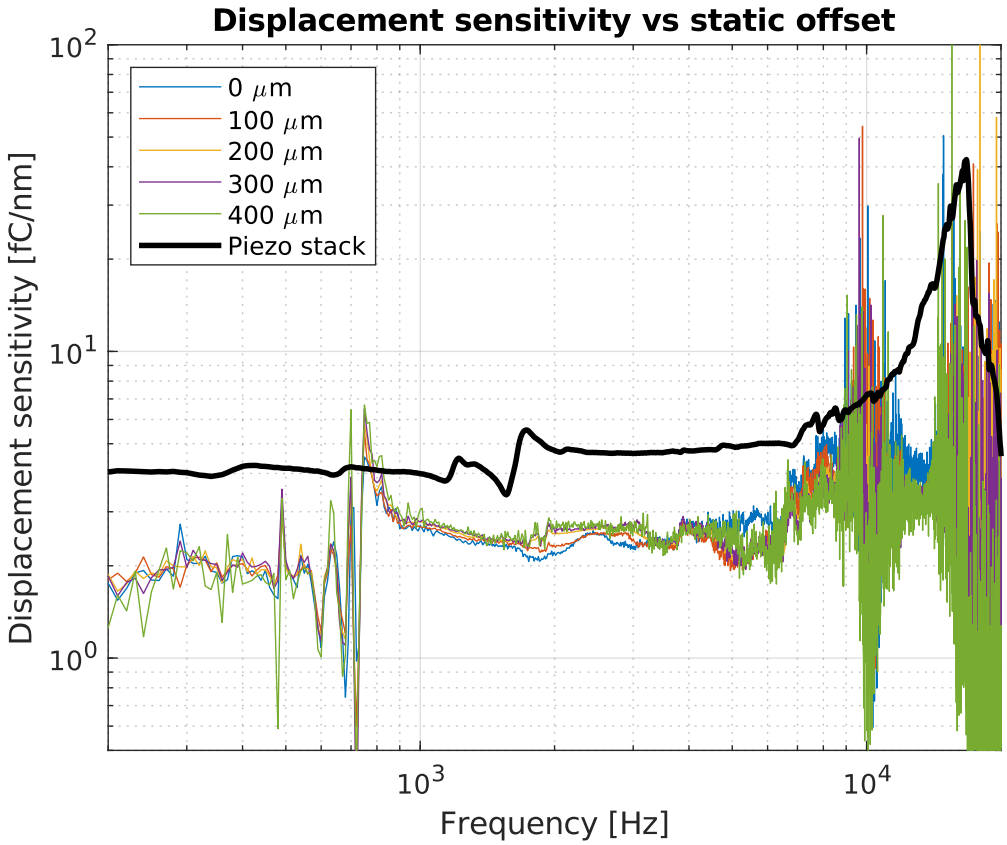

Figure 7-7: The displacement sensitivity is virtually constant across static offset.

leakage conductance. These parasitics have a significant impact on the noise floor. The ground shield and differential output of the cantilever-mic also offer significant EMI reduction. We show here the cantilever-mic's performance compared to our lab's previous designs—the drum-mic and the coch-mic. The drum-mic directly detects umbo motion while the coch-mic is implanted inside the scala tympani.

In Figure 7-8, all three noise curves were taken with the same amplifier, although the drum-mic and coch-mic are single-ended devices and were only connected to one terminal of the differential amplifier. The drum-mic and the coch-mic both have substantially more PVDF area than the cantilever-mic, leading to more noise at higher frequencies from a higher device capacitance. The coch-mic is coated in silicone and has low leakage conductance even when immersed in saline—however the drum-mic is not coated and is suffering from high leakage from the moisture in the middle ear cavity, causing excessive low-frequency noise. In both the drum-mic and

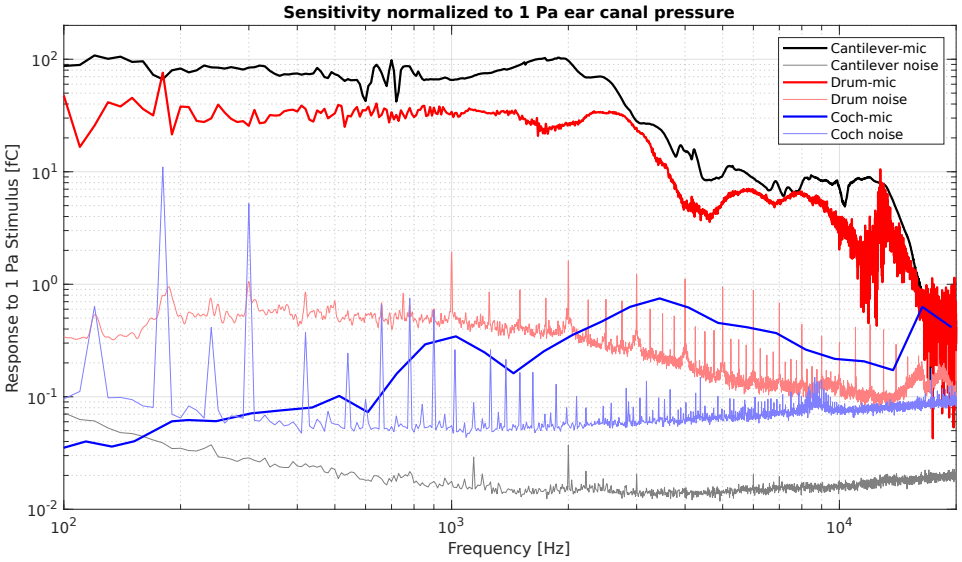

Figure 7-8: Comparison between the cantilever-mic, drum-mic, and coch-mic. Noise floor is normalized to a 1/3–octave bandwidth.

coch-mic noise floors, we can see sharp peaks that appear to be from external electrical stimuli, especially at mains frequency and harmonics. These peaks are not present in the cantilever-mic's noise floor. This difference in EMI susceptibility is especially stark when comparing the measured EMI coupling capacitance between the cantilever-mic and the drum-mic, as the drum's EMI capacitance is measured to be 300 to 400 fF—over 500 times higher than the cantilever's EMI capacitance of 0.6 fF, as shown in Figure 7-9. It is therefore apparent that shielding is required for a practical implantable piezoelectric microphone. While the cantilever-mic has this shielding intrinsic to its design, the drum-mic and coch-mic require substantial redesign to add shielding without introducing parasitic capacitance.

Figure 7-10 shows the cantilever-mic in comparison to the Sonion 65GC31 electret microphone used in commercial cochlear implants. The cantilever-mic stacks up well, although there is still substantial room for improvement. One caveat of this data is that the cantilever-mic EIN is referenced to *ear canal pressure* instead of free-field sound pressure. At 1 kHz and above, the pinna acts like a horn and amplifies the ear canal pressure by up to 20 dB, thereby boosting the cantilever-mic's sensitivity at these frequencies. We would thus expect its EIN to be reduced by 10 to 20 dB

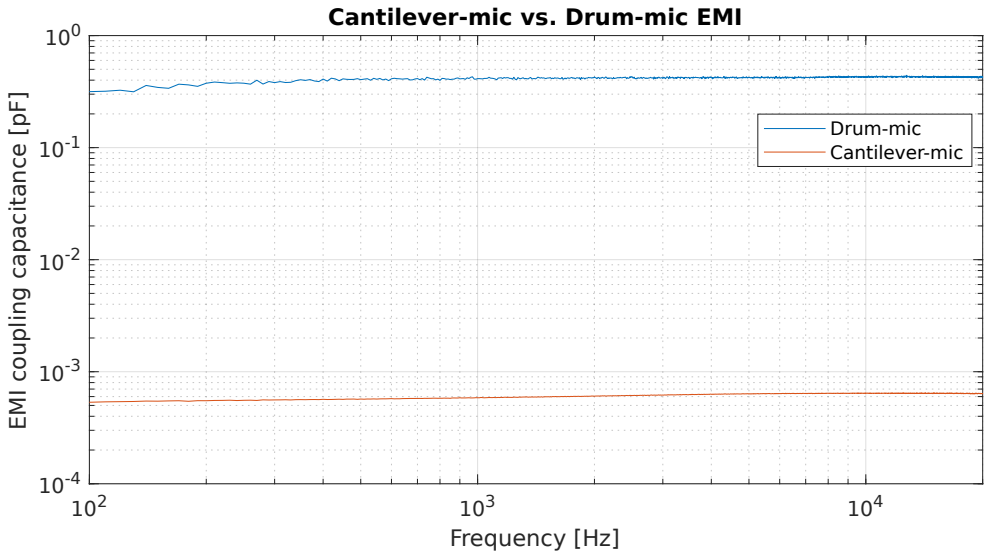

Figure 7-9: The drum-mic's EMI capacitance is over 500 times that of the cantilever-mic.

between 2 kHz and 8 kHz [17], bringing the cantilever-mic performance much closer to that of the Sonion device.

7.5 Summary

We determined that the cantilever-mic and amplifier maintain a consistently low noise charge of about 0.062 fC when implanted, indicating low parasitic capacitance and effective waterproofing preventing leakage current. Furthermore we measured exceptionally low EMI sensitivity of 0.6 fF—over 500 times lower than that of our lab's previous PVDF microphone designs. Our pressure sensitivity was also quite high, measuring 80–100 fC/Pa referenced to ear canal pressure below 2 kHz, falling off to 8–10 fC/Pa above 4 kHz. This signal to noise ratio puts us within striking range of commercial hearing aid microphones, especially when pressure gain provided by the pinna above 1 kHz is considered.

This microphone design still has substantial room for improvement. Figures 7-5 and 7-7 show that the cantilever-mic's implanted displacement sensitivity is substantially lower than when measured on a lab bench with a piezo stack. This indicates that the current method for fixturing the cantilever is insufficient, and highlights the

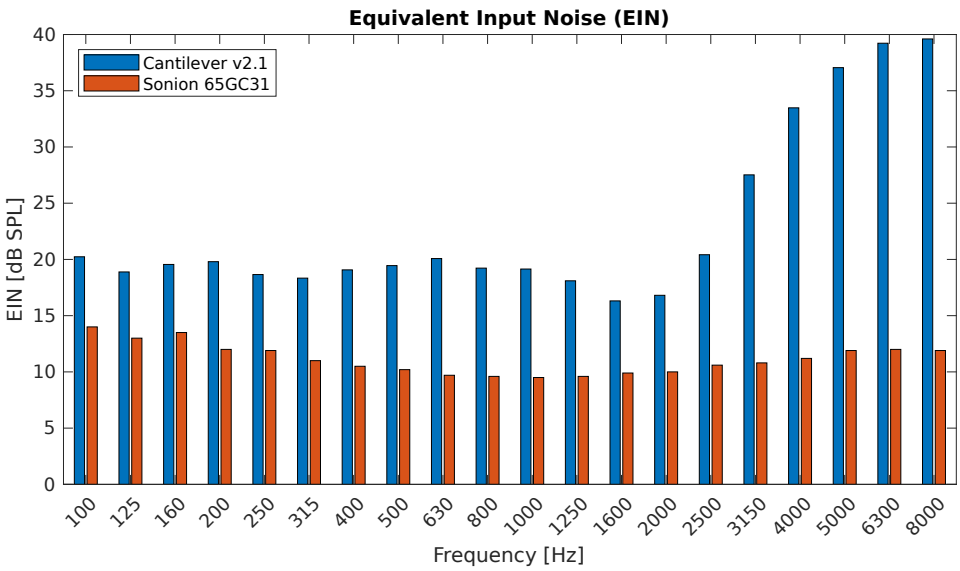

Figure 7-10: Equivalent input noise of the cantilever-mic next to a high quality hearing aid microphone—the Sonion 65GC31.

importance of developing a rigid surgical support. We also need to take measurements of the cantilever-mic pressure sensitivity with a simulated pinna in order to reinforce our claim that the pinna improves high-frequency performance.

Chapter 8

Discussion

8.1 Summary of work

From February 2021 to May 2022 we successfully designed, built, and tested a functional implantable PVDF microphone as well as a very low noise preamplifier. Chapter 1 describes the designs that formed the foundation of the cantilever-mic, Chapter 2 describes the fabrication process in detail, especially the process for layering PVDF. Appendix A provides much greater detail on the multiple failures to effectively layer PVDF and the myriad of failed designs leading up to the current form of the cantilever-mic. We explore the criteria for optimizing implantable microphone design in Chapter 3 and apply these principles to the cantilever-mic. Chapter 4 is relatively self-contained, outlining the design of the differential charge amplifier as well as a comparison to other charge amplifiers both commercially available and in literature. In Chapter 5 we outline the requirements for implantation and provide some of our early attempts to build an adjustable rigid platform for holding the cantilever-mic under the umbo. Chapter 6 describes the measurement techniques used to characterize our microphone and finally, these measurement results are detailed in Chapter 7.

Overall our cantilever-mic design showed performance comparable to a commercial hearing aid microphone and implantation was demonstrated to be feasible. This design provided higher sensitivity and lower noise than any design our lab had previously tested, with only 385 e^- (0.062 fC) of charge noise and 80–100 fC/Pa sensitivity below

2 kHz. The noise floor of our device is equivalent to only 0.015 nm of RMS deflection. We also demonstrated orders of magnitude less electromagnetic interference than our previous designs by building a differential output and a complete ground shield into the design, with no unshielded cables hanging off acting as antennas, achieving an EMI coupling capacitance of 0.6 fF compared to the drum-mic's 300–400 fF. This low susceptibility to EMI allows the cantilever-mic to tolerate electrical noise from mains power, switched-mode power supplies, as well as from a nearby cochlear implant.

8.2 What we learned

This section tries to provide a compreshensive summary of the key features and pieces of information we learned while building the cantilever-mic and amplifier. These findings encompass the whole project from esoteric fabrication tips to fundamentals of piezoelectric sensor design.

8.2.1 Core piezoelectric sensor design principles

In Chapter 3 we establish the rationale for the shape, dimensions, location, and construction of the cantilever-mic from the perspective of increasing sensitivity and reducing noise. We establish that tracking umbo motion has the largest available acoustic energy as well as more lenient design constraints than inside the cochlea. This enables the mechanical impedance of the sensor to be tuned to match that of the eardrum and attached ossicles, maximizing transfer from acoustic energy to elastic strain energy. Chapter 3 shows that elastic strain in the PVDF should be distributed as evenly as possible parallel to the piezoelectric axis and both Chapters 3 and 4 demonstrate the importance of reducing parallel parasitic capacitance. Therefore we chose a triangular cantilever to ensure a relatively even stress distribution and chose a stackup of PVDF layers sandwiching a Kapton core with thickness ratio 1:2:1 to optimize for high charge out with minimum parasitic capacitance. The rest of the sensor was also designed to avoid unnecessary capacitance—we avoid unnecessary PVDF in the sensor and use short cables between the sensor and amplifier. The dimensions

of the cantilever are chosen to achieve a rough mechanical impedance match with the umbo—this is possible since the mechanical impedance of the cantilever is linear and controllable. Finally, the cantilever achieves a rough electrical noise impedance match with the amplifier. These factors all contribute to high sensitivity.

8.2.2 Fabrication quirks

The cantilever-mic solves the connectorization problems and EMI sensitivity that plagued our lab's previous microphone designs. Chapter 2 details the novel process for layering PVDF onto an inert Kapton flex PCB substrate. We demonstrate a strong mechanical bond and suitably tough electrical connection by capacitively coupling the PVDF to a sputter-coated chrome or aluminum electrode on the substrate through a microns-thin glue layer. We sputter-coat the substrate with 200 nm of aluminum or chrome—these metals are highly reactive and bond strongly to the substrate. After this, the substrate is spin-coated on both sides with AZ3312 positive photoresist and exposed to UV light through a contact photolithograpny mask. The UV-exposed photoresist is dissolved in a TMAH solution, leaving the underlying metal to be etched away.

Gluing the layers together is accomplished by masking off portions of the substrate to remain glue-free, aligning the PVDF so that d_{31} is parallel to the strain induced by bending, and taping the PVDF in place on both sides of the substrate. To mark the poling of the PVDF, we typically cut the top right corner off of each rectangular PVDF piece—this helps keep track of which side is up, since pen rubs off very easily. Devcon Plastic Steel epoxy is added between the layers and then squeezed out as much as possible with a doctor blade. After the sensors are glued, they are cut to shape, leaving a small gap between the edge of the electrode and the edge of the sensor, and then they are sputter-coated with a ground shield.

Surface quality is important to achieving strong adhesion between substrate and sputter-coated layer. The ideal surface is very clean and microscopically rough. Prior to sputter-coating the sensors we run a one minute oxygen substrate etch to promote adhesion. We use 3000-grit sandpaper to roughen the PVDF before gluing.

Steps where layers start and stop are one of the biggest sources of problems in the fabrication process. We found that sputter-coated aluminum had poor step coverage, and so used chromium, even though aluminum is easier to etch and has significantly better electrical conductivity. Future development should focus on eliminating or mitigating the problems that result from 90° steps, including fragile sputter-coated layers and uneven photoresist layer thickness.

In general, the sputter-coated layers are resilient to normal handling and implantation. However the coating is very easily scratched or abraded especially by metal objects. A thin parylene or silicone layer encapsulating the sensor would help protect the surface against abrasion and oxidation.

8.2.3 Charge amplifier design ideas

Chapter 4 outlines a few important design principles about low noise amplifier design for very high impedance sensors like our cantilever. Parasitic capacitance has a direct effect on charge amplifier noise—this capacitance interacts with the amplifier's input-referred voltage noise to produce additional charge noise. This charge noise typically manifests as white noise across the bandwidth of the amplifier, and as little as 5 pF of additional capacitance can have a noticeable effect on the noise floor at frequencies above 1 kHz. Parasitic leakage resistance generates substantial Johnson current noise, which manifests as charge noise with a -20 dB/decade frequency characteristic. Leakage resistance as high as 10 GΩ can have a noticeable impact on the noise floor at low frequencies.

The feedback capacitors of the charge amplifier also contribute to the charge noise, so making these capacitors as small as possible is beneficial to performance. We found that 1 pF capacitors between PCB interconnect in a four-layer PCB are highly practical and effective. Similarly, we choose 10 GΩ feedback resistors to reduce low frequency noise charge from Johnson noise current. Surface contamination on the amplifier, especially from solder flux under the U.FL connectors is a major problem—great care should be taken to solder the high-impedance front end of the board with minimal flux, especially the U.FL connectors. In newer versions of the charge amplifier

we added 1 nF decoupling capacitors between the inputs and the first-stage op amp to give more flexibility grounding the sensor.

8.2.4 Implantation constraints

We evaluate the practicality of implanting the cantilever-mic in Chapter 5 and found that surgical implantation is likely feasible. A practical device for holding the cantilever-mic under the umbo must be mechanically rigid—otherwise the supporting structure will bend instead of the piezoelectric material in the sensor. We are targeting a stiffness several times the roughly 1 N/mm stiffness of the cantilever, effectively dictating that the supporting structure is made of a strong biocompatible metal like titanium. The anatomy of the mastoid cavity, facial recess, and middle ear is highly variable, and the surgeon must be able to adjust the fit of the supporting structure to the anatomy. We succeeded in designing a 3D-printed plastic supporting structure prototype with a lockable ball joint that fit the majority of temporal bones we tested it in. Further surgical flexibility can be achieved by having a small selection of supporting structure components with different shape, allowing the surgeon to pick the components that most closely match the patient's anatomy. Once we have built a satisfactory plastic prototype we plan to 3D-print the supporting structure components in titanium.

8.2.5 Metric selection

To compare the cantilever-mic to our lab's previous microphone designs and to microphones in the outside world, we need a standard set of metrics. These metrics must be reliably measurable and sufficiently universal to draw useful comparisons. For instance, EMI coupling capacitance is a useful metric as it offers a good way to characterize the effect of external electric fields on a sensor and compare the response of different sensors to a changing external potential. We choose to use sensor charge output as a sensitivity metric for comparing our different piezoelectric microphone designs. Charge output is unaffected by leakage resistance and capacitance from variable experimental conditions, and it can be accurately measured with a charge

amplifier. Noise floor is measured by recording several seconds of amplifier output without stimulus and is reported as input-referred noise charge. For comparison with sensitivity, the noise floor is normalized to 1/3–octave bandwidth.

Our most useful data comes from composite metrics, especially from pressure sensitivity. We use charge output relative to ear canal pressure close to the umbo since both of these signals are straightforward to measure. Displacement sensitivity, defined as charge out per unit umbo displacement, can be measured both implanted with laser doppler vibrometry (LDV) and on the lab bench with a piezo stack. The most universal composite metric for microphone performance is equivalent input noise (EIN), which is defined as the ratio of noise floor (normalized to 1/3–octave bandwidth) to pressure sensitivity. The EIN thereby represents quietest sound distinguishable from background noise at that particular frequency.

8.3 Avenues for improvement

The differential cantilever PVDF microphone design improved performance substantially over our previous PVDF microphone designs and solved multiple design issues. However there is still a good deal of work to be done to improve the design and make it practically implantable. The first major design challenge we have yet to solve is the rigid fixture for the microphone. As discussed in Chapter 5, we have built several prototype designs for the sensor, including a promising ball and socket-based design. We intend to develop both the hardware and practical surgical procedure necessary to implant a cantilever-mic with requisite mechanical stiffness.

The cantilever-mic is not suitable for implantation in its current form for several reasons. Materials would have to be carefully chosen to ensure biocompatibility without compromising electrical and mechanical performance. Fortunately, PVDF is biologically inert, but materials in the flex PCB base like the solder mask or copper traces are likely not acceptable. The aluminum or chromium ground shield would also likely have to be replaced or at least coated with titanium, which has substantially poorer electrical conductivity than either of the two original metals.

The current charge amplifier design is optimized for laboratory bench testing and would need to be retooled for implantation, especially in terms of size and power consumption. While 15 mW of quiescent power consumption may seem small, it is gigantic compared to the 20 µW of the Sonion 65GC31. One solution is to replace the charge amplifier with a differential JFET-based voltage amplifier similar to the internal JFET amplifier within most commercial electret microphones. This solution is suitable for a microphone application where precise measurement of charge output is not required, and simulation demonstrates noise floor similar to the charge amplifier with under 50 µW of power consumption. Another possible solution is to design a custom charge amplifier IC similar to [13] and optimize for low power consumption. Either a JFET amplifer or a charge amplifier IC could be directly integrated with the cantilever-mic's structure, further reducing noise floor by cutting down on parasitic capacitance from cabling.

8.4 Conclusion

Overall we are quite happy with the differential cantilever-mic design and charge amplifier. The combination of these designs offers compelling noise performance as well as insensitivity to unwanted electromagnetic interference. The cantilever-mic also shows tolerance to variation in surgical implantation, so long as the surgical hardware is sufficiently rigid. We think this work represents an exciting development in implantable microphone technology, and we hope to see it in use someday, improving the lives of deaf and hearing-impaired individuals.

Appendix A

Evolution of design and fabrication

A.1 Introduction

The sensor blueprint described in Chapter 2 is the direct result of very many failed designs and fabrication processes. We felt that while these failures did not necessarily fit into the main body of the thesis, they were nevertheless important to include both as a testament to effort expended and to help future readers interested in replicating our results.

A.2 Motivation for multiple layers

A efficient piezoelectric sensor must achieve a good impedance match between the sensor and the surrounding biological structures. The drum microphone achieves a rough impedance match through its shape—a small displacement force generates a large tensile force in the drum, with this mechanical advantage roughly proportional to the ratio of the drum diameter to the total static offset. This variable mechanical advantage means that the drum microphone must have a carefully controlled nonzero static offset to function, making a flat drum design unsuitable for applications where this static offset cannot be effectively controlled.

A.2.1 Evolution of the layered drum

Although I eventually settled on a cantilever sensor for a number of reasons, the path I took to reach that conclusion was anything but straightforward. My initial goal was to solve the impedance matching problem inherent to the intracochlear microphone by developing a sealed bending–mode drum hydrophone. The initial motivation for a layered design was to produce a drum that would operate with zero static offset, rendering it suitable for a hydrophone. The concept of a drum hydrophone worked reasonably well in simulation, but it had several glaring issues. Firstly, the drum head would have to be extraordinarily small—most likely less than 1.5 mm in diameter in order to fit through the round window. While this constraint alone made the device borderline impractical, the connection between the drum body and head had to be flexible. This was achievable in simulation by using a specially shaped silicone gasket, but fabricating and gluing such a gasket would be obnoxiously difficult. Even if these fabrication constraints could be overcome, the end result would be less than spectacular. Achieving a good impedance match with the cochlea would require a larger area of drum surface and the amount of acoustic energy available in the scala tympani is not particularly impressive if the round window is not blocked.

The layered drum is not a compelling replacement for Ben's drum design. It adds additional complexity without solving the fundamental issue of linearity, although it can be operated with very little static offset where mechanical impedance is more linear. The problem of nonlinear mechanical impedance is inescapable for drums—the nature of the boundary conditions causes the radius of curvature of the drum to change. If the structure is designed such that it can bend along a single axis, this nonlinearity disappears. I developed the idea of an arch-like design that could be formed by folding and layering PVDF. While this process was certainly more feasible than making a drum hydrophone, it still presented difficulties. Making precise folds would require jigs, and attaching leads was still an unresolved problem. The arch design may well have been possible to build, but it was not the clean solution I was looking for.

A.2.2 Cantilevers and cables

Cantilevers are quite possibly the simplest and most straightforward form of bending–mode transducer. Unlike drum designs, they do not require a flexible boundary condition. I had intentionally avoided cantilevers because of the stability problems with the diving board design, but it became apparent that the diving board's problems were mostly a consequence of its mounting system. The diving board was anchored to the skull nearly a centimeter away from the umbo—this meant that the diving board itself was long and narrow. This form factor gave the diving board an unfortunate tendency to twist and fall off of the umbo. The mechanical impedance was also far too low, although this problem was partially fixed by gluing a metal strip to the back of the device to stiffen it.

An ideal cantilever would be relatively wide and short to present a good acoustic impedance match with the ear without twisting excessively. Making the cantilever triangular would give a relatively uniform stress distribution along the length of the cantilever, increasing the device efficiency. Such a design could be implanted by installing a flat platform between the skull and the promontory and mounting the cantilever to that platform. This implantation scheme is not overly complicated and allows relatively easy adjustment of the cantilever's exact fit in the ear.

Throughout this discussion we have avoided the surprisingly thorny issue of how to electrically connect the sensor to the amplifier. The connection needs to be mechanically robust, electrically shielded, and waterproof without introducing significant parasitic capacitance or leakage resistance, and designing an effective interface is critical for making a practical implantable sensor.

Our lab had previously attached ribbon cables to PVDF sheets by using a hot bar bonder, but this method leaves exposed ribbon cables that effectively act like antennas for electromagnetic interference. The cantilever sensor provides an elegant way to combine the mechanical design of the sensor with an electrical connection. By layering PVDF onto a flexible polyimide printed circuit board, the flex-PCB becomes an integral part of the cantilever and gives a convenient place to attach connectors

and cables.

A.3 Layered cantilever fabrication

The layered cantilever offers several advantages over previously discussed designs. It has mechanical impedance and sensitivity that are both controllable and linear. The sensitive charge-collecting inner electrodes are sealed from water ingress and completely shielded from electromagnetic interference, and the differential output further reduces EMI. The integration into a flex-PCB allows for cables or connectors to be soldered directly to the sensor without damaging the heat-sensitive PVDF.

These design advantages are of course conditional on the ability to fabricate such a sensor. In particular, we needed a mechanically shear-resistant adhesive layer between the PVDF and the flex-PCB that could effectively transfer charge from the piezoelectric to the amplifier.

A.3.1 Conductive adhesive tapes: disappointing shear strength

Since we had previously used a hot bar bonder to attach leads to PVDF sensors, we reasoned that a similar technique could be applied to produce layered sensors. Hot–bar bonders are used widely in the electronics industry, especially for attaching fine–pitch connectors on components like LCD panels. In the hot–bar bonding process, a heat–sensitive epoxy tape containing conductive micro–beads is placed between the two surfaces, and then the bonder applies pressure at 140 to 200 °C to cure the tape and produce a relatively permanent bond. Unfortunately PVDF depoles well below this temperature. While hot bar bonding works fine for connecting wires to depoled regions of PVDF, it is ineffective for layering poled PVDF.

There exist tapes with a weaker pressure sensitive adhesive that use a similar principle to the hot bar bonder. 3M Z–Axis tape (9703) uses a similar principle to hot bar bonder tape, but it does not require heat to set. This tape promised to greatly simplify the layering process—no mucking around with adhesives necessary! The tape was used both to layer the PVDF and to attach the sensor to the flex-PCB.

Tragically it became clear that the tape was not up to the task.

While the tape performed its primary role of properly electrically connecting the system, it failed to prevent the layers of PVDF from shearing past each other. When a single layer of PVDF was deflected and released, a faint audible snap could be heard, but the taped device exhibited a sluggish response indicative of significant damping. Likewise the sensor exhibited a disappointingly small output. It was evident that the tape was allowing the PVDF layers to shear past each other. Testing some scrap PVDF with generic double sided tape yielded a similar damped response—it seems as though pressure–sensitive adhesives do not stick particularly well to PVDF. Back to the drawing board…

A.3.2 Conductive epoxy: adventures in delamination

If conductive tape did not work, surely a conductive adhesive would function as desired. Such an adhesive would perform triple duty as an electrode for the PVDF, an electrical connection to the flex-PCB, and a mechanical connection capable of resisting shear stress between the polyimide base and the PVDF.

The problems began almost immediately. Most conductive epoxies, such as the MG Chemicals 8330S silver epoxy paste that we used, consist of conductive particles suspended in a matrix. This means that a minimum layer thickness is necessary to retain conductivity. However, we also needed the adhesive layer to be as thin as possible—otherwise, the epoxy would add unwanted mechanical impedance. Cutting the epoxy with a few percent acetone was helpful for controlling the viscosity. After some experimentation with a doctor blade to spread the epoxy as desired, the glued sample was cured at 65 °C for several hours. Once the epoxy had fully cured, the sensor was cut to shape with scissors.

The epoxied sensors were sensitive but had several serious flaws. Firstly, the epoxy got absolutely everywhere when putting together the sensor. This was mostly mitigated by trimming the sensor, but some judicious scraping was still required to remove short circuits. The inner electrodes formed by the epoxy were exposed at the edges of the sensor. This created serious problems with parasitic leakage resistances

and leakage currents that ruined the sensor's noise floor. In high humidity, these DC leakage currents were large enough to rail the charge amplifier. Finally, the conductive epoxy delaminated all too easily, even if both surfaces were sanded prior to gluing.

Ultimately, we abandoned conductive adhesives and settled for capacitive coupling as discussed in Chapter 2, and the template for the cantilever-mic was born.

Appendix B

Mechanics of Hearing

The ear consists of three main subsystems, shown in Figure B-1. The outer ear serves to funnel sound into the ear canal, the inner ear transduces sound into neural impulses in the auditory nerve, and the middle ear serves as a mechanical impedance match between the air in the ear canal and the fluid of the cochlea—the portion of the inner ear dedicated to sound processing.

The outer ear consists of the pinna (the visible part of the ear) and the ear canal. The pinna's primary role is to act as an acoustic horn that focuses sound into the ear canal. It also affects the frequency response of incoming sound depending on direction, playing an important role in sound localization. The ear canal transmits sound from the pinna to the eardrum.

The middle ear, shown in Figure B-2 serves as an impedance match between the air in the ear canal and the fluid inside the cochlea. The malleus is attached to the eardrum and the end of the malleus forms the point of the eardrum's conical shape, called the umbo. The malleus transmits acoustic energy from the eardrum through the incus to the stapes, which is connected to the oval window of the cochlea. Since the eardrum is several times larger than the oval window, a volume displacement at one pressure at the eardrum is translated into a smaller volume displacement at higher pressure in the cochlea. This impedance matching allows much more acoustic energy to be transmitted into the cochlea instead of reflecting off of interface between air and cochlear fluid.

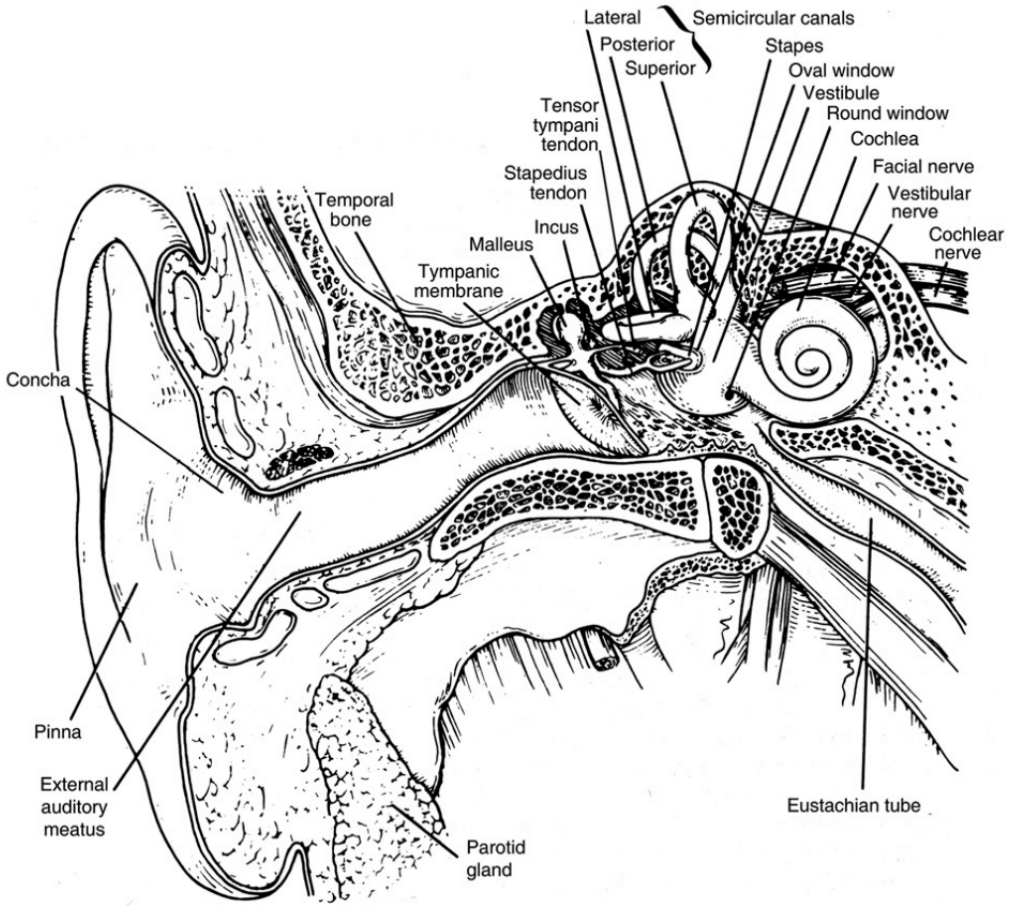

Figure B-1: A cross section of the human ear [17]. The pinna and ear canal act like a horn at high frequencies, causing a pressure gain inside the ear canal.

The inner ear consists of the cochlea, which handles sound processing, and the vestibular system, which detects linear and rotational motion. The cochlea is shaped like a spiral with approximately 2.5 turns, with this spiral separated into the scala vestibuli and the scala tympani by the basilar membrane and the cellular sensory tissue of the organ of Corti. Figure B-3 illustrates a cross-section of the cochlea showing these structures. While the scala tympani is connected to ambient middle ear pressure through the low-impedance round window, the pressure in the scala vestibuli is driven by the ossicular chain through the oval window. This pressure difference causes motion of the basilar membrane, which is detected by the organ of Corti and transduced into neural impulses. Different longitudinal of the basilar membrane have different resonant frequencies—the resonant frequency decreases towards the apex of

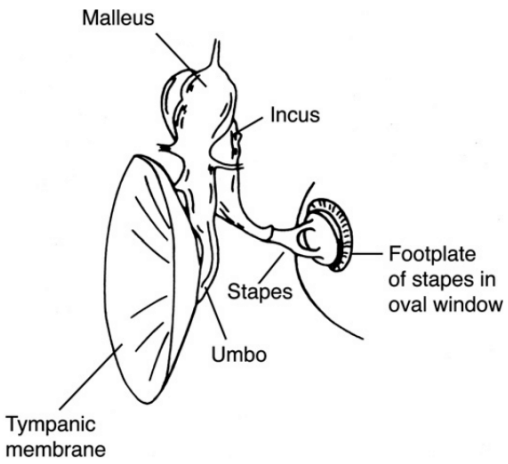

Figure B-2: A diagram of the middle ear from [17]. The oval window is much smaller than the eardrum, causing a pressure gain from eardrum to scala vestibuli.

the spiral (away from the stapes). This is a primary mechanism by which the cochlea can detect differences in pitch.

Cochlear implants take advantage of this spatial layout of pitch sensitivity. The cochlear implant is a flexible electrode array inserted into the scala tympani through the round window or through a hole drilled in the cochlea. The cochlear implant processor maps different frequency ranges to electrical stimulation of particular electrodes along the implant's electrode array. These electrical impulses stimulate the auditory nerve directly, substituting for the mechano-electrical transduction process of the organ of Corti.

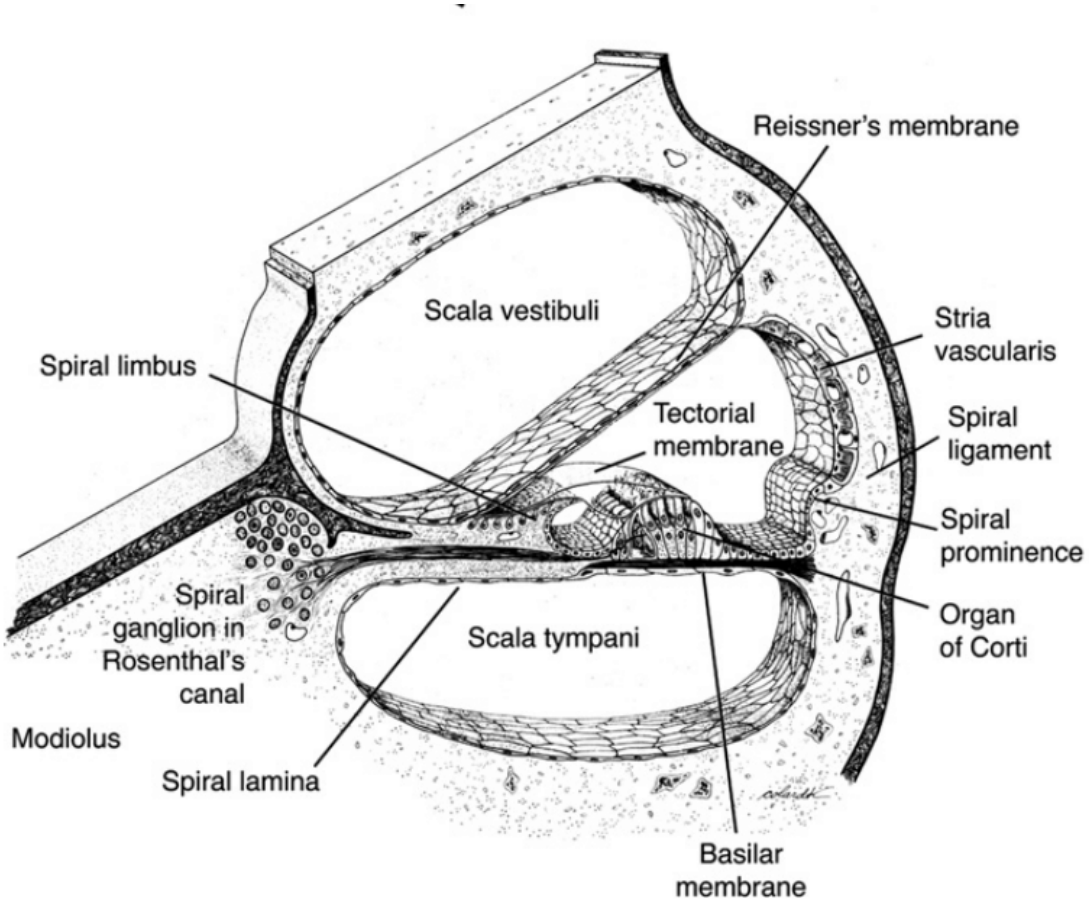

Figure B-3: A cross section of the cochlea from [17]. The basilar membrane vibrates in response to the differential pressure between the scala vestibuli and scala tympani. The organ of Corti transduces this vibration into neural impulses in the auditory nerve.

Bibliography

[1] Emad Alnasser. A Novel Low Output Offset Voltage Charge Amplifier for Piezoelectric Sensors. *IEEE Sensors Journal*, 20(10):5360–5367, May 2020. Conference Name: IEEE Sensors Journal.

[2] Analog Devices, LTC6240/LTC6241/LTC6242. *Single/Dual/Quad 18MHz, Low Noise, Rail-to-Rail Output, CMOS Op Amps*. Rev. 624012fe.

[3] Diego Calero, Stephan Paul, André Gesing, Fabio Alves, and Júlio A. Cordioli. A technical review and evaluation of implantable sensors for hearing devices. *BioMedical Engineering OnLine*, 17(1):23, December 2018.

[4] Lorenzo Capineri, Pietro Giannelli, Giacomo Calabrese, Maurizio Granato, and Giovanni Frattini. Design, Realization and Characterization of a Differential Charge Amplifier for Ultrasonic Piezopolymer Transducers. In *2018 IEEE International Ultrasonics Symposium (IUS)*, pages 1–5, October 2018. ISSN: 1948-5727.

[5] Benjamin Cary. Design of a surgically-viable umbo microphone for implantable assistive hearing devices, June 2021.

[6] CEC Vibration Products, 1-328. *Differential Charge Amplifier*. Rev. 16008.

[7] Jianghua Chen, Xuewen Ni, and Bangxian Mo. A low noise CMOS charge sensitive preamplifier for MEMS capacitive accelerometer readout. In *2007 7th International Conference on ASIC*, pages 490–493, October 2007. ISSN: 2162-755X.

[8] Francis (Pete) X. Creighton, Xiying Guan, Steve Park, Ioannis (John) Kymissis, Hideko Heidi Nakajima, and Elizabeth S. Olson. An Intracochlear Pressure Sensor as a Microphone for a Fully Implantable Cochlear Implant. *Otology & Neurotology*, 37(10):1596–1600, December 2016.

[9] N. Duncan, Z. Sternovsky, E. Grün, S. Auer, M. Horanyi, K. Drake, J. Xie, G. Lawrence, D. Hansen, and H. Le. The Electrostatic Lunar Dust Analyzer (ELDA) for the detection and trajectory measurement of slow-moving dust particles from the lunar surface. *Planetary and Space Science*, 59(13):1446–1454, October 2011.

[10] FEMTO, HQA-15M-10T. *High Frequency Charge Amplifier*, Feb 2012. Rev. 3.

[11] Darcy L. Frear, Xiying Guan, Christof Stieger, John J. Rosowski, and Hideko Heidi Nakajima. Impedances of the inner and middle ear estimated from intracochlear sound pressures in normal human temporal bones. *Hearing Research*, 367:17–31, September 2018.

[12] Y. Hu, J.L. Solere, D. Lachartre, and R. Tutchetta. Design and performance of a low-noise, low-power consumption CMOS charge amplifier for capacitive detectors. *IEEE Transactions on Nuclear Science*, 45(1):119–123, February 1998. Conference Name: IEEE Transactions on Nuclear Science.

[13] S. Kelz, T. Veigel, M. Grözing, and M. Berroth. A Fully Differential Charge-Sensitive Amplifier for Dust-Particle Detectors. In *2018 14th Conference on Ph.D. Research in Microelectronics and Electronics (PRIME)*, pages 13–16, July 2018.

[14] Hideko Heidi Nakajima, Michael E. Ravicz, Saumil N. Merchant, William T. Peake, and John J. Rosowski. Experimental ossicular fixations and the middle ear's response to sound: Evidence for a flexible ossicular chain. *Hearing Research*, 204(1-2):60–77, June 2005.

[15] Steve Park, Xiying Guan, Youngwan Kim, Francis (Pete) X. Creighton, Eric Wei, Ioannis(John) Kymissis, Hideko Heidi Nakajima, and Elizabeth S. Olson. PVDF-Based Piezoelectric Microphone for Sound Detection Inside the Cochlea: Toward Totally Implantable Cochlear Implants. *Trends in Hearing*, 22:233121651877445, January 2018.

[16] Flurin Pfiffner, Lukas Prochazka, Ivo Dobrev, Adrian Dalbert, Jae Hoon Sim, Francesca Harris, Jeremie Guignard, Joris Walraevens, Christof Röösli, and Alex Huber. A MEMS condenser microphone based acoustic receiver for totally implantable cochlear implants. *The Journal of the Acoustical Society of America*, 143(3):1778–1778, March 2018.

[17] James O. Pickles. *An Introduction to the Physiology of Hearing*. Emerald Group Publishing, 2012.

[18] J. Xie, Z. Sternovsky, S Auer, K. Drake, E. Grün, M. Horanyi, H. Le, and R. Srama. Laboratory testing and data analysis of the Electrostatic Lunar Dust Analyzer (ELDA) instrument. *Planetary and Space Science*, 89:63–70, December 2013.