Ultra-low-power Circuits and Systems for Wearable and Implantable Medical Devices

by

Marcus Yip

B.A.Sc., University of Toronto (2007)
S.M., Massachusetts Institute of Technology (2009)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2013

Author .....
Department of Electrical Engineering and Computer Science
August 30, 2013

Certified by .....
Anantha P. Chandrakasan
Joseph F. and Nancy P. Keithley Professor of Electrical Engineering
Thesis Supervisor

Accepted by .....
Leslie A. Kolodziejski
Chair, Department Committee on Graduate Students

Ultra-low-power Circuits and Systems for Wearable and Implantable Medical Devices

by

Marcus Yip

Submitted to the Department of Electrical Engineering and Computer Science
on August 30, 2013, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Advances in circuits, sensors, and energy storage elements have opened up many new possibilities in the health industry. In the area of wearable devices, the miniaturization of electronics has spurred the rapid development of wearable vital signs, activity, and fitness monitors. Maximizing the time between battery recharge places stringent requirements on power consumption by the device. For implantable devices, the situation is exacerbated by the fact that energy storage capacity is limited by volume constraints, and frequent battery replacement via surgery is undesirable. In this case, the design of energy-efficient circuits and systems becomes even more crucial.

This thesis explores the design of energy-efficient circuits and systems for two medical applications. The first half of the thesis focuses on the design and implementation of an ultra-low-power, mixed-signal front-end for a wearable ECG monitor in a 0.18μm CMOS process. A mixed-signal architecture together with analog circuit optimizations enable ultra-low-voltage operation at 0.6V which provides power savings through voltage scaling, and ensures compatibility with state-of-the-art DSPs. The fully-integrated front-end consumes just 2.9μW, which is two orders of magnitude lower than commercially available parts.

The second half of this thesis focuses on ultra-low-power system design and energy-efficient neural stimulation for a proof-of-concept fully-implantable cochlear implant. First, implantable acoustic sensing is demonstrated by sensing the motion of a human cadaveric middle ear with a piezoelectric sensor. Second, alternate energy-efficient electrical stimulation waveforms are investigated to reduce neural stimulation power when compared to the conventional rectangular waveform. The energy-optimal waveform is analyzed using a computational nerve fiber model, and validated with *in-vivo* ECAP recordings in the auditory nerve of two cats and with psychophysical tests in two human cochlear implant users. Preliminary human subject testing shows that charge and energy savings of 20-30% and 15-35% respectively are possible with alternative waveforms. A system-on-chip comprising the sensor interface, reconfigurable sound processor, and arbitrary-waveform neural stimulator is implemented in a 0.18μm high-voltage CMOS process to demonstrate the feasibility of this system.

The sensor interface and sound processor consume just 12μW of power, representing just 2% of the overall system power which is dominated by stimulation. As a result, the energy savings from using alternative stimulation waveforms transfer directly to the system.

Thesis Supervisor: Anantha P. Chandrakasan
Title: Joseph F. and Nancy P. Keithley Professor of Electrical Engineering

Acknowledgments

Deciding to come to MIT was probably one of the best decisions I've ever made. I've had the great fortune to learn from some of the smartest and most down-to-earth people I've ever met, and the experience has brought me a lot of joy and fulfillment. But the best things I will take away from MIT are the relationships that I've built with the wonderful people here, many of whom I can safely say will be lifelong friends and colleagues. MIT and Boston will always have a special place in my heart, and for that, I have too many people to thank.

First, I would like to thank my advisor, Prof. Anantha Chandrakasan, for his guidance, support, and encouragement, and for allowing me to pursue research topics that interested me. Anantha has been a wonderful mentor and role model, and he has taught me how to shape and present ideas, and how to always keep the impact and big picture in mind. No matter how busy Anantha gets, his priority has always been his students, and I am grateful to him for always putting me in the best position to succeed.

I would like to thank my thesis committee members, Prof. Konstantina Stankovic, Prof. David Perreault, and Prof. Heidi Nakajima for all their contributions and advice, and I am grateful for the generosity they've shown with their time and energy. I would especially like to thank Prof. Stankovic who was instrumental in shaping the direction of the CI work right from the beginning. Her enthusiasm and expertise with CIs and the auditory sciences have added so much value to this thesis.

I would like to thank the numerous people at the Massachusetts Eye and Ear Infirmary across the river who have been so open and willing to help with the CI work. I'd like to thank Dr. Ken Hancock and Prof. Bertrand Delgutte for fitting in cat ABR and ECAP measurements for us. I'd like to thank Prof. Don Eddington for providing us with the nerve fiber model. Don and Victor Noel were also extremely generous in providing their time and resources and helping us out with the human subject measurements at the CIRL. I am also grateful to Prof. Heidi Nakajima, without whom the middle ear measurements would not have been possible. She spent countless hours with me in the chambers, always willing to lend a hand with positioning the laser.

When I first arrived at MIT, I was fortunate to have so many people within ananthagroup (and fellow Canadians!) to turn to for help. I'd like to thank Denis, Naveen, Brian, Yogesh, Joyce, Viv, Payam, Nathan, Fred, and Jose as a part of the "older" group that I looked up to and constantly bugged for help when I first arrived. In particular, Joyce Kwong and Vivienne Sze showed me the ropes of being a grad student, and they showed me that it was indeed possible to be laughing in the lab all day, every day. I got to know Brian Ginsburg after he graduated from MIT, and even though we never overlapped as students, he has never hesitated to offer help and advice on my work. Brian's depth and breadth of knowledge on everything circuits is a reminder of how little I know, and he inspires me to keep reading and learning. Fred Chen has helped me with random CAD issues numerous times over the years, often in the wee hours of the morning over gchat. I am still amazed at how much he knows about everything, and often times just off the top of his head. Finally, I'd

like to thank Jose Bohorquez who I had the pleasure of working with during the first generation of the ECG front-end. I am grateful that he trusted me enough to let me help tape out a chip that was in the critical path of his thesis. I am indebted to him for showing me how to better think of and analyze circuits, for showing me Cadence tricks that I had no idea existed, for introducing me to Opal Kelly, and mostly, for being a great mentor.

I would also like to thank present and recent members of ananthagroup for making it a great place to come to work, and in particular, Patrick, Masood, Mahmut, Kailiang, Saurav, Rahul, Bonnie, Yildiz, Phil, Arun, and Rui. Even though I came a year after Patrick, Masood, and Mahmut, they took me under their wings and paved the way for me, showing me exactly what I needed to do, and when I needed to do it by. I'd especially like to thank Rui Jin, who I had the pleasure of working with during his UROP and Super-UROP with the group. Even though I was supposed to be his mentor, I'm pretty sure I learned more from him than the other way around. I am also grateful to Rui for his efforts with the nerve fiber simulations and stimulator circuit design which contributed significantly to this thesis.

I would like to thank Michael Georgas for our coffee breaks, over which we've shared successes and failures, celebrated and commiserated. It has been a great ride, since the open house in 2007, and I'm happy that we'll be able to complete our theses within months of each other.

I also want to thank Margaret Flaherty for keeping the lab running so smoothly over the years. There has not been a single socket, PCB, chip, or Digi-Key order that has not come or gone through Margaret, and she has taken away so much stress by taking care of all the little things.

This thesis would not have been possible without the funding over the years from Texas Instruments and NSERC. I'd also like to thank the TSMC University Shuttle Program and MOSIS for chip fabrication support.

Boston is such a special place to me because of all the friends that I've made here. In no particular order (other than from the recipient list of our e-mail chains), George, Jackie, Andy, Jason, Daniel, Jon, Mike, Doris, Melinda, Courtney, Keith, Cathy, Cody, Fred, Jes, thanks for the support, dancing, and fun times over the years. I am also lucky to have a great group of friends to go home to in Toronto: Bryan, Edward, Alan, Cynthia, Anita, Hannah, and Polly, thanks for always welcoming me back like I never left. I could not ask for a better group of friends.

This thesis would not have been possible without Erica. I'm lucky to have Erica in my life to support me and cheer me on, and I am grateful for her patience and sacrifice. She has been with me every step of the way over the last couple years, and she probably knows more about tapeouts than any dentist should. I'd also like to thank Erica's parents for feeding me Chinese food from the best little restaurant.

Finally, I'd like to thank my sister Martha, my grandma Po-po, and my Mom and Dad for their love and support, without which I would not be here today.

Contents

List of Figures

List of Tables

List of Acronyms

ADC	analog-to-digital converter
AFE	analog front-end
ASIC	application-specific integrated circuit
BJT	bipolar junction transistor
BPF	band-pass filter
CA	charge amplifier
CIC	cascaded integrator-comb
CIS	continuous interleaved sampling
CI	cochlear implant
CMFB	common-mode feedback
CMOS	complementary metal-oxide semiconductor
CVD	cardiovascular disease
DAC	digital-to-analog converter
DDS	direct-digital synthesizer
DNL	differential non-linearity
DSP	digital signal processor or digital signal processing
ECAP	electrically-evoked compound action potential
ECG	electrocardiogram
ENOB	effective-number-of-bits
FFT	fast Fourier transform
FICI	fully-implantable cochlear implant
FIR	finite impulse response
FOM	figure-of-merit
FPGA	field-programmable gate array
GA	genetic algorithm

HPF	high-pass filter
IIR	infinite impulse response
INL	integral non-linearity
LDV	laser Doppler vibrometry
LNA	low-noise amplifier
MCL	most-comfortable-level
MEI	middle ear implant
MEMS	micro-electro-mechanical systems
MSFE	mixed-signal front-end
NEF	noise efficiency factor
NMOS	n metal-oxide semiconductor
OSR	oversampling ratio
PCB	printed circuit board
PEF	power efficiency factor
PGA	programmable-gain amplifier
PLI	power-line interference
PMOS	p metal-oxide semiconductor
PSRR	power supply rejection ratio
PW	phase width
PZFE	piezoelectric sensor front-end
QFP	quad flat package
SAAF	*SINC* anti-aliasing filter
SAR	successive approximation register
SFDR	spurious-free dynamic range
SNDR	signal-to-noise and distortion ratio
SoC	system-on-chip
SPL	sound pressure level
THD	total harmonic distortion
THR	threshold

Chapter 1

Introduction

Advances in electronics have opened up many new possibilities in health care. In the realm of *wearable* medical devices and activity monitors, small and non-invasive devices with rich sensing, processing, and communication capabilities are an enabling technology for *connected health*, which aims to shift health care from the hospital toward the home [3, 4]. Wearable sensor nodes can be used for long-term health or lifestyle monitoring, where the focus is on *prevention* rather than cure. As the size and cost of semiconductor memory decrease with technology scaling, sensor nodes can store an increasing amount of data which the user or a physician can later analyze. Alternatively, wearable sensor nodes can also leverage energy-efficient digital signal processing to perform computations locally at the node [5, 6]. By extracting meaningful features from raw data (for example, detecting arrhythmias from an electrocardiogram), the amount of data transmitted and power consumed by the last-meter wireless link (often the most power-hungry component in a sensor node [7–11]) can be significantly reduced.

Through the miniaturization of electronics, devices can also be *implanted* inside the body leading to a breadth of other applications such as neural prostheses or implantable cardioverter-defibrillators (ICD). Examples of neural prostheses include retinal implants which bypass damaged photoreceptor cells in the eye to directly stimulate the retinal neurons [12], or cochlear implants (CIs) which bypass damaged hair cells in the cochlea to directly stimulate the auditory nerve [13]. An ICD is an

implantable battery-powered device used in patients who are at risk of sudden cardiac death. It detects cardiac arrhythmias and delivers electric shocks to correct them. Its battery life is typically 5 to 7 years, at which point surgery is required to replace the device. Therefore, ultra-low-power electronics can make a big impact by extending the time between surgeries. Going beyond implantables, devices like an ingestible sensor from Proteus Digital Health can be swallowed to monitor patient medication compliance [14].

Regardless of the usage scenario, wearable or implantable medical devices all have similar design requirements such as maximizing functionality, intelligence, and lifetime, while minimizing size and cost. The miniaturization of wearable devices is driven by factors such as aesthetics, wearability, and comfort, whereas the size of implantable devices is constrained by anatomy. Ultimately, the size constraint combined with the limited energy density of today's batteries necessitate the design of ultra-low-power circuits and systems which is the main focus of this thesis. In order to continue the recent advances in medical device technology, innovations have to be made at all levels of the design, from physiology, to sensors and materials, to algorithms and architectures, and of course, the underlying circuit subsystems.

This thesis examines the design of ultra-low-power circuits and systems for two different medical applications. The first is the design of a front-end system for a *wearable* electrocardiogram monitor, while the second is the design of a fully-*implantable* cochlear prosthesis. The rest of this chapter provides background that will provide context for the rest of the thesis. Section 1.1 discusses the general requirements of wearable monitoring devices, with particular emphasis on cardiac monitoring for cardiovascular disease. Section 1.2 provides background on cochlear implants which are used to restore hearing to individuals who are profoundly deaf. Section 1.3 discusses the general system design principles and methodology used in this thesis. Finally, the thesis contributions and organization are summarized in Section 1.4.

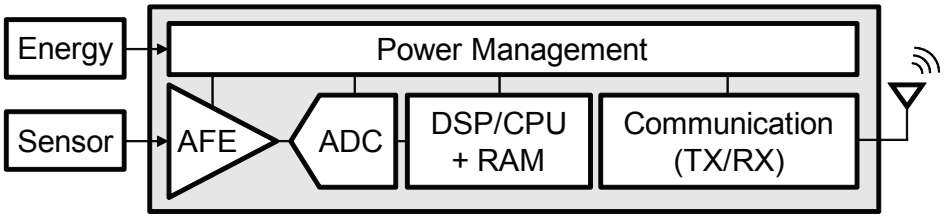

Figure 1-1: Block diagram of a typical wearable sensor node.

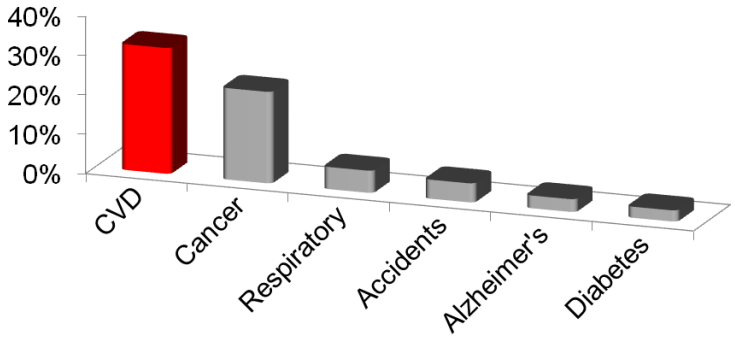

Figure 1-2: Leading causes of death in the United States for 2009 from the National Vital Statistics Reports, Centers for Disease Control (CDC) [15]. Cardiovascular disease (CVD) accounted for 32.3% of all deaths.

1.1 Wearable ECG Monitoring

Figure 1-1 shows a block diagram of a typical wearable sensor node. It consists of an analog front-end (AFE) and an analog-to-digital converter (ADC) to interface the analog signal from the sensor to the digital signal processor (DSP) on the sensor node. The DSP can further condition the data and store it in local memory (RAM), or communicate with a nearby basestation (such as a smart phone) with a short-range communication link. Furthermore, a power management subsystem interfaces with the energy source such as a battery, super-capacitor, or energy-harvester, and delivers the appropriate system voltages and load currents to the circuit blocks in an efficient manner.

The AFE can interface with sensors for various vital signs such as temperature, blood pressure, respiratory rate, and blood oxygen content. However, the electrocar-

Table 1.1: Typical power breakdown of a modern wearable sensor node. †AFE power is included in the calculation of the effective energy per conversion-step.

Component	Commercial Products	Academic Prototypes
†AFE + ADC	$1 - 5$ pJ/conv-step [23–25]	$0.3 - 3$ pJ/conv-step [26–28]
DSP	$10 - 50$ pJ/bit [29, 30]	$0.3 - 2$ pJ/bit [6, 31]
Wireless Radio	$20 - 90$ nJ/bit [32, 33]	$0.2 - 3$ nJ/bit [34–37]

diogram (ECG) is by and large the most important vital sign because cardiovascular disease (CVD) is the leading cause of death in the United States as shown in Figure 1-2. According to the American Heart Association, in 2009, CVD accounted for 32.3% of all deaths in the United States, leading to an estimated $312.6 billion in health expenditures [16].

Studies have shown that many abnormal symptoms can be detected in an ECG prior to a heart attack or sudden cardiac death [17]. For patients with CVD, a common practice is to wear a Holter monitor to continuously log their ECG for 24 to 48 hours. However, since symptoms of CVD such as cardiac arrhythmias are often very intermittent, a 24 to 48 hour window is often inadequate [18]. Furthermore, Holter monitors are heavy, bulky, and not conducive to long-term monitoring. Therefore, in recent years, there have been significant research efforts in developing wearable sensor nodes for ECG monitoring [7, 10, 19–22].

Table 1.1 shows the power breakdown of a typical sensor node with numbers taken from both commercial products and academic prototypes. As Table 1.1 suggests, recent advances in AFE/ADC and DSP design have pushed their power consumption (10's of pJ/bit) orders of magnitude below the power required for wireless data transmission (10's of nJ/bit). The most effective way to reduce the radio power consumption is to minimize the number of bits to be transmitted to permit highly duty cycled operation. This has been accomplished with data reduction schemes like feature extraction in the analog domain [10] or digital domain [8, 9] by performing local processing directly on the sensor node. Another recent approach that offers significant data reduction with a different set of tradeoffs between generality and data compression is compressed sensing [11]. Wearable sensor nodes also allow for the

possibility of using electronic textiles as a wired communication medium to further reduce the cost of communication to the pJ/bit range [38]. As a result of these recent developments, the power consumption of the AFE and DSP are now once again part of the bottleneck. Therefore, the first focus of this thesis is the design of a front-end for ECG monitoring with *micro-Watt* power consumption. Furthermore, in order to demonstrate compatibility with state-of-the-art ultra-low-voltage DSPs operating below 1V [6, 39], another requirement is to operate the front-end from the same digital supply (0.6V) in order to simplify the power management subsystem.

1.2 Cochlear Implants

As of 2010, over 30 million people in the United States suffer from hearing loss [40]. For those who suffer from conductive hearing loss in which there is damage to the ossicular chain (bones of the middle ear), treatment options include medical or surgical treatment, or various types of hearing aids and prosthetics. A more common form of hearing loss is sensorineural hearing loss in which there is damage to the cochlea (inner ear). In this case, damage to the hair cells in the cochlea degrades the transduction of acoustic information to electrical impulses in the auditory nerve. For mild cases in which there is still a sufficient amount of functional hair cells remaining, a hearing aid can help. However, for those who are profoundly deaf (unable to detect sounds less than 90dB SPL[1]), an implantable cochlear prosthesis remains the only option to stimulate viable auditory nerves.

Cochlear implants (CIs) use electronics to directly stimulate the auditory nerve fibers, thus bypassing the damaged hair cells in the cochlea. Current state-of-the-art CIs consist of an external and internal unit as shown in Figure 1-3. The external unit comprises a microphone to pick up sound, a sound processor to digitize, analyze and compress the sound into coded signals, and a transmitter to send data wirelessly to the internal unit via a coil. The implanted unit comprises a receiver and stimulator unit

[1]Sound pressure level (SPL) in units of dB SPL is a logarithmic measure of sound pressure with respect to a reference of $P_{ref} = 20\mu\text{Pa}_{\text{rms}}$, and is calculated by $P[\text{dB SPL}] = 20\log\left(\frac{P[\text{Parms}]}{P_{ref}}\right)$.

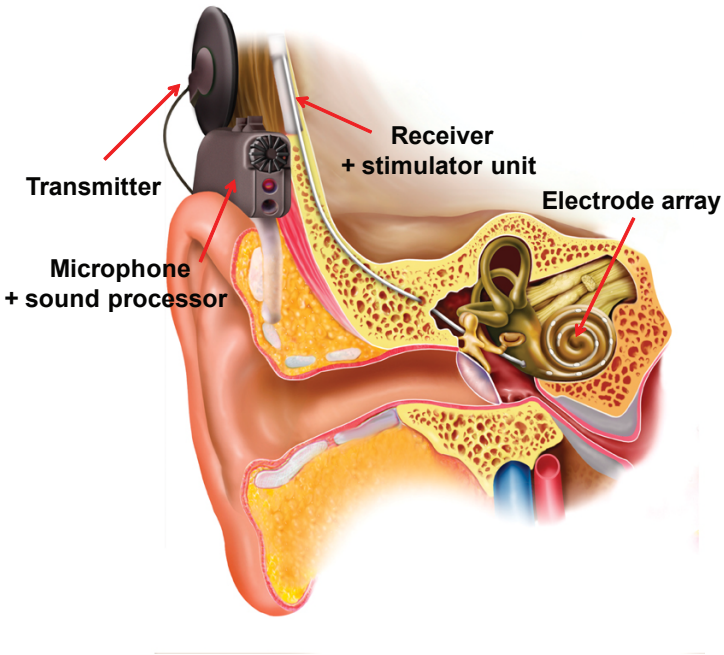

Figure 1-3: Block diagram of a conventional cochlear implant.

embedded in the skull, and an electrode array which contacts the cochlea. Electrical current stimulus is modulated by the received codes and delivered to the electrode array, triggering action potentials in the auditory nerve which are interpreted by the brain as sound.

Although today's CIs are quite successful in restoring hearing in the profoundly deaf, the external components present a number of concerns. Practically, the external unit cannot be worn in the shower or when participating in water sports. Socially, there is a certain degree of social stigma attached to wearing such a device [40, 41]. Finally, external microphone-based CI users are unable to use the sound localization cues provided by the scattering and filtering properties of the natural outer ear (known as pinna cues) [42]. These reasons motivate the development of a fully-implantable cochlear implant (FICI). Among the many challenges of developing a FICI, there are several major obstacles that stand out:

1. **Wireless power delivery and storage**: In today's CIs, the absence of an implanted battery means that power must be transferred *continuously* from an external source, requiring the user to wear the external component at all times.

To address this issue, a FICI must be able to operate untethered (i.e., without a coil that continuously transfers power). One possible approach is to rapidly and wirelessly charge up an implanted ultra-capacitor or battery [43] from which the implanted unit can operate for a full day to enable a one-charge-a-day usage model. This would allow the FICI user to be essentially autonomous, aside from re-charging while asleep.

2. **Implantable acoustic sensor**: A fully-implantable solution requires an implantable acoustic sensor that can replace the conventional external microphone. Previous work has looked at MEMS sensors [40, 44], and piezoelectric sensors [45, 46], in which the principle is to detect the mechanical vibration of the middle ear and process the vibration as sound. The sensor must be small and light weight as to not appreciably alter the vibration of the ossicles, and it must also possess the large dynamic range of speech.

3. **Ultra-low-power sound processing and stimulation**: Using the rapid wireless charging usage model proposed above, a back-of-the-envelope calculation suggests that the total power consumption of the FICI must be limited to 1 mW (assuming 12 hours of operation/day from a 5g ultra-capacitor with an energy density of 5 W·hr/kg and 50% power conversion efficiency). The majority of the 1 mW power budget will be taken up by the process of electrically stimulating the auditory nerve. Nerve fiber stimulation is often the most power-hungry operation because the threshold for action potential initiation is determined by biology, and significant amounts of power is consumed in the electrode impedance and electrode drivers [47]. Assuming typical stimulation power of 150μW to 750μW [48,49], this leaves approximately 250μW for the implantable sensor interface and sound processing. Therefore, a FICI requires very energy-efficient sound processing and stimulation circuits to maximize lifetime given a stringent energy storage constraint.

The second half of this thesis will address the design and implementation of the core circuit blocks of a fully-implantable (i.e., *invisible*) CI, namely items (2) and (3)

above. A system-on-chip (SoC) comprising the sensor interface, sound processor, and energy-efficient neural stimulator is designed to demonstrate feasibility of the system.

1.3 System Design Principles and Methodology

This thesis focuses on the design of ultra-low-power, energy-efficient circuits and systems for medical applications. Throughout the thesis, this is realized by applying four general low-power system design principles to a top-down system design methodology for medical applications. A complete background on ultra-low-power design fundamentals for bio-electronics can be found in [50]. This thesis focuses on the four principles below:

1. **Extending beyond IC design**: Medical devices interact with the human body, and thus it is important for engineers to extend the scope of the design beyond integrated circuits, and examine and understand the interface to the body. Doing so will 1) influence the design of the system architecture, and 2) provide new opportunities for innovation at multiple layers of the design. For example, in Chapter 2, knowledge of the ECG signal aggressors on the body drives the design of a robust system architecture. In a second example, in Chapters 3 and 4, alternate waveforms for electrical neural stimulation are investigated and a middle ear sensor is developed for implantable acoustic sensing. This work drives the design of the CI system-on-chip in Chapter 5.

2. **Highly-digital, mixed-signal circuits and systems**: Since medical devices interface with the body and real world, power-efficient, low-noise analog design is always necessary. However, smart partitioning of the system between analog and digital computation, and the use of highly-digital mixed-signal circuits and systems can improve energy-efficiency, reduce die size, and increase system robustness. For example, in Chapter 2, digital techniques like oversampling and $\Delta\Sigma$-modulation are applied to improve the performance of an analog system, and a mixed-signal feedback loop is used to increase robustness against

power-line interference.

3. **Voltage scaling**: Medical applications are generally low-bandwidth, and are therefore quite suitable for low voltage operation. Scaling the supply voltage of digital circuits down to moderate inversion (i.e., slightly above the threshold) provides a good balance between digital energy efficiency and performance bandwidth [51]. Analog power consumption, on the other hand, is a complex function of factors such as dynamic range, gain-bandwidth, topology, and supply voltage. However, highly-digital mixed-signal circuits (e.g., successive approximation ADCs) can benefit from voltage scaling if the benefits of digital voltage scaling outweigh any increase in analog power. In this thesis, there are numerous examples of digital and mixed-signal circuits operating at 0.6V.

4. **Integration, customization, and optimization**: Maximizing the amount of integration onto a system-on-chip can simplify the overarching system, reduce the amount of I/O between multiple components, and provide the opportunity to optimize the interface between circuits and subsystems. Customization of the design is another effective strategy to reduce power by avoiding general purpose functions which may be inefficient. Finally, the application of fine-grained circuit-level optimization to the entire system can result in significant overall power savings.

These principles will be applied to a design methodology where we first examine the allocation of power in the system, and look for solutions at all layers of the design to reduce the power of the most power-hungry processes. For example, Chapter 3 of this thesis discusses alternate waveforms for neural stimulation in neural prostheses. This aims to reduce the stimulation power which can often dominate the entire power budget in neural prostheses. After addressing the most power-hungry processes (i.e., the "lowest hanging fruit"), we will then smartly design the system architecture and optimize the circuit implementation to further reduce the system power. This will become evident in Chapters 2 and 5, which present two examples of custom ultra-low-power ICs for a wearable and implantable application respectively.

1.4 Thesis Contributions and Organization

This thesis investigates the design of ultra-low-power mixed-signal circuits and systems for two separate medical applications. The first is an ECG front-end system for *wearable* ECG monitoring which is covered in Chapter 2. The second is an invisible, fully-*implantable* cochlear implant that is covered in Chapters 3 to 5. While these two medical applications are seemingly distinct, the common goal of this work is to minimize the power of the overall system in order to meet the stringent energy (or equivalently, volume) limitations that are typical of personal medical devices. This is accomplished by holistically considering the requirements of the system, and optimizing the design at the sensor, architecture, and circuit levels together. The main contributions of this thesis are in the following four areas:

1. **Ultra-low-voltage mixed-signal ECG front-end**: Chapter 2 presents a mixed-signal front-end that leverages a highly-digital architecture and analog circuits optimized for low-voltage in order to operate at 0.6V without sacrificing robustness and dynamic range. Aggressive voltage scaling improves power-efficiency and demonstrates compatibility with low-voltage DSPs for future system-on-chip development. General principles for low-voltage analog design are discussed, and many of the circuit techniques developed in this work are applied to the cochlear implant system. Overall, the front-end consumes just 2.9μW, which is two orders of magnitude lower than commercially available parts. At 0.6V, it also achieves the lowest voltage operation of any ECG front-end to date.

2. **Energy-efficient neural stimulation waveforms**: The power consumption of neural prostheses like cochlear implants is typically dominated by the power required for neural stimulation. Chapter 3 addresses this issue by investigating energy-efficient non-rectangular electrical neural stimulation waveforms to reduce stimulation power. Biphasic waveforms with cochlear-specific parameters are studied with a computational model of a single nerve fiber. The simulation results are first validated with *in-vivo* measurements of ECAP in cats, and then

with threshold and loudness perception in two human cochlear implant users. Preliminary human subject measurements show that energy savings of 15-35% with alternative non-rectangular waveforms are possible within the comfortable hearing range.

3. **Implantable acoustic sensing**: The external microphone of a conventional cochlear implant must be eliminated in a completely invisible system. In Chapter 4, a method for implantable acoustic sensing using a piezoelectric sensor mounted on the middle ear is presented. Using a discrete prototype of the sensor front-end, characterization results from a sensor mounted on two human cadaveric temporal bones indicate that the sensor is capable of sensing sound from 300Hz to 10kHz over a 50dB dynamic range from 40 to 90dB SPL which is adequate for speech.

4. **System-on-chip for a fully-implantable cochlear implant**: In order to incorporate the contributions of Chapters 3 and 4, and to demonstrate the feasibility of a fully-implantable system, Chapter 5 presents a complete system-on-chip that comprises a piezoelectric sensor front-end, an arbitrary waveform stimulator, and a low-voltage reconfigurable digital sound processor to complete the signal path from the sensor to the stimulator. The proof-of-concept system-on-chip is interfaced to a sensor mounted on a human cadaveric temporal bone, and measurement results demonstrate the effectiveness of alternate stimulation waveforms and reconfigurable sound processing on reducing the power of the overall system.

Finally, Chapter 6 summarizes the results of this thesis, provides concluding remarks, and suggests potential directions for future work.

Chapter 2

An Ultra-Low-Voltage Mixed-Signal Front-End for ECG Monitoring

The American Heart Association's 2020 impact goal is stated as "by 2020, to improve the cardiovascular health of all Americans by 20%, while reducing deaths from cardiovascular diseases and stroke by 20%" [16]. This goal is a daunting yet worthy one because cardiovascular disease (CVD) still accounts for over 30% of all deaths in the United States [16]. One focus of the AHA going forward is the emphasis on CVD prevention and promotion of positive cardiovascular health, rather than on treating previously established conditions. However, for patients who have (or are at risk for) CVD, it is critical to monitor their cardiac health over long periods of time in order to detect markers before serious conditions are established [52]. The electrocardiogram (ECG) is one of the most prescribed tools for monitoring overall cardiac health and diagnosing CVD. Holter monitors have conventionally been used for ambulatory ECG monitoring, but they are ultimately limited by their bulky size and limited lifetime of 1 to 2 days [17].

Recent advances in electronics have spurred the development of a wealth of long-term cardiac monitors, mostly in academia [7, 20, 52], but also in the industry [21, 53, 54]. However, there is opportunity to further extend the lifetime and miniaturize the

size of wearable ECG monitors.

This chapter focuses on the design and implementation of an ultra-low-power front-end system for wearable ECG monitoring. One of the goals of this work is to develop an "analog" front-end system capable of operating from an ultra-low "digital" supply voltage of 0.6V without sacrificing robustness or dynamic range. This would ensure compatibility with state-of-the-art DSPs for future system-on-chip development which would provide power, size, and cost benefits. A second complementary goal of this work is to develop the lowest power front-end to aid in maximizing the lifetime of the overall monitor. The main contribution of this work is the fine-tuning of a highly-digital system architecture together with ultra-low-voltage analog circuit optimizations to create a fully-integrated ECG front-end system that meets the required specifications for ambulatory monitoring, and achieves the lowest-in-class power consumption to date. General principles for low-voltage analog design applicable to this work are also provided. A complete background on general low-voltage and low-power design principles can be found in [50, 55].

This chapter is organized as follows: Section 2.1 provides background on low-voltage front-end systems for ECG monitoring. Section 2.2 presents system specifications, approaches for low-voltage systems, and the architecture of the ECG front-end in this work. Section 2.3 describes the design of each of the circuit blocks in detail. Sections 2.4 and 2.5 present measurement results from the fabricated prototype, and Section 2.6 provides a summary of the chapter.

2.1 Background

Large and bulky ambulatory ECG monitors with long cables typically have poor patience acceptance and limited wearability. One way to overcome this issue is to develop a thin and flexible wearable monitor with the electrodes directly on a flexible PCB where the electronics sit [52]. Extending this idea, another concept is the "wearable ECG band-aid" depicted in Figure 2-1, where the goal is to develop an ultra-thin (a few mm in thickness) device that can be easily adhered to the body. Its

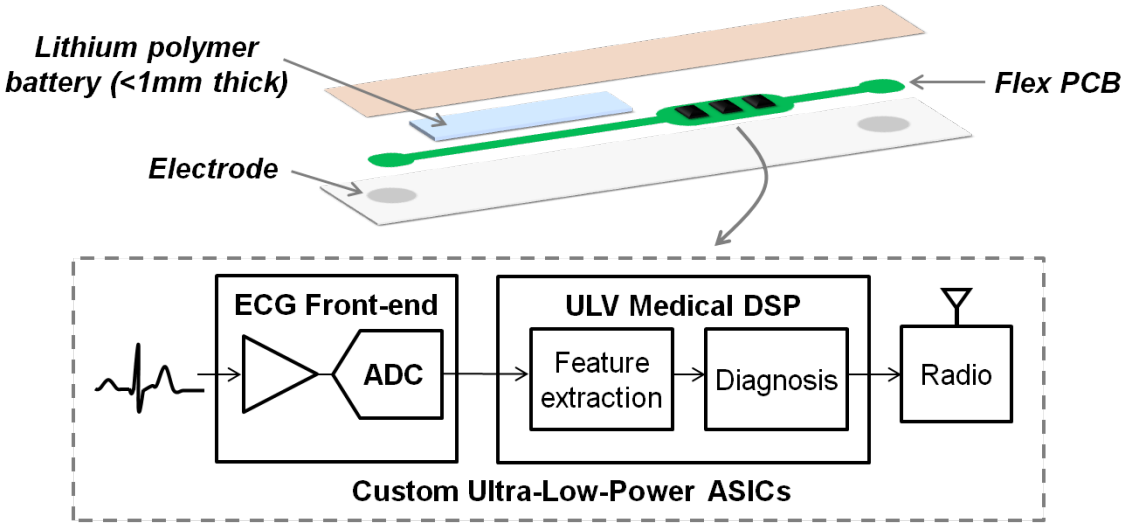

Figure 2-1: A conceptual wearable ECG band-aid with two electrodes for a single-lead measurement. The electronics (block diagram shown) sit on a flexible PCB and are powered from an ultra-thin lithium polymer battery.

low profile and small size would improve patient comfort, increasing the likelihood of patience acceptance.

In order to achieve a thin and small form factor, this conceptual[1] band-aid consists of two electrodes for a single-lead ECG measurement, with the electronics sitting on a flexible PCB, all powered by an ultra-thin lithium polymer battery [56]. Due to the limited storage capacity of the battery, ultra-low-power ASICs are required. The electronics in the band-aid comprise the ECG front-end to acquire the signal, and a custom ultra-low-voltage medical DSP (designed by Joyce Kwong) [6] optimized for sub-1V operation to perform feature extraction and classification in order to compress the amount of data transmitted by the radio. The focus of the work in this chapter is the design of the front-end that operates from 0.6V to improve energy-efficiency through voltage scaling, and also to demonstrate compatibility with state-of-the art DSPs to enable future system-in-package or system-on-chip development for the band-aid system.

Supply voltage scaling is an effective way to achieve quadratic power reduction in digital circuits [57]. For analog circuits however, the situation is more complicated

[1]The band-aid is a vision for a wearable ECG monitor and it is not actually realized in this work.

because power consumption is a complex function of factors such as dynamic range, gain, bandwidth, circuit topology, and supply voltage [50]. For example, in order to maintain a desired dynamic range under a scaled supply voltage, the noise in the circuit must be scaled proportionately which requires an increase in current consumption for purely noise-limited designs [51]. However, if the supply range is used more efficiently at low voltages to maximize signal swing (e.g., minimizing the required headroom), then it is possible for analog power to scale with voltage. Furthermore, in highly-digital mixed-signal systems such as the one in this work, the benefits of voltage scaling on the digital circuits may outweigh any increase in analog power.

There are existing commercial off-the-shelf components for ECG recording like the ADS1298 (Texas Instruments) [25] or the ADAS1000 (Analog Devices) [58] which are very high performance front-end systems that are targeted for diagnostic-quality clinical ECG applications. As a result, they consume mW's of power from a 3V supply and are not well suited for ambulatory monitoring applications. The current state-of-the-art front-ends found in academic literature are able to achieve power consumption that is about 2 orders of magnitude lower than the commercial parts because they are customized for ECG monitoring applications which have less stringent specifications than clinical ECG applications [10, 28, 59]. However, they typically require high supply voltages to perform signal conditioning and accommodate aggressors like electrode offset and 50/60Hz power-line interference. Therefore, this work focuses on the design of a mixed-signal front-end for ECG monitoring that leverages a highly-digital architecture in order to operate from a 0.6V supply which improves power-efficiency through voltage scaling, and facilitates integration with low-voltage DSPs. Particular emphasis is placed on the design of ultra-low-voltage front-end analog circuits aided by configurable and energy-efficient digital processing. A highly-integrated solution is developed to demonstrate feasibility of a 0.6V system.

The next section briefly highlights off-chip and on-chip signal aggressors that can corrupt ECG signals, making voltage scaling challenging for front-end systems.

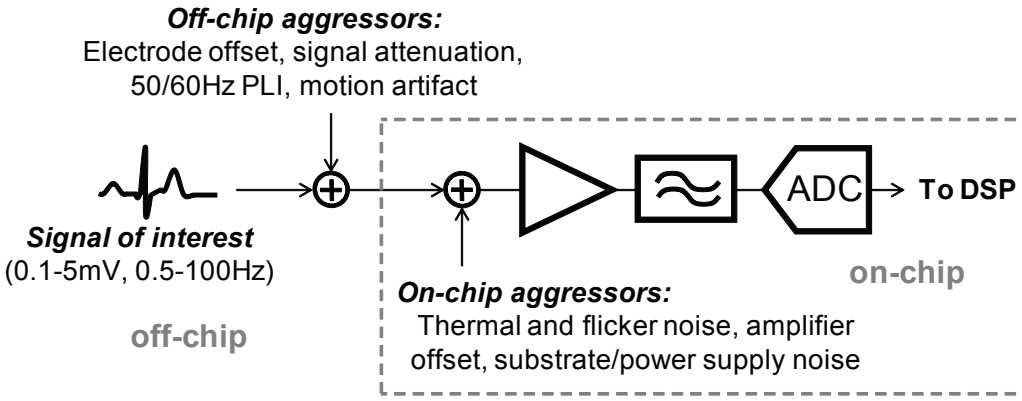

Figure 2-2: Block diagram of a typical sensor signal chain including off-chip and on-chip aggressors which can corrupt the signal of interest.

2.1.1 Sources of Signal Aggressors

The design of the front-end for any system is crucial because it typically determines the noise and dynamic range in the system, and it must faithfully extract the signal of interest often in a noisy environment. Since ECG signals are on the order of 0.1mV to 5mV in amplitude and reside in the 0.5 to 100Hz frequency band, they are susceptible to a number of off-chip and on-chip aggressors which can often be larger than the signal itself as shown in Figure 2-2. Off-chip aggressors encompass all interference from the surrounding environment external to the chip, as well as from changes in the body-sensor interface. On-chip aggressors refer to those that originate from the semiconductor devices within the chip itself.

The most common forms of off-chip aggressors include DC electrode offset which arises from mismatch in the electrode half-cell potential [60], 50/60Hz power-line interference (PLI) coupled from nearby power lines [61], signal attenuation due to non-zero electrode impedances, and motion artifact from random movement which disturbs the skin-electrode interface. The most common types of on-chip aggressors include noise from semiconductor devices such as thermal or flicker ($1/f$) noise, amplifier offset due to threshold voltage mismatch (from random dopant fluctuation), and substrate and power supply noise. Further details on both categories of aggressors are provided in Appendix A.1. It is crucial that the front-end must be able to deal with the numerous off-chip and on-chip aggressors to reliably acquire the signal

of interest. This requires a combination of low-noise design techniques at the circuit level, and robust design choices at the architectural level.

The next section presents the system specifications and architecture. In order to gain more context for the architecture overview, the reader may first want to review Appendix A.1 which discusses the challenges of ECG signal acquisition, with particular emphasis on 50/60Hz power-line interference which is discussed in Appendix A.2.

2.2 Architecture Overview

In this section, we first summarize the required system specifications for ECG recording systems. Then, three different system approaches based on the requirements are considered, and an overview of the system architecture of the front-end is presented.

2.2.1 System Specifications

In this work, the system specifications for ECG recording systems from two separate international standards are considered. The first standard is the IEC 60601-2-47 specification from the International Electrotechnical Commission (IEC) which regulates requirements for ambulatory ECG systems [62]. The second is the ANSI/AAMI EC13 specification from the American National Standards Institute (ANSI) for cardiac monitors [63]. Table 2.1 lists the specifications from both standards.

Table 2.1: Minimum performance specifications for ECG recording systems from international standards committees.

Standard	IEC 60601-2-47 [62]	ANSI/AAMI EC13 [63]
Maximum input signal	$6mV_{p\text{-}p}$	$\pm 5mV$ ($10mV_{p\text{-}p}$)
Input-referred noise	$50\mu V_{p\text{-}p}$ over 10 sec	$30\mu V_{p\text{-}p}$ over 10 sec
Differential input impedance	$10M\Omega$	$2.5M\Omega$
CMRR	60dB @ 50/60Hz 45dB @ 100/120Hz	89dB @ 50/60Hz
Bandwidth	$0.67 - 40Hz$	$0.67 - 40Hz$
Electrode offset tolerance	$\pm 300mV$	$\pm 300mV$

Table 2.2: Summary of target performance specifications for the ECG front-end in this work.

Maximum input signal	$\geq$8mV$_{\text{p-p}}$
Input-referred noise	$\leq 4\mu$V$_{\text{rms}}$
Dynamic range	$\geq$55dB
Common-mode input impedance (Z_C)	$\geq$50MΩ @ 50/60Hz
Differential-mode input impedance (Z_D)	$\geq$20MΩ @ 1Hz
CMRR	$\geq$60dB @ 50/60Hz
Gain	40 to 70dB
Bandwidth	0.5Hz to 150Hz
PLI tolerance	$\geq$10mV$_{\text{p-p}}$
Electrode offset tolerance	$\geq \pm$300mV
Supply voltage	0.6V
Power consumption (2 week lifetime)	$\leq 10\mu$W

The metrics from Table 2.1 are discussed in detail in Appendix A.3, where numerical examples are provided to either justify these specifications, or to suggest alternative requirements. Based on the IEC and ANSI/AAMI standards and the analysis in Appendix A.3, a summary of the performance specifications for the ECG front-end in this work is provided in Table 2.2. These specifications will guide the design of the system architecture.

2.2.2 Approaches for Low-Voltage Systems

It is well known that off-chip aggressors such as electrode offset and PLI can interfere with ECG recording systems if not properly accounted for. Both electrode offset and PLI can be limiting factors for low-voltage systems. For example, just 10mV of electrode offset and 5mV$_{\text{p-p}}$ of PLI can easily saturate a sub-1V analog front-end with 40dB of gain before digitization by the ADC. Even if the problem of signal saturation is avoided, these signal aggressors can also significantly increase the front-end dynamic range requirements. As a different example of the worst case, assume that the desired ECG signal of 0.1mV sits on top of a 300mV DC electrode offset, with 1mV$_{\text{p-p}}$ of PLI coupled onto it. Furthermore, a signal dynamic range of 55dB is required (see Table 2.2) for ECG monitoring applications. This translates to roughly 124dB (20log(300/0.1)+55) or 20 bits of required dynamic range in the system because

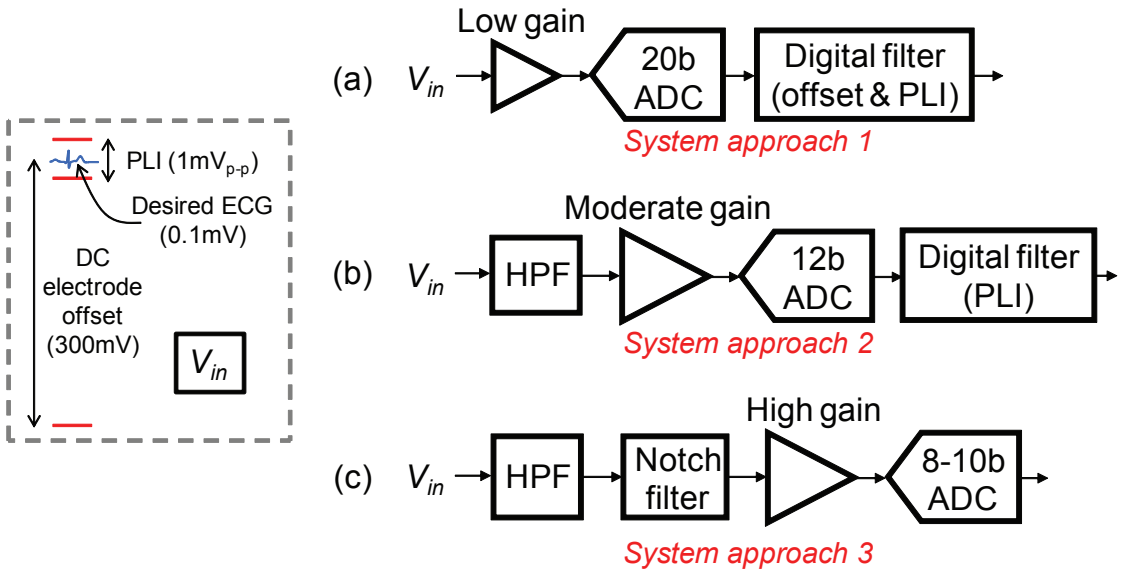

Figure 2-3: System approaches for acquiring the desired ECG in the presence of large aggressors using a (a) high resolution, (b) medium resolution, and (c) low resolution ADC. V_{in} is the sum of the desired ECG signal with the DC electrode offset and PLI.

most of that is wasted on the large electrode offset. This situation is depicted on the left of Figure 2-3, and three system approaches are considered.

The first system approach shown in Figure 2-3(a) could be to use a low-gain amplifier to prevent the electrode offset from saturating the front-end, and a 20-bit ADC to digitize the signal and interferers. Both offset and PLI can then be filtered out in the digital domain. However, the design of an ADC with 20 effective bits of resolution is extremely challenging, and only a few parts from Texas Instruments or Analog Devices actually satisfy the performance requirements. The ADS1222 (Texas Instruments) or the AD7799 (Analog Devices) are both 24-bit ADCs with approximately 20 bits of effective resolution in a 200Hz bandwidth [64, 65] which would satisfy the dynamic range requirements. However, they consume about 1mW from a 3V supply (which is the state-of-the-art for commercially available ADCs in this range) which is about 2 orders of magnitude more power than the budget allows for in this work (see Table 2.2). Therefore, this first system approach proves to be too power hungry, and the design of a 20-bit ADC is extremely challenging and requires

very specialized expertise[2].

A second system approach shown in Figure 2-3(b), could be to place a high-pass filter (HPF) at the front-end of the system to reject the 300mV electrode offset, while allowing the ECG (0.1mV) and PLI ($1mV_{p\text{-}p}$) to pass through. This relaxes the dynamic range requirements after the HPF to approximately 75dB ($20\log(1/0.1)+55$) or 12 bits of resolution. With this approach, the rest of the signal chain would comprise a moderate gain amplifier and a 12-bit ADC which would be much easier to design, and consume much less power. As an estimate, assuming a reasonable ADC figure-of-merit (FOM) of 100fJ/conversion-step [66], a 12-bit ADC with a 500Hz sampling rate would require approximately $P = FOM \cdot 2^{ENOB} \cdot f_S = 0.2\mu W$ of power, which would be within the system budget. However, the amount of supply voltage scaling would still be limited by the expected amount of PLI (which can be larger than the ECG). Therefore, this second system approach could be feasible from a power consumption standpoint, but may still not be suitable for ultra-low-voltage systems as required in this work.

Finally, a third approach shown in Figure 2-3(c) could be to add a 50/60Hz notch filter at the front-end in addition to the HPF to remove PLI right at the front end, which further relaxes the dynamic range requirements of the amplifier and ADC, and allows the supply voltage for those blocks to be scaled more aggressively [59]. In this approach, the forward path blocks would only have to satisfy the dynamic range requirements of the desired signal itself, which allows the use of an ADC with 8 to 10 bits of resolution. This simplifies the block level design, uses the system dynamic range more efficiently, and reduces the power consumption further. For these reasons, this final approach is chosen for the ECG front-end in this work, and the system architecture including the implementation of the HPF and notch filter is discussed in the next section.

[2]The ADS1298 from Texas Instruments is a 24-bit analog front-end for bio-potential recording that uses this first system approach, requiring approximately $750\mu W$/channel.

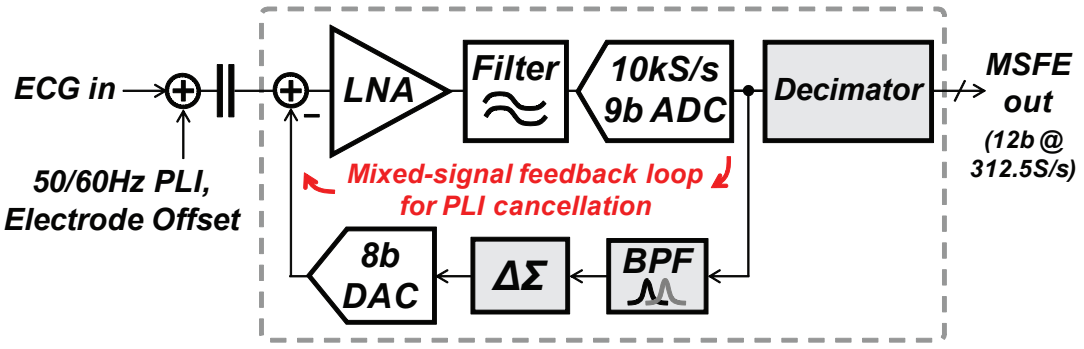

Figure 2-4: System architecture of the mixed-signal front-end that passively rejects electrode offset with *ac*-coupling, and actively rejects PLI with a mixed-signal feedback loop.

2.2.3 Mixed-Signal Front-End System Architecture

A simplified block diagram of the system architecture for the proposed ECG front-end is shown in Figure 2-4, with the analog circuit blocks shown in white, and digital circuit blocks shaded in gray. Because of the tight integration between analog and digital circuits in this front-end system, it is appropriately called a mixed-signal front-end (MSFE) for ECG monitoring. The system is based on the first generation system designed by Jose Bohorquez [59]. This architecture is robust to off-chip interferers by passively rejecting DC electrode offset through *ac*-coupling (i.e., the inputs are capacitively coupled), and actively canceling PLI with a mixed-signal feedback loop annotated in red.

Ignoring the PLI cancellation loop for now, the operation of the forward signal path of the MSFE is as follows: The ECG signal itself is *ac*-coupled and gained up by a low-noise amplifier (LNA), and conditioned by an anti-aliasing filter with programmable gain. The signal is then digitized by a 9-bit ADC which is oversampled at f_S=10kS/s to relax the filtering requirements of the anti-aliasing filter, and also improve its effective dynamic range. For example, f_S=10kS/s provides an oversampling ratio (OSR) of 33.3 (assuming an ECG bandwidth of 150Hz), improving the dynamic range of the ADC by 15dB (10log(OSR)). Assuming that the 9-bit ADC has an ENOB of 8.5 bits, the effective dynamic range of the ADC becomes 68dB (DR =

6.02×ENOB + 1.76 + 10log(OSR) [dB]). This improvement in dynamic range is only realized by using a digital decimation filter to remove out-of-band noise and reduce the data rate back down to the Nyquist rate. The 2-stage decimation filter in this system provides the necessary filtering and downsampling by 32×, producing 12-bit output data at 312.5Hz.

Next, the operation of the mixed-signal PLI cancellation loop is described. Any 50/60Hz PLI content in the signal at the ADC output is captured in the feedback path by a programmable digital band-pass filter (BPF) which can be tuned to 50Hz or 60Hz. The 50/60Hz content is then negatively fed back to the input of the system through a $\Delta\Sigma$-modulated 8-bit digital-to-analog converter (DAC) to servo it out. The BPF in the feedback path creates a sharp notch at 50/60Hz in the closed-loop response of the system, implementing the notch filter required in the third system approach of Figure 2-3. By canceling out the 50/60Hz PLI right at the front-end, the dynamic range requirements of the forward path blocks are relaxed. This is one of the key enablers of ultra-low-voltage and low-power operation [59].

A detailed block diagram of the ultra-low-voltage MSFE implemented in this work is shown in Figure 2-5. It is suitable for either a two-electrode system, or three-electrode system with a patient ground reference electrode as shown in Figure A-2 of Appendix A. In addition to the blocks described above, the MSFE also includes peripheral circuits like current and voltage references, as well as an oscillator and clock generation circuits on-chip, providing a fully-integrated solution at 0.6V.

Overall, this mixed-signal system shifts the burden of signal processing to the digital domain which is suitable for low-voltage systems. For example, oversampling is leveraged and digital decimation filtering is used to improve the effective resolution of the system. Digital $\Delta\Sigma$-modulation is also used in the feedback path to improve the dynamic range of the DAC by shaping its quantization noise out-of-band (made possible by first oversampling the system). These digital dynamic range enhancements permit the use of relatively low resolution analog components (9-bit ADC, 8-bit DAC) without limiting the system dynamic range, resulting in a very energy-efficient system.

Compared to the first generation system in [59], the new contributions and im-

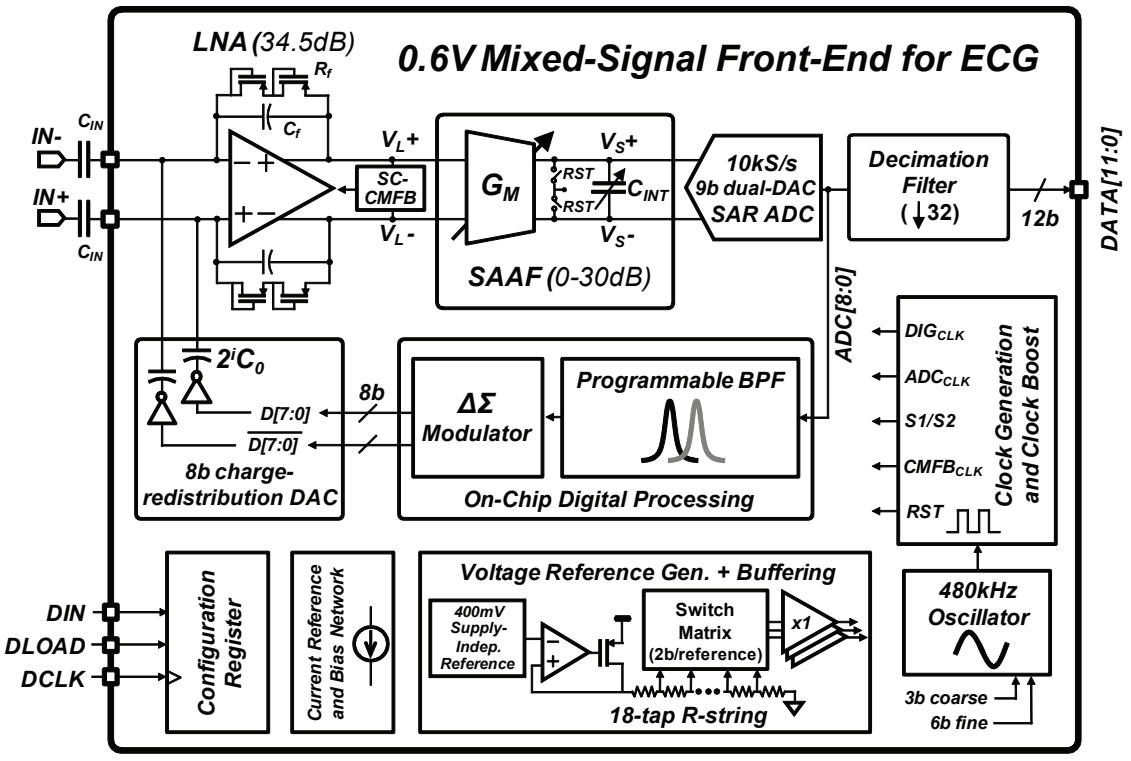

Figure 2-5: Detailed system block diagram of the ultra-low-voltage fully-integrated mixed-signal front-end.

provements of this work are listed below:

- **Aggressive supply voltage scaling**: The supply voltage is reduced from 1.5V to 0.6V by optimizing the design of the analog circuits for ultra-low voltage. General principles for low-voltage analog design are discussed in Section 2.3.1.
- **Digital ΔΣ-modulation**: Digital ΔΣ-modulation is added to the feedback DAC to simultaneously improve the PLI cancellation range, while maintaining low input-referred noise.
- **Ultra-low-power SAR ADC**: A buffer-less ultra-low-power successive approximation register (SAR) ADC is designed to achieve the oversampling in the system without consuming excessive power.
- **Integration**: The first generation prototype required an off-chip ADC and FPGA to implement the digital blocks and feedback system [59]. This work fully integrates all components on a single die including all required current/voltage

references and clock generation. A decimation filter is also integrated to filter out high frequency noise, reduce the data rate, and improve the effective resolution of the system.

- **Performance**: By fine-tuning the system architecture and optimizing individual circuit blocks, the overall system in this work achieves better performance, a higher level of integration, and lower power consumption.

2.3 Description of Circuit Blocks

This section presents details on the design of each of the core circuit blocks in the MSFE shown in Figure 2-5. All analog circuits were designed within the Cadence environment using Spectre or Hspice. All digital blocks were coded in Verilog, synthesized using Synopsys Design Compiler, and place-and-routed using Astro. The digital flow used a nominal supply voltage of 1.8V, and significant margin based on a simulated scaling factor was added to ensure operation at 0.6V. Top-level mixed-signal verification was completed with Synopsys NanoSim as well as Magma FineSim.

2.3.1 Principles for Low-Voltage Analog Circuit Design

Before delving into the details of the circuit description, we first highlight some general design principles that are followed in this work to enable the analog circuits to operate from an ultra-low supply voltage of 0.6V. A more complete background on general low-voltage design techniques can be found in [50, 55].

1. **Folding and sub-threshold biasing to enable cascoding**: Many conventional analog circuit topologies rely on cascoding (i.e., stacking) to achieve high gain and output resistance. This becomes difficult at low voltage because the reduced voltage headroom limits the number of stacked devices that can remain in saturation. Here, two strategies are used to enable cascoding at low voltage: 1) using folded circuit topologies, and 2) biasing transistors in sub-threshold when possible. First, by using folded circuits (e.g., folded-cascode op-amps),

cascoding can still be used while limiting the transistor stack to at most 4 devices. Secondly, biasing transistors in sub-threshold reduces their saturation voltage (V_{dsat}) to just 4-6× the thermal voltage, or roughly 100-150mV [55, 67]. This allows for all 4 stacked devices to remain in saturation at a supply voltage of 0.6V. Another benefit of sub-threshold biasing is that g_m/I_D is maximized which can help minimize circuit noise if applied to the appropriate devices.

2. **Cascaded topologies**: When stacking is not possible, another approach to achieve high gain at low voltage is to cascade multiple stages. Cascading is also necessary to ensure that large signal swings only occur where it is possible (e.g., in a common-source stage with 2 stacked devices), and avoid large signal swings at nodes with many stacked devices. For example, a popular strategy is to cascade a folded-cascode first stage with a common-source second stage. The common-source stage can handle large signal swings (up to 300mV), while the gain of the common-source stage limits the signal swing at the folded-cascode stage to at most a few tens of mV as required.
3. **Fully-differential operation**: Fully-differential operation is used to increase the signal swing at low voltage, and also to improve robustness against common-mode interference (e.g., from the substrate or power supply). However, fully-differential circuits typically require common-mode feedback which must be both stable and linear. For cascaded circuits (e.g., 2- or 3-stage op-amps), the use of local common-mode feedback at each stage helps to avoid placing multiple poles in a larger loop that is difficult to compensate, especially at low voltages.
4. **Avoid signal path switches**: The use of MOS switches should be avoided in the signal path or at any node that is close to the mid-rail because of their degraded *on*-resistance at low voltage. When possible, MOS switches are placed at the rails (i.e., either at ground or the supply voltage) in order to maximize their V_{GS} to improve their conductance. If placing switches in the signal path is unavoidable (e.g., sampling switches in an ADC), bootstrap circuits can be used to boost the gate voltage as long as it does not present a reliability concern [68].

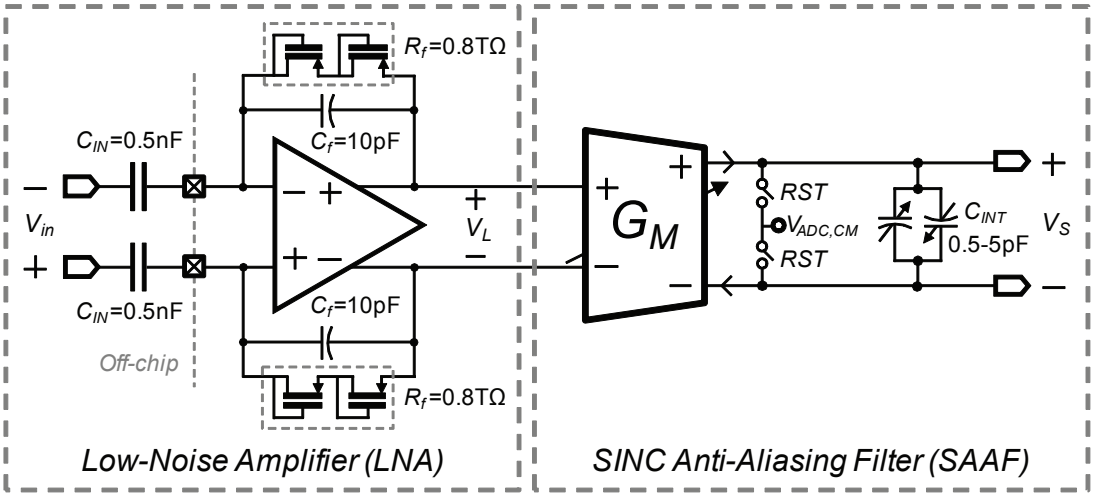

Figure 2-6: Architecture of the LNA with passive feedback, and discrete-time SAAF. The input capacitors (C_{IN}) of the LNA in this work are implemented off-chip, and V_S is the input to the ADC.

2.3.2 Front-End Analog Circuit Overview

The schematic of the low-noise amplifier (LNA) and anti-aliasing filter based on [59] and used in this work is shown in Figure 2-6. The signal chain is fully-differential to maximize signal range at low voltage. The LNA uses passive feedback to accurately set the gain and ensure good linearity at 0.6V. The mid-band gain is determined by the ratio C_{IN}/C_f and is set to 34dB, and *ac*-coupling through C_{IN} achieves greater than ±300mV of electrode offset rejection as required (see Table 2.2). The anti-aliasing filter comprises a transconductor and integration capacitor that is periodically reset at the system sampling frequency of f_S=10kHz. This creates an integrate-and-dump function at f_S, resulting in a *SINC*-shaped frequency response with deep notches placed at f_S and its integer multiples which are precisely in the center of the aliasing bands [59]. The LNA and the *SINC* anti-aliasing filter (SAAF) are described in detail over the next two sections.

2.3.3 Low-Noise Amplifier

The design and sizing of components in the LNA of Figure 2-6 is constrained by a number of specifications such as noise, common-mode rejection, input impedance,

and linearity (Table 2.2). A complete analysis of the LNA architecture in this work is provided in Section 4.3 of [69], and the results are used here to drive the optimization of the LNA in this work.

Transfer Function

In the frequency range of interest where the loop gain of the feedback loop is large, the input-output transfer function from v_{in} to v_L is given by

$$H_{LNA}(s) = \frac{v_L}{v_{in}}(s) = \frac{C_{IN}}{C_f}\frac{s}{s+\omega_f} \tag{2.1}$$

where $\omega_f = \frac{1}{R_f C_f}$ is the high-pass corner frequency, and $G_{LNA} = \frac{C_{IN}}{C_f}$ is the mid-band gain of the LNA. $H_{LNA}(s)$ provides the desired high-pass characteristic to filter out the large DC electrode offset. In this work, the nominal value of G_{LNA} was designed to be 50V/V (34dB), where C_{IN}=500pF and C_f=10pF. The choice of these values will be justified below with noise, input impedance and CMRR considerations. Finally, the high-pass corner was designed to be $f_f = \frac{\omega_f}{2\pi} = 20$mHz, requiring R_f to be 800GΩ, which is achieved with the widely used MOS-bipolar pseudo-resistor [70].

Noise

The two main contributors to the input-referred noise of the LNA is the noise from the op-amp ($V^2_{ni,oa}(f)$), and the noise from the feedback resistor R_f ($V^2_{ni,R_f}(f)$). The total input-referred noise power density of the LNA is given by

$$\begin{aligned} V^2_{ni}(f) &= \overbrace{V^2_{ni,oa}(f)}^{} + V^2_{ni,R_f}(f) \\ &= \overbrace{\left(\frac{C_{IN}+C_f}{C_{IN}}\right)^2 \left[\underbrace{\frac{4kT}{\kappa g_m}\left(\frac{\text{NEF}}{2.02}\right)}_{\text{thermal noise}} + \underbrace{\frac{K_f}{WLC_{ox}f}\alpha_f}_{1/f\text{-noise}}\right]}^{} + \overbrace{\frac{4kT}{R_f}\left(\frac{1}{2\pi f C_{IN}}\right)^2}^{} \end{aligned} \tag{2.2}$$

where the thermal noise and $1/f$-noise contributions from the op-amp are shown with braces, k is the Boltzmann constant, T is the temperature in Kelvin, g_m is

the transconductance of the op-amp input pair, W and L are the op-amp input pair dimensions, and K_f, κ, and C_{ox} are all process constants. The extra factors NEF/2.02 and α_f account for the additional thermal and $1/f$ noise contributed by other transistors in the op-amp beyond the input pair [69]. The $(C_{IN} + C_f)/C_{IN}$ factor in $V^2_{ni,oa}(f)$ represents the transfer function that refers the noise of the op-amp to the input of the LNA.

The second term in Equation 2.2 is the input-referred noise power spectral density of R_f. Even though resistor noise is typically white, the op-amp and feedback network shape the noise to give a $1/f^2$ shape when referred to the input of the LNA. In the first generation system [59], this $1/f^2$ noise from R_f was dominant below 5Hz (see Fig. 13 of [59]) and limited the noise performance of the LNA at low frequencies. In this work, the sizing of components is optimized so that the contribution of $V^2_{ni,R_f}(f)$ becomes negligible, and the total LNA noise is set by the op-amp noise only.

Based on Equation 2.2, assuming the noise terms in the square bracket are determined by the op-amp topology and design, there are two knobs that can be used to minimize $V^2_{ni}(f)$. First, a large C_{IN}/C_f ratio will limit the contribution of the op-amp noise when input-referred (so that $(C_{IN} + C_f)/C_{IN}$ is close to unity). Secondly, a large C_{IN} will limit the $1/f^2$ noise from R_f. However, C_{IN} cannot be made arbitrarily large because of size/volume constraints, and we will also see that C_{IN} is inversely proportional to the differential-mode input impedance Z_D.

A design choice was made to set $G_{LNA} = C_{IN}/C_f$ to 50 so that $(C_{IN} + C_f)/C_{IN}$ is 1.02. Setting G_{LNA} too large would lead to saturation of the LNA with large input signals, and so G_{LNA}=50 is a good tradeoff for a supply voltage of 0.6V. In order to determine the required value of C_{IN} to minimize $V^2_{ni,R_f}(f)$, consider that $\omega_f = \frac{1}{R_f C_f}$ is typically a fixed design parameter. Assuming that both G_{LNA} and ω_f are fixed, R_f is expressed as $R_f = \frac{G_{LNA}}{\omega_f C_{IN}}$. By substituting R_f into $V^2_{ni,R_f}(f)$, the constraint for C_{IN} given a maximum value of $V^2_{ni,R_f}(f)$ becomes

$$C_{IN} \geq \frac{4kT\omega_f}{G_{LNA}} \cdot \frac{1}{(2\pi f)^2} \cdot \frac{1}{V^2_{ni,R_f}(f)|_{\max}}. \tag{2.3}$$

For example, if we want to limit $V_{ni,R_f}^2(f)$ to just 10% of $V_{ni,oa}^2(f)$ at f=20mHz, then $V_{ni,R_f}^2(f)|_{\max}$=9×10^{-12} V^2/Hz ($V_{ni,oa}^2(f)$=9×10^{-11} V^2/Hz at f=20mHz from simulation). Evaluating Equation 2.3 shows that the minimum required value of C_{IN} is 293pF, and so a value of 500pF was chosen for adequate margin. Therefore, given that G_{LNA}=C_{IN}/C_f=50 and C_{IN}=500pF, the noise of the LNA is essentially set by the noise of the op-amp.

CMRR and Input Impedance

As discussed in Appendix A.3, the CMRR and common-mode input impedance Z_C need to be large to limit the amount of differential PLI, while the differential-mode input impedance Z_D needs to be large to avoid signal attenuation. In the bandwidth of interest ($\omega > \omega_f$), the CMRR is determined by the mismatch in C_{IN}, C_f, and any non-zero common-mode gain of the op-amp, $A_{cm}^{'}$. Each source of mismatch can be analyzed separately, and the overall CMRR of the LNA is determined by superposition according to

$$\frac{1}{CMRR} = \frac{1}{CMRR_{\Delta C_{IN}}} + \frac{1}{CMRR_{\Delta C_f}} + \frac{1}{CMRR_{\text{op-amp}}}. \tag{2.4}$$

It can be shown that the overall CMRR of the LNA in the bandwidth of interest ($\omega > \omega_f$) is

$$CMRR = \frac{1 + \frac{C_{IN}}{C_f}}{\frac{\Delta C_{IN}}{C_{IN}} + \frac{\Delta C_f}{C_f} + A_{cm}^{'}}. \tag{2.5}$$

Due to the large size of C_{IN} in this work, it is implemented with an off-chip ceramic capacitor, while C_f=10pF is implemented with metal-insulator-metal (MIM) capacitors on-chip. Therefore, for typical technologies, expected values of $\Delta C_{IN}/C_{IN}$ and $\Delta C_f/C_f$ are approximately 1% and 0.1% respectively, and $A_{cm}^{'}$ can be assumed to be much less than 1% for a fully-differential op-amp. This results in a theoretical CMRR of 20log($\frac{1+50}{0.01+0.001}$) = 73dB, which satisfies the IEC specification of 60dB (but not the ANSI/AAMI specification). Despite this relatively low value, the mixed-signal PLI

notch increases the effective CMRR at the PLI frequency by removing any residual differential PLI.

It can also be shown that the common-mode input impedance of this LNA is

$$Z_C = Z_{IN} + Z_f \tag{2.6}$$

where $Z_{IN} = \frac{1}{j\omega C_{IN}}$ and $Z_f = \frac{R_f}{1+j\omega R_f C_f}$. Given that in the bandwidth of interest ($\omega > \omega_f$) $Z_f \approx \frac{1}{j\omega C_f}$, and $C_{IN} \gg C_f$, then $Z_C \approx \frac{1}{j\omega C_f}$.

Finally, because of the virtual ground, the differential-mode input impedance is simply the series impedance of the input capacitors,

$$Z_D = \frac{2}{j\omega C_{IN}}. \tag{2.7}$$

For the chosen values of C_{IN}=500pF and C_f=10pF, the theoretical value of Z_C at 60Hz and Z_D at 0.5Hz are 265MΩ and 1.27GΩ respectively, which are both well above the required specification in Table 2.2.

LNA Op-Amp Schematic Design

The schematics of the LNA op-amp and its common-mode feedback (CMFB) circuits are shown in Figure 2-7. The design of the op-amp is guided by the general principles for low-voltage analog design as discussed earlier. The op-amp comprises two cascaded stages. The first is a folded-cascode stage to achieve high gain, and PMOS input devices are used to minimize $1/f$-noise. By using a folded structure, cascoding can still be used to achieve high gain while limiting the transistor stack to at most 4 devices. This permits approximately 150mV of V_{DS} across each transistor which is adequate for keeping all devices in saturation at 0.6V. Feedback ensures that the signal swing at the output of this stage (at V_P and V_M) is just a few mV. The output stage is a common-source stage with only 2 stacked devices to allow up to 250mV of swing on each side (500mV differentially).

As a result of the fully-differential signal path, the LNA op-amp requires common-

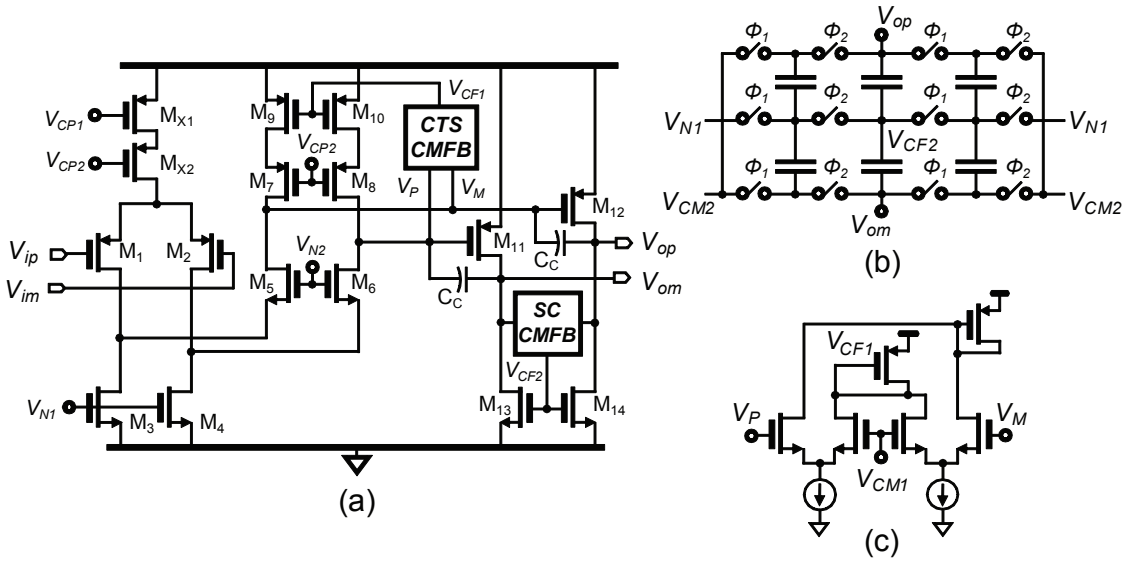

Figure 2-7: (a) Schematic of the fully-differential two-stage op-amp used in the low-noise amplifier. The switched-capacitor CMFB (SC-CMFB) and continuous-time CMFB (CTS-CMFB) circuits are shown in (b) and (c) respectively.

mode feedback that is both stable and linear which can be challenging at low voltage. At the output stage where the signal swing is large, switched-capacitor CMFB (Fig. 2-7(b)) is used because it is inherently linear and does not introduce distortion in the differential signal path. A consequence of using switched-capacitors is that switching noise is injected into the system at the switching frequency of 10kHz. However, as we will see, these noise spikes will be filtered out by the SAAF described later.

At the output of the folded-cascode stage, a more area-efficient continuous-time CMFB (Fig. 2-7(c)) can be used without worry of distortion because the signal swing is small. By separating the CMFB into two local feedback loops, we avoid placing multiple poles in a larger loop that is difficult to compensate at low voltage.

2.3.4 *SINC* Anti-Aliasing Filter

Following the LNA, the signal is then filtered and gained by the *SINC* anti-aliasing filter (SAAF) shown in Figure 2-6. Anti-alias filtering is required in any system prior to sampling at f_S to prevent the unwanted folding of signal aggressors back onto the baseband signal. Fig. 2-8(a) shows the continuous-time spectrum of a

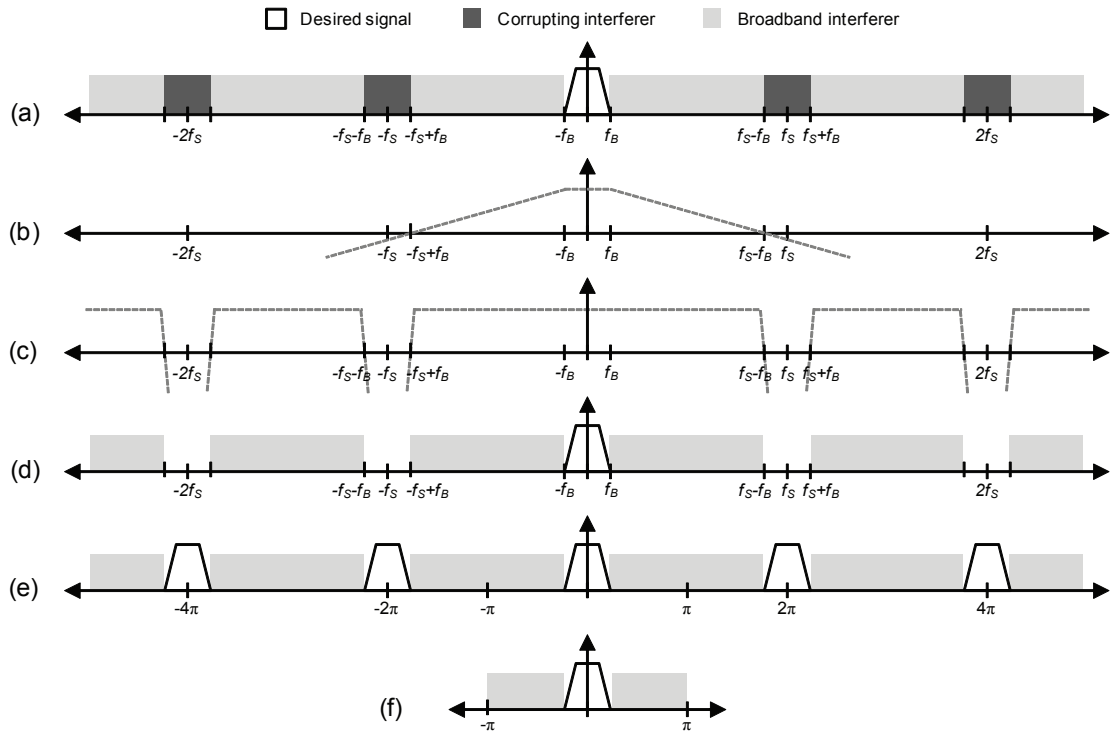

Figure 2-8: (a) Continuous-time spectrum of a signal prior to sampling, with the desired component in white, broadband interferer in light gray, and the corrupting interferer in dark gray at multiples of the sampling frequency. (b) A conventional anti-aliasing filter (AAF) with a low-pass response. (c) An alternative frequency response of an AAF that rejects only the corrupting interferers, and allows aliasing outside the bandwidth of the desired signal. (d) Continuous-time spectrum of the signal in (a) after passing through the AAF shown in (c). (e) Discrete-time spectrum of the signal in (d) after sampling. (f) Digital baseband of the signal in (e).

generic signal prior to sampling, where the desired signal (with bandwidth f_B) is in white, broadband interferers are in light gray, and the corrupting interferers are in dark gray. Fig. 2-8(b) shows the frequency response of a conventional low-pass anti-aliasing filter, where anything beyond $f_S - f_B$ is filtered out to avoid aliasing. However, an alternative frequency response (proposed by Jose Bohorquez for the first generation system in [69]) for anti-aliasing is provided in Fig. 2-8(c), where it rejects only the corrupting aggressors situated at f_S and its integer multiples, but allows aliasing everywhere else outside the desired signal bandwidth. Fig. 2-8(d) shows the spectrum of the original signal after passing through the alternative anti-aliasing filter, and Fig. 2-8(e) is the corresponding discrete-time spectrum after sampling.

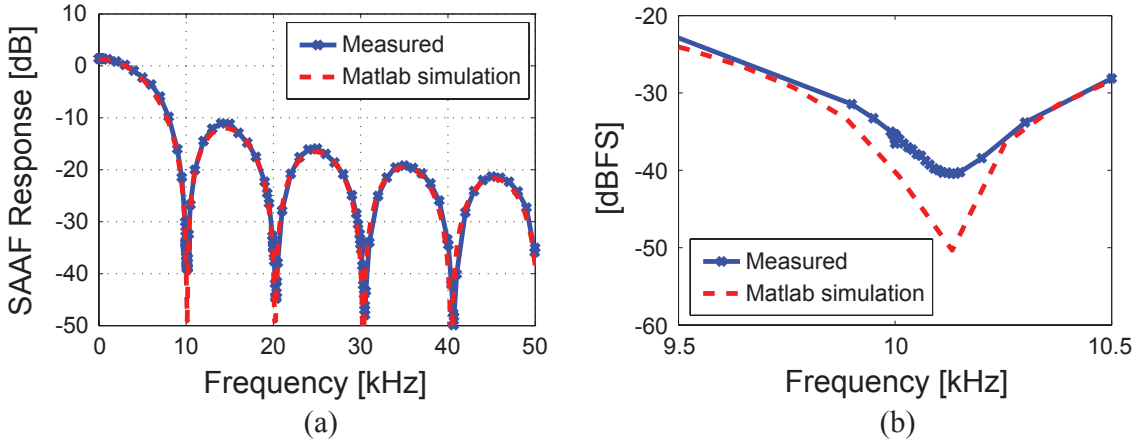

Figure 2-9: (a) Simulated and measured frequency response of the SAAF. (b) Zoom in of the notch in the frequency response about 10kHz.

Finally, Fig. 2-8(f) shows the digital baseband of the signal after sampling. As a result of the notches in the alternative filter at f_S and its integer multiples, there is no aliasing within the signal band.

This alternative approach for anti-aliasing in combination with an oversampled system allows for a very robust and area-efficient circuit implementation based on *charge sampling*. The SAAF shown in Figure 2-6 works by using a transconductor to integrate the output signal current onto an integration capacitor C_{INT} over a period $T_S = 1/f_S$, and then the voltage on the capacitor is sampled at the end of the period. After the charge has been sampled, C_{INT} is reset (i.e., discharged), and the process repeats. This process of charge sampling (i.e., integrate-and-dump), results in a *SINC*-shaped frequency response with deep notches placed at f_S and its integer multiples, which satisfies the frequency response of the alternate anti-aliasing filter approach shown in Figure 2-8(c). The frequency response of the SAAF is shown in Figure 2-9, and a zoom in of the notch at 10kHz shows that better than 30dB of anti-aliasing is achieved over the small 150Hz ECG bandwidth around 10kHz.

Figure 2-10 shows time-domain waveforms illustrating the operation of the SAAF, with the input to the SAAF shown in blue (V_L), output of the SAAF shown in green (V_S), and the ADC samples at f_S=10kHz in red. Fig. 2-10(a) shows the waveforms for a low-frequency signal, where the integration period (T_S=$1/f_S$=100μs) of the SAAF

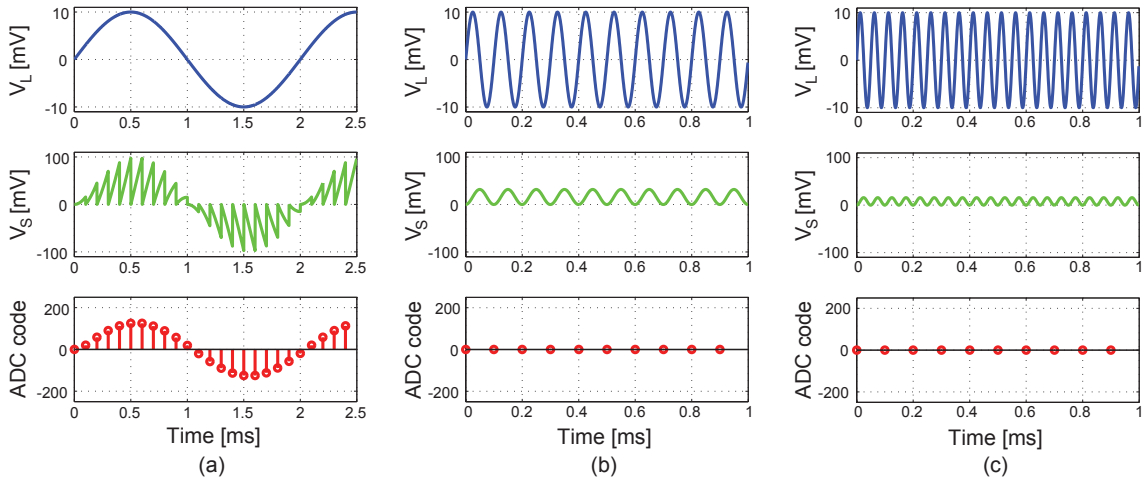

Figure 2-10: Time-domain waveforms V_L, V_S, and ADC code which are the outputs of the LNA, SAAF, and ADC respectively. The SAAF gain is set to 20dB, and the integration period is 100μs corresponding to an ADC sampling rate of 10kS/s. Signal tones of 500Hz, 10kHz, and 20kHz are shown in (a), (b), and (c) respectively.

is small compared to the period of the signal, and the ADC output represents the input signal with a certain amount of gain. Fig. 2-10(b) shows the waveforms when the input signal frequency is 10kHz. Since the integration period T_S exactly matches one period of the signal, the SAAF output V_S integrates to zero each time, resulting in zero ADC output. This is the time-domain illustration of the notch at 10kHz in the SAAF response. A similar thing happens when the input signal is at 20kHz, and the waveforms are shown in Fig. 2-10(c).

In addition to anti-aliasing, the SAAF also provides a low-frequency gain of

$$G_{SAAF} = \frac{G_M}{2f_S C_{INT}} \tag{2.8}$$

which is digitally tunable through G_M and C_{INT}. C_{INT} is the sum of a 4-bit switched-capacitor, as well as the input capacitance of the ADC (discussed in Section 2.3.5). G_M is also set digitally with the transconductor circuit shown in the next section. In this work, C_{INT} is implemented as a differential capacitance between the differential outputs of the transconductor with a maximum value of 5pF. Compared to the implementation in [59] which uses two single-ended 10pF capacitors, it requires 4$\times$ less

capacitance (and area) to achieve the same gain.

This *charge sampling* SAAF approach was proposed by Jose Bohorquez and is analyzed in detail in [69], and the three main benefits are summarized here:

1. **Area-efficiency**: A conventional low-pass AAF would require much more chip area due to larger passives. As a simple example, consider a signal bandwidth of 100Hz, with a 10kHz sampling frequency. Assuming that we require 40dB of attenuation at 9.9kHz (roughly 2 decades in frequency), a first order filter with a low-pass corner at 100Hz is required. A simple RC filter would require R=15.9MΩ and C=100pF, which are much larger than the passives used in the SAAF (C_{INT}=5pF, and G_M is generated from a 5MΩ resistor).
2. **Low power**: Compared to other filter topologies such as biquad filters, active-RC filters, or switched-capacitor filters, the SAAF is much simpler and requires only one active component (e.g., transconductor) which can be made ultra-low-power.
3. **Programmable gain**: The gain of the SAAF is easily programmable by digitally controlling the value of C_{INT} and G_M.
4. **Robust frequency response**: Any variation in C_{INT} and G_M affects only the gain, and not the frequency response. In fact, the placement of the notches relies only on the clock frequency at which the SAAF is reset. In contrast, with conventional filters, variation in the passives usually results in a change in the shape of the frequency response.

In addition to these three benefits, this work also leverages the SAAF in two other ways. First, the SAAF has the added benefit of filtering any spikes from the switched-capacitor CMFB circuit of the LNA which is switched at 10kHz by design. Second, the input capacitance of the ADC is merged with C_{INT}, eliminating the need for a power-hungry ADC buffer. This will be discussed in Section 2.3.5.

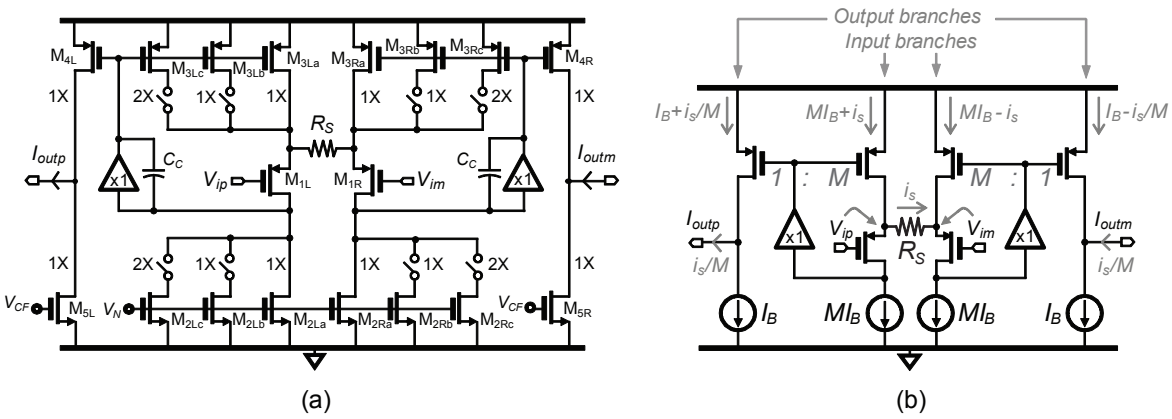

Figure 2-11: (a) Schematic of the fully-differential transconductor used in the SAAF. (b) Simplified schematic of the transconductor shown in (a).

SAAF Transconductor Schematic Design

The full schematic of the transconductor used in this work is shown in Figure 2-11(a). It is based on the implementation in [59] but optimized here for smaller area and low-voltage operation. The circuit uses negative feedback to implement a transconductance with a wide linear range. The feedback loops (one on the left side, one on the right side) comprise the input devices M_{1L} and M_{1R}, the unity-gain buffers, and transistors M_{3L} and M_{3R}. For example, if V_{ip} rises, the drain node of M_{1L} falls and the unity-gain buffer pulls the gate of M_{3L} down, causing the source of M_{1L} to rise accordingly, essentially following V_{ip}. This negative feedback lowers the effective source resistance of $M_{1L,R}$ which operate as source followers, forcing the entire differential input signal $V_i = V_{ip} - V_{im}$ to fall across R_S resulting in a signal current $i_S = V_i/R_S$ flowing through R_S.

A simplified schematic is shown in Figure 2-11(b). Transistors $M_{2L,R}$ and $M_{5L,R}$ act as fixed current sources with value MI_B and I_B respectively, where M=1, 2, 3 or 4. Transistors $M_{3L,R}$ are devices that carry the bias current MI_B from $M_{2L,R}$, as well as $\pm i_S$ depending on the polarity of V_i. The currents of $M_{3L,R}$ get mirrored to $M_{4L,R}$ with a ratio of M:1, and so $M_{4L,R}$ carry a current of $I_B \pm i_S/M$. The output currents are simply $I_{outp,m} = \pm i_S/M$ which is the difference in current between $M_{4L,R}$ and the current sources $M_{5L,R}$.

Given that the differential input signal is $V_i = V_{ip} - V_{im} = i_S R_S$ and the differential output current is $I_{out} = I_{outp} - I_{outm} = 2i_S/M$, the transconductance of this circuit is given by

$$G_M = \frac{I_{out}}{V_i} = \frac{2}{MR_S}. \tag{2.9}$$

By using the current mirror ratio M and R_S=5MΩ, G_M can be set between 100nA/V and 400nA/V with 2 bits. Compared to the design in [59] which uses a switched-resistor up to 20MΩ and a fixed M=1, this implementation reduces the resistor area by 4×.

Three optimizations are made here for ultra-low-voltage operation. First, the G_M of this circuit is made tunable by changing M by using switches that are closer to the supply rails, avoiding the use of switches near mid-rail as in [59][3]. Second, unity-gain buffers are used in the feedback loop (instead of PMOS source followers as in [59]) because the limited headroom does not permit analog level shifting. Lastly, cascoding is avoided to allow for input signals up to 400m$V_{p\text{-}p}$. A consequence of not using cascoding in the output branch is the reduced output impedance of the transconductor. This has the effect of making the filter act more like a leaky integrator (i.e., a first-order filter), which limits the depth of the notches in the SAAF frequency response. However, since the current in the output branch of the transconductor is quite low (10's of nA), the output impedance is still adequate.

2.3.5 Dual-DAC SAR ADC

Oversampling in this system provides benefits such as reduced filtering and dynamic range requirements, as well as the opportunity to use alternative anti-aliasing strategies like the SAAF. The main tradeoff is the increased power consumption in the ADC. However, recent state-of-the-art ADC designs achieve figures-of-merit (FOM) that are less than 10 fJ/conversion-step [71–75], making the ADC power consumption a tiny fraction of the overall system power. For example, a 10 fJ/conversion-step

[3]The implementation in [59] uses a switched-resistor up to 20MΩ with a fixed M=1 in order to make the G_M tunable. This requires mid-rail switches in series with R_S which should be avoided at low voltage because of the large series switch resistance.

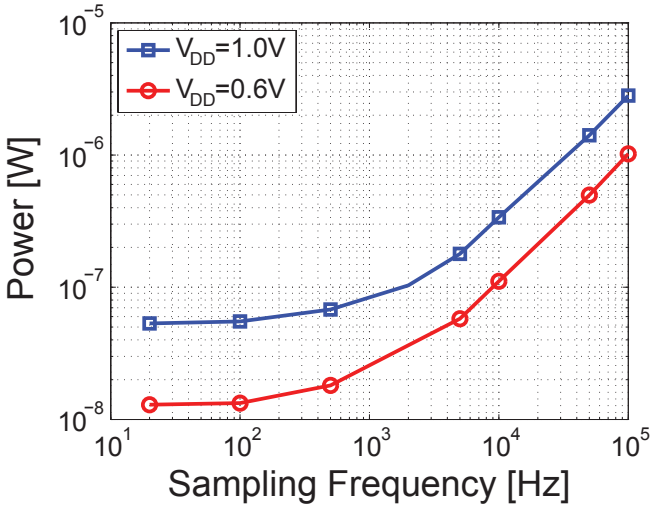

Figure 2-12: Power versus sampling frequency of the SAR ADC from [1] which achieves a minimum FOM of 22.4 fJ/conversion-step.

ADC with an ENOB of 10 bits at 10kS/s would only require $P = FOM \cdot 2^{ENOB} \cdot f_S =$ 100nW of power. Furthermore, low-speed ADCs (below 1kS/s) are typically leakage-dominated and have worse energy-efficiency [1]. This is illustrated in Figure 2-12 which shows the power versus sampling frequency of the ADC in [1], where the power scales linearly with frequency above a few kS/s, but flattens below 1kS/s because of leakage power. At 0.6V, the ADC power at 500S/s and 10kS/s are approximately 20nW and 100nW respectively, representing only a 5× increase in power despite a 20× oversampling ratio. In this case, oversampling at 10kS/s can be achieved with a power overhead of only 80nW.

The ADC FOM typically accounts only for the power consumption of the core ADC circuits. However, in a system, we must also consider the power consumed by the ADC buffer which must drive the ADC input capacitance to the desired accuracy within the allotted settling time. In this work, a dual-DAC 9-bit SAR ADC at 10kS/s is designed to be tightly integrated with the SAAF so that it avoids a power-hungry fast-settling ADC buffer. The ADC block diagram and its interface with the SAAF is shown in Figure 2-13. The input capacitance (of value C_S) of the ADC is added to the 4-bit switched-capacitor (C_D) in the SAAF such that the total integration capacitance for the SAAF is $C_{INT} = C_S + C_D$. To achieve this using a SAR ADC architecture, the

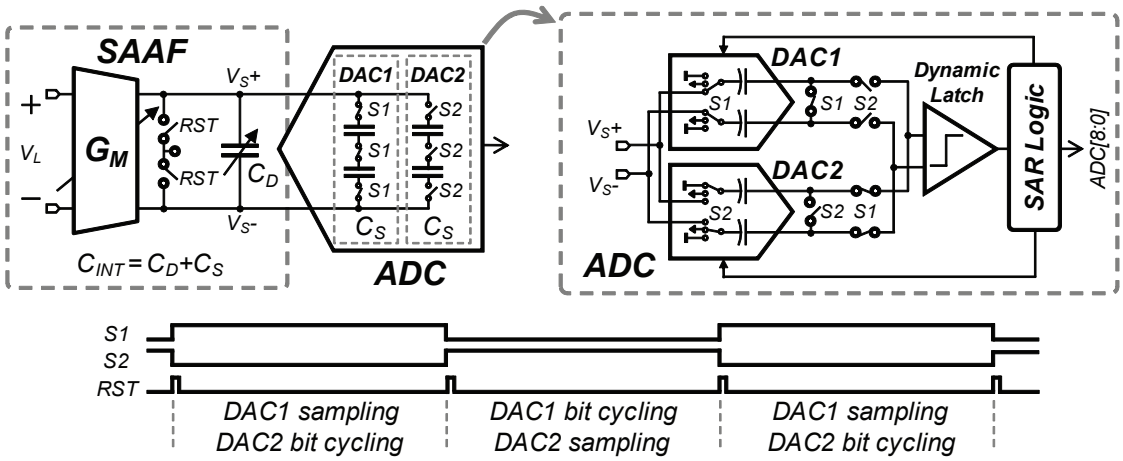

Figure 2-13: Simplified schematic of the interface between the SAAF and ADC. The input capacitance (C_S) of the ADC is merged with the SAAF integration capacitor to eliminate the need for an ADC buffer. The ADC uses dual interleaved DACs to alternate between input sampling and bit cycling.

ADC uses dual capacitive DACs, where the two interleaved DACs alternate between input sampling and bit cycling. For example, while DAC1 is sampling and adding its capacitance C_S to C_D, DAC2 is being bit-cycled and resolving the bits from the previous sample. By interleaving the DACs, the settling requirements of both the sampling and bit cycling phases are relaxed. However, any mismatch in the two DACs can result in spurs in the output spectrum of the ADC [76]. An offset mismatch manifests itself as a spur at $f_S/2$=5kHz, while a gain mismatch results in a spur that is dependent on the signal frequency. From measurements, gain mismatch was negligible and a small offset spur was observed at $f_S/2$. However, it is a non-issue in this system because it is subsequently filtered out by the decimation filter.

Overall, the SAR ADC uses a highly-digital architecture that enables operation at 0.6V. Low-voltage techniques are summarized as follows: a constant-V_{GS} bootstrap circuit is used in the input sampling network to achieve adequate linearity [68] in order to address the issue of degraded *on*-switch conductance at low voltage. The DAC switches are placed near the supply rails in order to improve their conductance. Furthermore, the ADC is fully dynamic (i.e., no static bias currents) which provides very energy-efficient operation. At 10kS/s, measurements confirm that it consumes

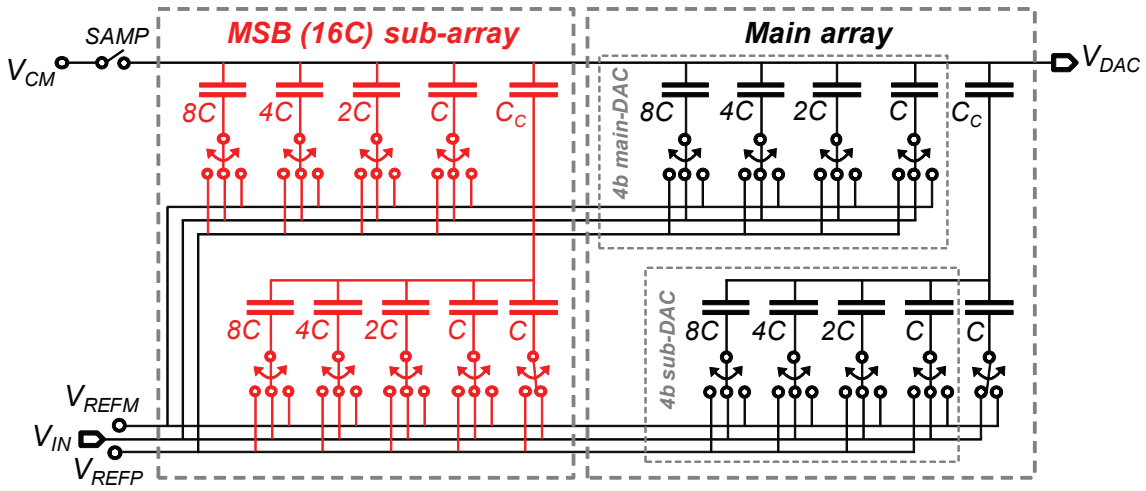

Figure 2-14: Schematic of the DAC in the ADC which takes the form of a split-capacitor array comprising the MSB sub-array and the main array, both segmented into a 4-bit main-DAC and 4-bit sub-DAC. It is drawn as a single-ended DAC for simplicity, but a differential version is implemented in this work.

just 87nW (not including the reference buffer power) thus demonstrating that the power overhead of oversampling is minimal in this system.

DAC Architecture

The SAR ADC is able to achieve ultra-low-power by minimizing the switching energy in the DACs. Figure 2-14 shows the schematic of the DAC used in the ADC and details on its design can be found in [1]. First, it avoids the large ratio between the MSB and LSB capacitors by using a 4-bit sub-DAC. Second, the MSB capacitor in the main DAC is split into a separate MSB sub-array, which is identical in structure to the main array [77]. The split-capacitor array is able to reduce the average switching energy by 37% by improving the efficiency of the "down" transitions during bit cycling. Finally, the positive and negative references for the DAC (V_{REFP} and V_{REFM}) can be set on-chip to provide $V_{REFP} - V_{REFM}$ =200mV, 250mV, 300mV, or 350mV. This is equivalent to providing up to 4.9dB of embedded gain in the ADC.

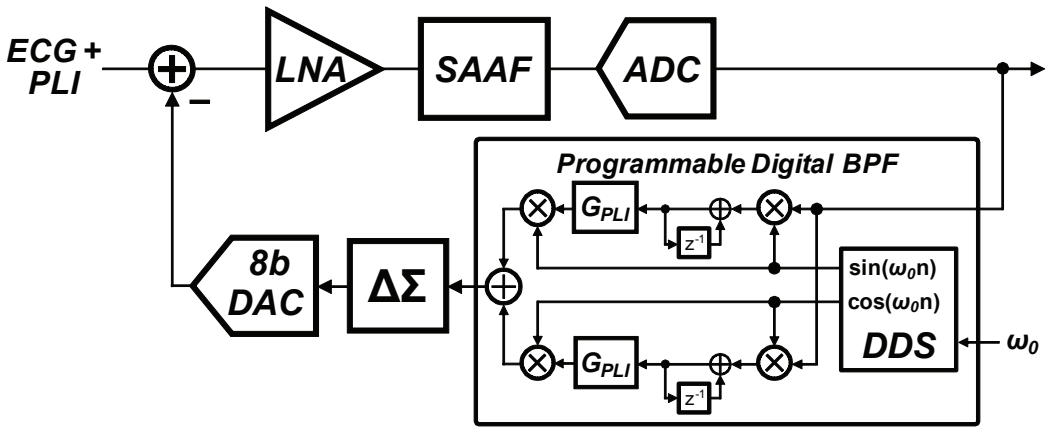

Figure 2-15: Block diagram of the programmable digital band-pass filter in the feedback path of the MSFE.

2.3.6 Programmable Digital BPF Implementation

This section discusses the implementation details of the programmable digital band-pass filter (BPF) which is shown in Figure 2-15. The BPF is based on the design in the first generation system which was implemented on an FPGA using 16-bit multipliers, 32-bit accumulators, and a 16-bit wide SRAM [59]. In this work, the BPF is optimized for area and integrated on-chip to operate at 0.6V. The BPF takes the form of a frequency-translated accumulator, where the input is up-converted through multiplication with $\sin \omega_0 n$ and $\cos \omega_0 n$, accumulated, then down-converted and summed. The center frequency and pass-band width are digitally programmable through ω_0 and G_{PLI} respectively.

In this work, the sinusoids ($\sin \omega_0 n$ and $\cos \omega_0 n$) required for frequency translation are generated on-chip using a direct-digital synthesizer (DDS) with 9-bit outputs. On overview of the DDS operation is as follows: the input to the DDS is a digital word representing the frequency of oscillation (ω_0), and an accumulator computes the summation (i.e., integral) of the frequency which represents the phase ϕ. The phase ϕ is then used as the input address to a look-up table (LUT) which stores an output function such as $\sin(\phi)$ and acts as a phase-to-waveform generator.

Figure 2-16 shows a detailed block diagram of the implemented DDS, where two techniques are used to reduce the size of the mux-based LUT and the overall DDS. The

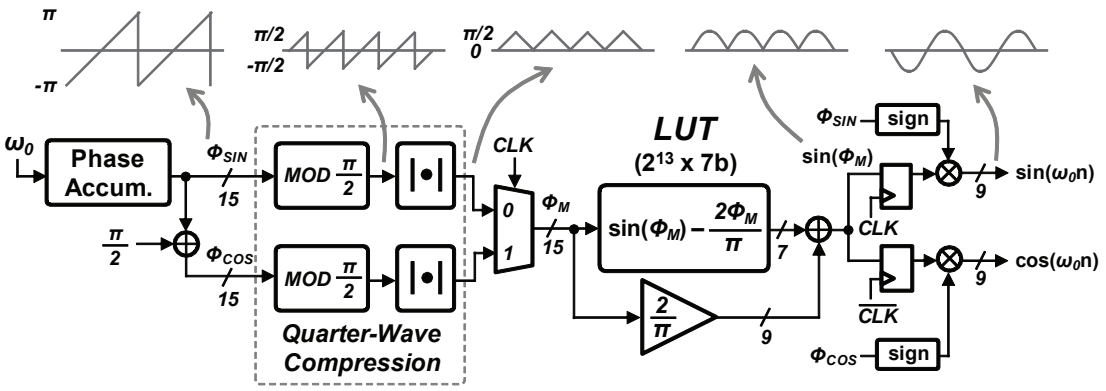

Figure 2-16: Block diagram of the direct-digital synthesizer (DDS) with 9-bit outputs.

input to the DDS is a 13-bit frequency word ω_0, and the phase accumulator generates the 15-bit phase ϕ_{SIN} for the sine waveform. Since a cosine is just a phase-shifted version of sine by a quarter period, $\pi/2$ is added to ϕ_{SIN} to generate the phase for the cosine waveform ϕ_{COS}. Then, both ϕ_{SIN} and ϕ_{COS} ($\in$ [-π,π]) go through a quarter-wave compression block which performs the function $|\phi \bmod \frac{\pi}{2}|$ that effectively limits the phase to just a quarter of the range ($\in [0,\frac{\pi}{2}]$). By doing so, instead of storing values for the entire waveform ($\phi \in$ [-π,π]), only a quarter wave ($\phi \in [0,\frac{\pi}{2}]$) of the waveform needs to be stored by the LUT. This allows the the number of words in the LUT to be reduced by 4$\times$ from 2^{15} to 2^{13}.

The second technique is known as the sine-phase difference technique. Instead of storing sin(ϕ) in the LUT, the sine-phase difference (sin(ϕ) - $2\phi/\pi$) is stored [78]. This is advantageous because its range in magnitude is smaller than sin(ϕ) for ϕ=0 to $\frac{\pi}{2}$ as shown in Figure 2-17, and so the same amount of quantization can be achieved with 2 fewer bits. This allows for a reduction in the stored word length from 9 bits to 7 bits in the LUT. The only overhead with the sine-phase difference technique is the need for an extra adder to add back $2\phi/\pi$ to recover sin(ϕ). Overall, the number of bits in the LUT is reduced by over 5$\times$ from $2^{15} \times 9$ to $2^{13} \times 7$, requiring only 4,000 gates which is suitable for implementation on-chip.

In this work, a single LUT is used for both sine and cosine by multiplexing at the LUT input and de-multiplexing at the LUT output with the clock as the control signal. Since the clock frequency is 10kHz, the extra phase delay in the cosine waveform

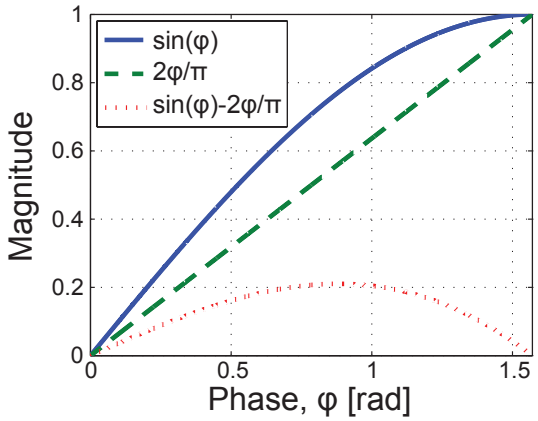

Figure 2-17: Magnitude versus phase for possible functions for the LUT of the DDS.

(typically set to 50Hz or 60Hz) from a half clock cycle in the output de-multiplexer is negligible. Finally, because of the quarter-wave compression of the phase, the output waveform of the LUT is actually $|\sin\phi|$ and $|\cos\phi|$. A simple multiplication with the sign of its phase recovers the complete sine and cosine waveforms as desired.

The frequency resolution of the DDS in this work is limited by the width of the phase accumulator and is given by

$$\Delta f = \frac{f_S}{2^B}, \tag{2.10}$$

where B is the width of the phase accumulator. In this implementation, f_S=10kHz and $B = 15$, and therefore $\Delta f = 0.3$Hz. Despite the limited frequency resolution of the BPF, the width of the filter can be be adjusted with G_{PLI} to compensate.

2.3.7 Digital ΔΣ-Modulation

This section discusses digital ΔΣ-modulation of the feedback DAC that completes the mixed-signal feedback loop for PLI cancellation. The feedback DAC is used to convert the digital 50/60Hz signal from the output of the digital BPF to an analog signal that is applied at the summing node of the LNA [59]. Figure 2-18 shows the single-ended equivalent of the LNA, together with the Thevenin circuit for the feedback DAC, which is a binary-weighted capacitive charge-redistribution DAC. The

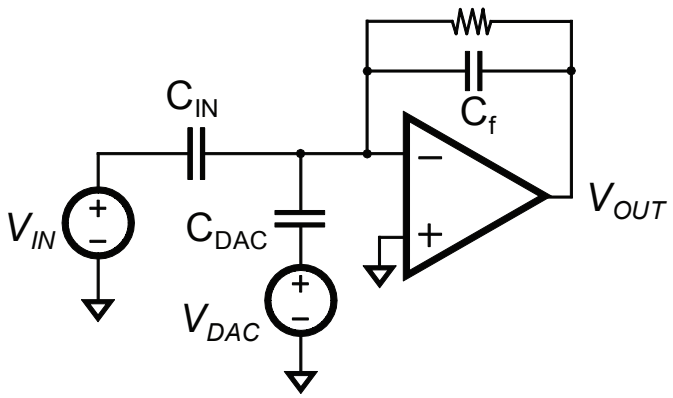

Figure 2-18: Single-ended equivalent of the LNA with the Thevenin equivalent circuit for the feedback DAC.

total capacitance of the DAC is $C_{DAC} = 2^N C_0$, where C_0 is the unit capacitance of the DAC and N is the resolution. The equivalent Thevenin voltage for the DAC is given by

$$V_{DAC} = 2V_{REF,D}\frac{D_{CODE}}{2^N}, \tag{2.11}$$

where $D_{CODE} \in [0, 2^N - 1]$ is the digital input code, $V_{REF,D}$ is the DAC reference voltage, and a factor of 2 accounts for the differential implementation. The input-referred DAC voltage is therefore

$$V_{DAC,in} = \frac{C_{DAC}}{C_{IN}}V_{DAC} = 2V_{REF,D}\frac{C_0}{C_{IN}}D_{CODE} = V_{LSB,in}D_{CODE}, \tag{2.12}$$

where $V_{LSB,in} = 2V_{REF,D}C_0/C_{IN}$ is the input-referred LSB voltage of the DAC. From Equation 2.12, the input-referred peak-to-peak voltage of the DAC for D_{CODE} from 0 to $2^N - 1$ follows as

$$V_{DAC,in,p-p} = 2^N V_{LSB,in} = 2V_{REF,D}\frac{2^N C_0}{C_{IN}} = 2V_{REF,D}\frac{C_{DAC}}{C_{IN}}. \tag{2.13}$$

Equations 2.12 and 2.13 can be used to understand the constraints for the design of the DAC. The first constraint is that $V_{DAC,in,p-p} \geq V_{PLI,max}$, where $V_{PLI,max}$ is the

maximum expected amount of PLI and is equal to $10\text{mV}_{\text{p-p}}$. In this work, $V_{REF,D}$ is set equal to $V_{DD} = 0.6\text{V}$ for convenience, and so the sizing for C_{DAC} must satisfy

$$C_{DAC} \geq C_{IN} \cdot \frac{V_{PLI,max}}{2V_{REF,D}}. \tag{2.14}$$

For C_{IN}=500pF, C_{DAC} was chosen to be 8.32pF for adequate margin.

The second constraint is that the quantization noise of the DAC must be below the noise floor of the LNA. The input-referred noise power spectral density of the DAC is given by

$$S_q(f) = \frac{V_{LSB,in}^2}{12 f_S}, \tag{2.15}$$

and so the total integrated noise power within the signal bandwidth of f_B is given by

$$v_{q,rms}^2 = \int_{-f_B}^{f_B} S_q(f)\mathrm{d}f = \frac{V_{LSB,in}^2}{12}\left(\frac{2f_B}{f_S}\right). \tag{2.16}$$

The total noise power of the LNA was simulated to be approximately 9×10^{-12} V_{rms}^2. Assuming that $v_{q,rms}^2$ must be kept below 10% of the LNA noise power (i.e., $v_{q,rms}^2 < 9\times10^{-13}$ V_{rms}^2), we can combine $V_{LSB,in} = 2V_{REF,D}C_0/C_{IN}$ and Equation 2.16 to calculate a maximum value for the DAC unit capacitance C_0:

$$C_0 \leq \sqrt{3} C_{IN} \frac{v_{q,rms}|_{\max}}{V_{REF,D}} \sqrt{\frac{f_S}{2f_B}}. \tag{2.17}$$

For this work, f_B=156Hz, f_S=10kHz, $V_{REF,D}$=0.6V, C_{IN}=500pF, and $v_{q,rms}|_{\max} = 0.95\mu\text{V}_{\text{rms}}$, and so $C_0 \leq 7.8\text{fF}$.

Equations 2.14 and 2.17 for the minimum C_{DAC} and maximum C_0 determine the required dynamic range for the DAC. For example, values of C_{DAC}=8.32pF and C_0=7.8fF mean that the DAC must have a resolution of at least $N = \log_2(\frac{8320}{7.8}) = 10$. However, a practical issue arises because this required value for C_0 can be below the minimum manufacturable size for capacitors in many typical CMOS technologies. Furthermore, such a small unit capacitance may not be able to meet the matching requirements of a 10-bit DAC. Therefore, another solution is needed to achieve the

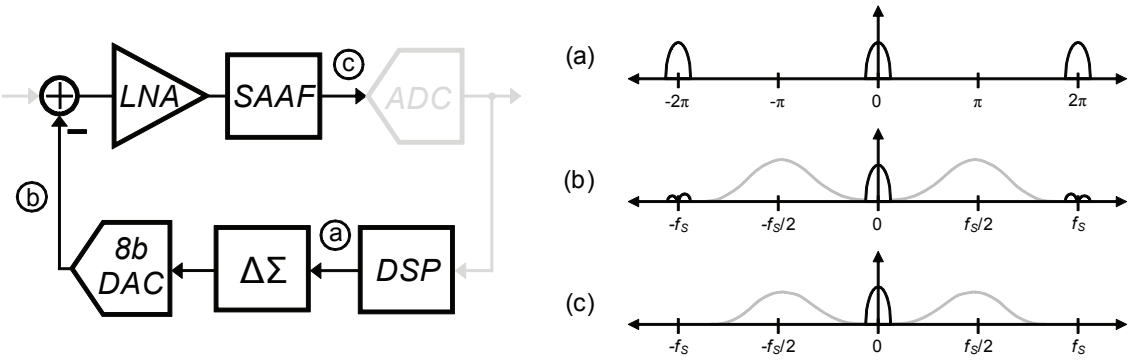

Figure 2-19: The signal chain of the feedback path is shown on the left, with associated spectra on the right. (a) Discrete-time spectrum of the DSP output (out-of-band noise is omitted for simplicity). (b) Continuous-time spectrum after passing through the $\Delta\Sigma$-modulator and feedback DAC. $\Delta\Sigma$-modulation shapes quantization noise (shown in gray) out-of-band, and the zero-order hold action of the DAC attenuates images at multiples of the sampling frequency. (c) Continuous-time spectrum after filtering by the LNA and SAAF. The deep notches of the SAAF remove any residual spectral content at multiples of the sampling frequency prior to sampling by the ADC.

required dynamic range while working within the limitations of the technology.

In this work, a design choice was made to use an 8-bit (N=8) DAC with C_0=32.5fF for adequate matching. Since the DAC alone does not meet the dynamic range requirements, $\Delta\Sigma$-modulation is added to the feedback DAC to increase its dynamic range. Figure 2-19 illustrates the signal spectra at various places in the feedback loop. Fig. 2-19(a) shows the discrete-time spectrum of the DSP output, and Fig. 2-19(b) shows the continuous-time spectrum after passing through the $\Delta\Sigma$-modulator and feedback DAC. The DAC introduces quantization noise shown in gray, but the $\Delta\Sigma$-modulator shapes it to higher frequencies in the spectrum which is made possible by oversampling. The zero-order hold action of the DAC attenuates the images of the original signal from the DSP. Fig. 2-19(c) shows the spectrum after filtering from the LNA and SAAF, where the deep notches of the SAAF remove any residual content at f_S and its integer multiples.

Figure 2-20(a) shows the block diagram of the digital $\Delta\Sigma$-modulator which truncates the incoming data (which can be from 9 to 12 bits wide) down to 8 bits. The process of truncation introduces truncation error modeled by e[n] in Figure 2-20(b).

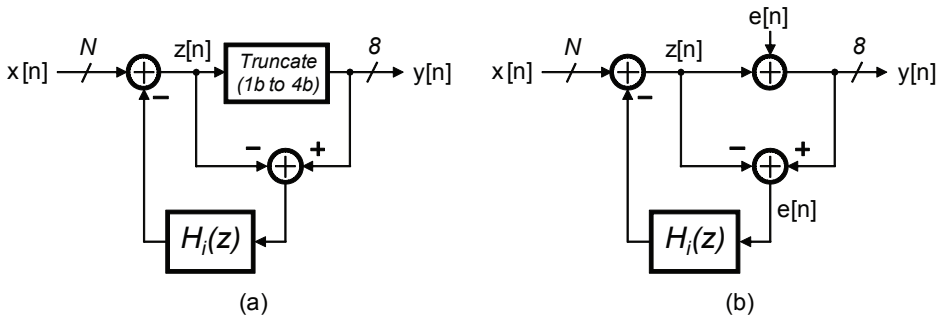

Figure 2-20: (a) Block diagram of the ΔΣ-modulator used in the feedback path. (b) Simplified block diagram where e[n] models the truncation error.

Table 2.3: Reconfigurable transfer function $H_i(z)$ in the ΔΣ-modulator to achieve higher-order noise-shaping.

Order, i	$H_i(z)$	$\lvert N_i(f)\rvert$ $(f \ll f_S)$
First $(i=1)$	$H_1(z) = z^{-1}$	$\lvert N_1(f)\rvert \approx \frac{2\pi f}{f_S}$
Second $(i=2)$	$H_2(z) = 1-(1-z^{-1})^2$	$\lvert N_2(f)\rvert \approx \left(\frac{2\pi f}{f_S}\right)^2$
Third $(i=3)$	$H_3(z) = 1-(1-z^{-1})^3$	$\lvert N_3(f)\rvert \approx \left(\frac{2\pi f}{f_S}\right)^3$

Solving the block diagram yields the following,

$$Y(z) = X(z) + E(z)[1 - H_i(z)] = X(z) + E(z)N_i(z) \tag{2.18}$$

where $N_i(z) = 1 - H_i(z)$ is the i^{th}-order noise transfer function determined by $H_i(z)$. In this work, $H_i(z)$ can be reconfigured to provide 1^{st}-, 2^{nd}-, or 3^{rd}-order noise shaping as shown in Table 2.3.

To see how $N_i(z)$ shapes the spectrum of e[n], using $N_1(z)$ as an example, we can evaluate the frequency response by solving for $|N_1(z)|$ for $z = e^{j\Omega}$ as follows,

$$|N_1(e^{j\Omega})| = |1 - e^{-j\Omega}| = \sqrt{(1-\cos\Omega)^2 + \sin^2\Omega} = \cdots = 2\sin\frac{\Omega}{2}. \tag{2.19}$$

Substituting $\Omega = \frac{2\pi f}{f_S}$, $|N_1(e^{j2\pi f/f_S})| = 2\sin\frac{\pi f}{f_S}$. Since the system is oversampled and $f \ll f_S$ for the frequencies of interest, the noise transfer function simplifies to $|N_1(f)| \approx \frac{2\pi f}{f_S}$ (and in general, $|N_i(f)| \approx \left(\frac{2\pi f}{f_S}\right)^i$). Therefore, with ΔΣ-modulation,

the total integrated noise power within the signal bandwidth is given by

$$v_{qDSMi,rms}^2 = \int_{-f_B}^{f_B} S_q(f)|N_i(f)|^2 \mathrm{d}f. \tag{2.20}$$

Since $|N_i(f)| \approx \left(\frac{2\pi f}{f_S}\right)^i$, the noise density at low frequencies $f < f_B$ is greatly attenuated and $v_{qDSMi,rms}^2$ can be made much lower than the total noise without $\Delta\Sigma$-modulation ($v_{q,rms}^2$). For example, if no $\Delta\Sigma$-modulation is used, the 8-bit DAC with C_0=32.5fF would contribute $v_{q,rms}^2$=1.58×10^{-11} $\mathrm{V}_{\mathrm{rms}}^2$ of input-referred noise to the system (using Equation 2.16). This is considerably larger than the 9×10^{-12} $\mathrm{V}_{\mathrm{rms}}^2$ of noise from the LNA itself. In contrast, from Equation 2.20, the total integrated noise for 1$^{\mathrm{st}}$- and 2$^{\mathrm{nd}}$-order $\Delta\Sigma$-modulation is 5.06×10^{-14} $\mathrm{V}_{\mathrm{rms}}^2$ and 2.92×10^{-16} $\mathrm{V}_{\mathrm{rms}}^2$ respectively (i.e., 0.6% and 0.003% of the LNA noise respectively). Therefore, $\Delta\Sigma$-modulation is used with a low-resolution 8-bit DAC in this work to simultaneously achieve a large PLI cancellation range and maintain low-noise performance.

2.3.8 Decimation Filter

In this system, the decimation filter performs three functions: 1) remove out-of-band noise from oversampling and $\Delta\Sigma$-modulation of the feedback DAC to increase the effective dynamic range, 2) downsample the data by 32× down to the Nyquist rate of 312.5Hz, and 3) remove the spur at $f_S/2$ that arises from any offset mismatch in the interleaved dual-DAC SAR ADC. Figure 2-21 shows the block diagram of the two-stage decimation filter which takes 9-bit data from the ADC at 10kS/s and outputs 12-bit data at 312.5S/s. The first stage is a 3$^{\mathrm{rd}}$-order (N=3) cascaded integrator-comb (CIC) filter that performs downsampling by D=16, and the second is a 19-tap FIR half-band filter that performs the remaining 2× of downsampling.

The CIC filter is a multiplier-less filter that requires only adders and delays and is therefore very hardware-efficient [79]. The integration section has a transfer function of $\left(\frac{1}{1-z^{-1}}\right)^N$ that operates at f_S=10kHz, while the comb section has a transfer function of $(1-z^{-1})^N$ operating at f_S/D. By the downsampling identity [80] shown in Figure 5-10 later in this thesis, the comb section has an effective transfer function of

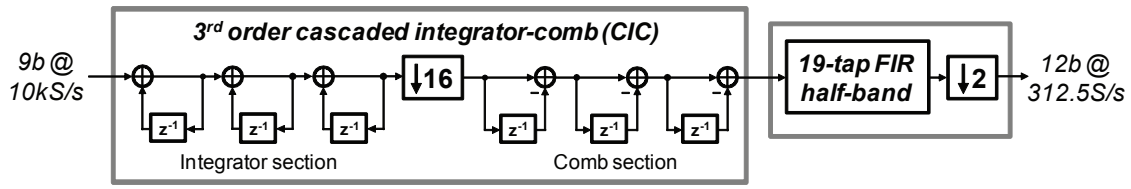

Figure 2-21: Block diagram of the two-stage decimation filter providing 32× downsampling.

$\left(1-z^{-D}\right)^N$ at f_S=10kHz. Therefore, the overall transfer function of the CIC filter (at f_S=10kHz) is given by

$$H_{CIC}(z) = \left(\frac{1-z^{-D}}{1-z^{-1}}\right)^N = \left(\sum_{k=0}^{D-1} z^{-k}\right)^N \tag{2.21}$$

where N=3 and D=16 in this work. From Equation 2.21, it can be seen that the CIC filter is functionally equivalent to a cascade of N FIR moving average filters of length D and therefore has a $SINC^N$-shaped frequency response as shown in Figure 2-22(a). Conceptually, the operation of the digital CIC filter is analogous to the SAAF discussed in Section 2.3.4 which performs its integration in the analog domain. The CIC filter has notches precisely at f_S/D and its integer multiples, which are exactly the frequencies that alias back to baseband following downsampling by D. For large downsampling ratios, CIC filters are much more economical than conventional FIR filters. However, the droop in the pass-band (0 to 150Hz for ECG) does increase with N and D. For this implementation, the droop at 50Hz and 100Hz are -0.27dB and -1.1dB which are acceptable.

Since D is set to 16 in the first stage in order to limit the pass-band droop, a second FIR filter stage is required for the remaining 2× of downsampling. Half-band filters have a transition band at a quarter of the sampling rate, and therefore are ideal for downsampling by 2. All odd filter coefficients (except for the middle coefficient) are zero, enabling N-tap performance with just $(N+1)/2+1$ non-zero coefficients. Furthermore, because it is a symmetric filter, the filter can be folded to reduce the number of multipliers by half. For this implementation, the 19-tap FIR half-band filter requires only 5 multipliers (middle coefficient is 0.5 and is implemented by a

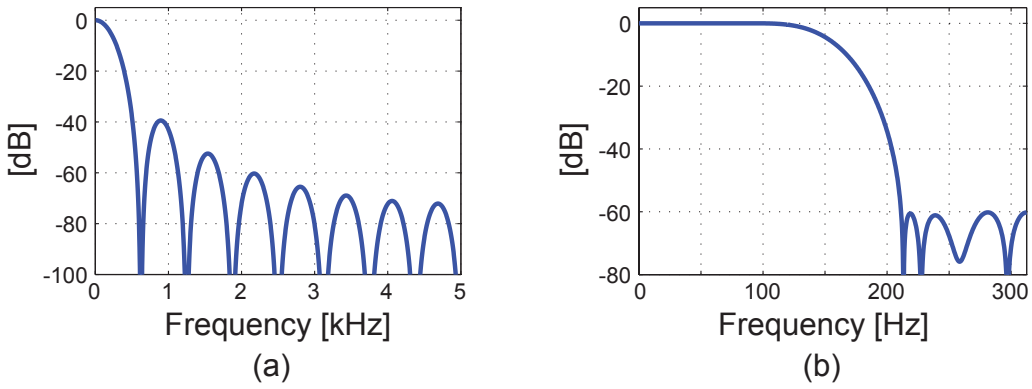

Figure 2-22: Frequency response of the (a) first stage CIC filter at f_S=10kS/s, and (b) second stage half-band filter at f_S=625S/s.

shift) and 10 adders to achieve a 60dB stop-band as shown in Figure 2-22(b).

Overall, due to the hardware-efficient implementation, the decimation filter requires approximately 6,000 gates and consumes just 34nW of power. Combined with the ultra-low-power ADC, the benefits of oversampling are achieved with minimal power overhead.

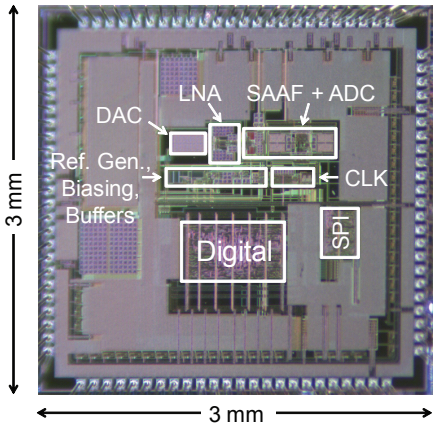

Technology	0.18µm CMOS
Supply voltage	0.6V
Active die area	Analog: 0.25mm^2 Digital: 0.20mm^2 Peripherals: 0.18mm^2

Figure 2-23: Die photo of the MSFE prototype implemented in a 0.18μm CMOS technology

2.4 Prototype Measurement Results

A prototype test chip of the MSFE was implemented in a 0.18μm CMOS technology, and the die photo is shown in Figure 2-23. The entire chip including pads measures

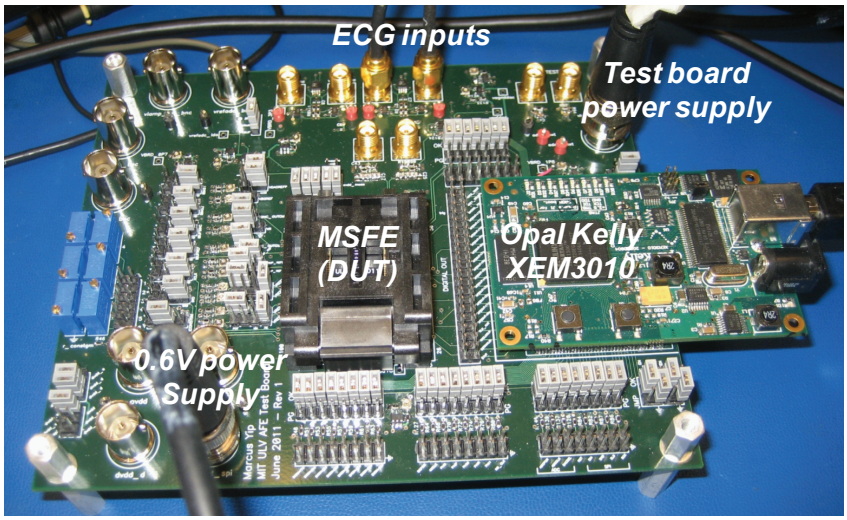

Figure 2-24: Printed circuit board for MSFE testing and characterization.

3mm × 3mm, but the active die area is 0.63mm^2. The chip was packaged in a 100-lead ceramic QFP package and the PCB used for testing and characterization is shown in Figure 2-24. The PCB connects to a daughter board (Opal Kelly XEM3010) which includes a Xilinx Spartan-3 FPGA and USB interface to facilitate communication between the FPGA and a Matlab API on a laptop computer. Matlab was used to control the XEM3010 daughter board in order to program the test chip, acquire output data, and perform analysis.

2.4.1 System Frequency Response

In order to measure the spectra and frequency responses of this mixed-signal system, the outputs of the analog circuits (LNA and SAAF) were buffered with high input-impedance op-amps (Analog Devices AD8603) and digitized by 16-bit 100kS/s ADCs (AD7684 from Analog Devices) on the test PCB. Swept-sine measurements were made by using a 12-bit DAC (Texas Instruments DAC7811) to drive the LNA input through a resistive divider to generate sinusoids with amplitudes from 5 to 10mV$_{\text{p-p}}$ as in [69]. The spectra and frequency responses were calculated by performing FFTs on the digitized output data in Matlab.

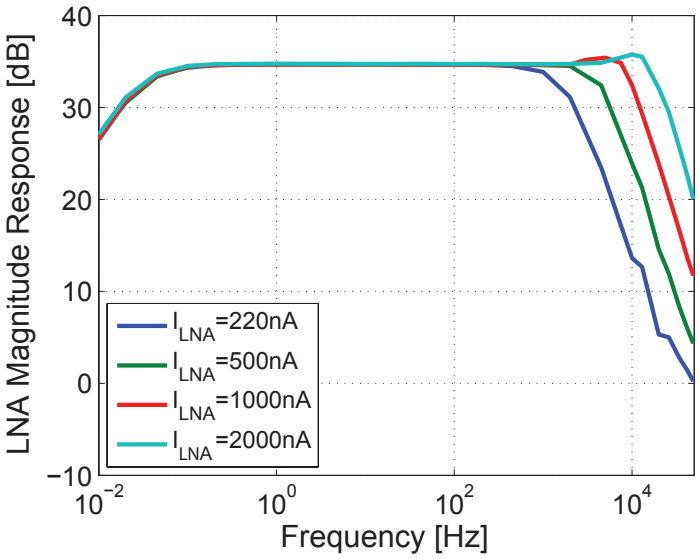

Figure 2-25: Measured magnitude response of the LNA for various total LNA current (V_{CM1}=275mV, V_{CM2}=150mV).

Figure 2-25 shows the measured magnitude response of the LNA for various total LNA current settings (I_{LNA}). The high-pass cut-off for electrode offset rejection is measured to be 20mHz, implying that the pseudo-resistors in the LNA feedback network have a value of 796GΩ which matches well with simulation. The low-pass cut-off is set by the op-amp's poles and increases with increasing I_{LNA}. In this system, I_{LNA} is actually determined by noise requirements, as long as the LNA bandwidth is greater than the ECG signal bandwidth of approximately 100Hz. The mid-band gain is measured to be 34.5dB, set by the ratio of C_{IN} (off-chip ceramic capacitors) to C_f (integrated on-chip).

Figure 2-26 shows the measured frequency response of the system taken at the output of the SAAF (blue line), where the system gain is set to 34.5dB (LNA gain 34.5dB, SAAF gain 0dB). The LNA provides the 20mHz high-pass corner, while the notches at 10kHz and its integer multiples are a result of the SAAF, providing the required anti-aliasing in the system. The sharp notch in the middle of the response is achieved by the mixed-signal feedback loop, with the frequency digitally tuned to 60Hz by setting the DDS frequency. Finally, the solid black line shows the response of the entire system taken at the output of the decimation filter, which limits the

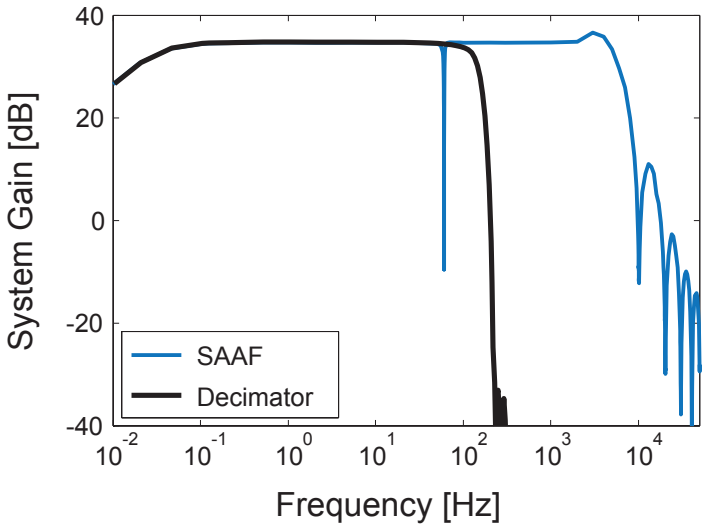

Figure 2-26: Measured frequency response of the system taken at the output of the SAAF (blue line) with the PLI notch digitally tuned to 60Hz. The solid black line shows the response taken at the output of the decimation filter. The SAAF gain is set to 0dB for this measurement.

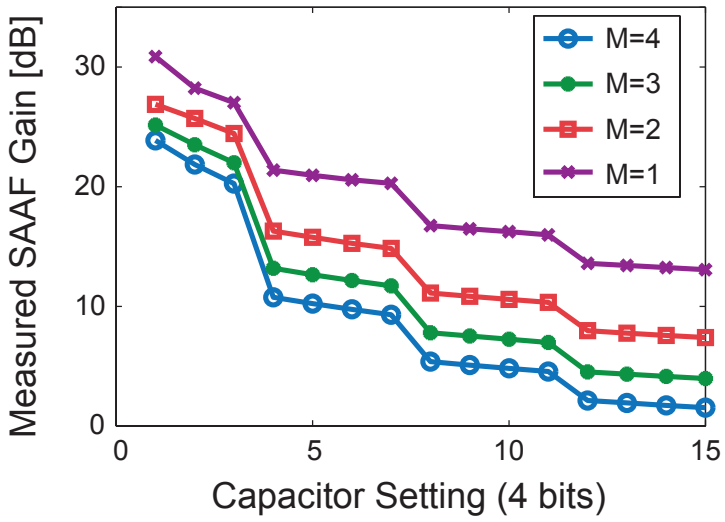

Figure 2-27: Measured SAAF gain for various capacitor and M settings.

system bandwidth to approximately 150Hz, suitable for ECG applications.

In order to accommodate a range of input ECG signal magnitudes (0.1mV to 5mV), the gain of the system can also be adjusted by setting the SAAF gain between 0dB and 30dB. Figure 2-27 shows the measured SAAF gain for various integration capacitor and M settings. Note that this measurement was made with the SAR ADC disabled in a test mode, and so the integration capacitance did not include the ADC input capacitance (500fF). Therefore, the actual SAAF gain in the system is slightly lower than the measured values shown in Figure 2-27. The total gain of the system can be set between 34.5dB and 69.4dB by selecting one of the 60 SAAF gain settings (0 to 30dB) and 4 ADC gain settings (0 to 4.9dB).

2.4.2 PLI Notch Filter

This section presents detailed measurements of the PLI notch filter and the performance of the feedback DAC with $\Delta\Sigma$-modulation. Figure 2-28 shows zoomed-in plots of the PLI notch, illustrating its digital programmability. Fig. 2-28(a) shows the notch digitally tuned to 50Hz or 60Hz by setting the DDS frequency (f_{PLI}), while Fig. 2-28(b) shows how the width of the notch can be digitally tuned by setting the gain of the feedback band-pass filter (G_{PLI}). In this system, the DDS frequency has a resolution of 0.3Hz ($f_S/2^{15}$), and so the notch width can be used to tradeoff selectivity for PLI attenuation.

Figure 2-29 shows the spectral density (input-referred) taken at the LNA output with an $8mV_{p\text{-}p}$, 60Hz input. The gray line and black line show the spectra with the PLI notch off and on respectively. With the PLI notch enabled, the 60Hz spur is attenuated by greater than 40dB, and the quantization noise of the feedback DAC is shaped to higher frequencies by the digital $\Delta\Sigma$-modulator in the feedback path, helping to maintain the same in-band noise floor.

The order of the digital $\Delta\Sigma$-modulator can also be reconfigured. The measurement shown in Figure 2-29 ($8mV_{p\text{-}p}$, 60Hz sinusoidal input) was repeated with different transfer functions in the $\Delta\Sigma$-modulator. Figure 2-30(a) shows the spectral density when no $\Delta\Sigma$-modulation is used with the notch enabled and it is clear that the in-band

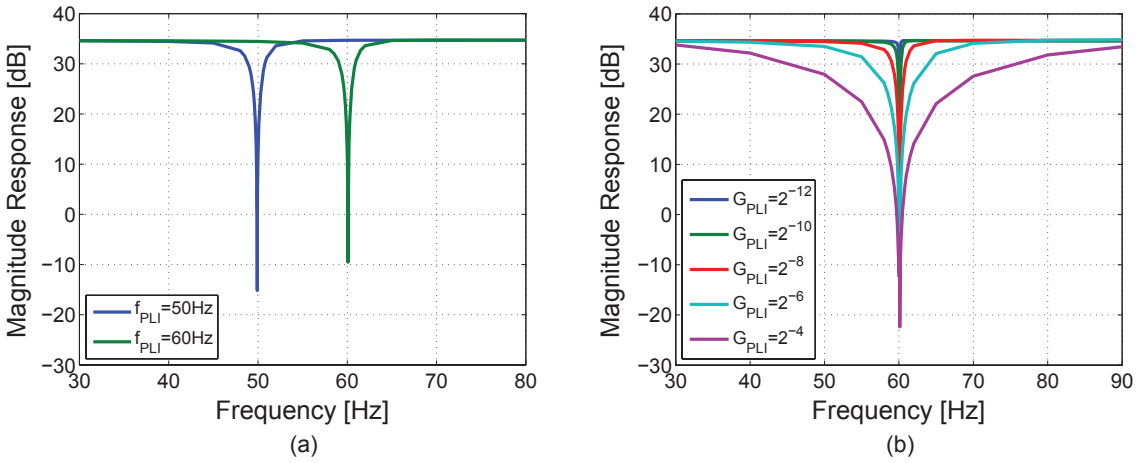

Figure 2-28: Magnitude response at the LNA output with digitally tunable (a) notch frequency (to 50Hz or 60Hz), and (b) notch width (for various G_{PLI} settings). Measurements are taken with the LNA biased with V_{CM1}=275mV, V_{CM2}=150mV, and a total current of I_{LNA}=1000nA.

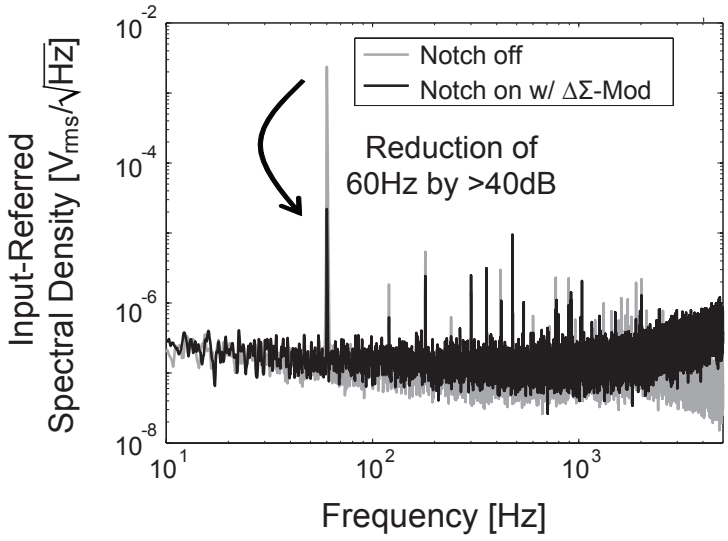

Figure 2-29: Measured input-referred spectral density at the LNA output with an $8mV_{p\text{-}p}$ 60Hz sinusoidal input. When the PLI notch is enabled, the 60Hz spur is attenuated by greater than 40dB, while the quantization noise of the DAC is shaped to higher frequencies.

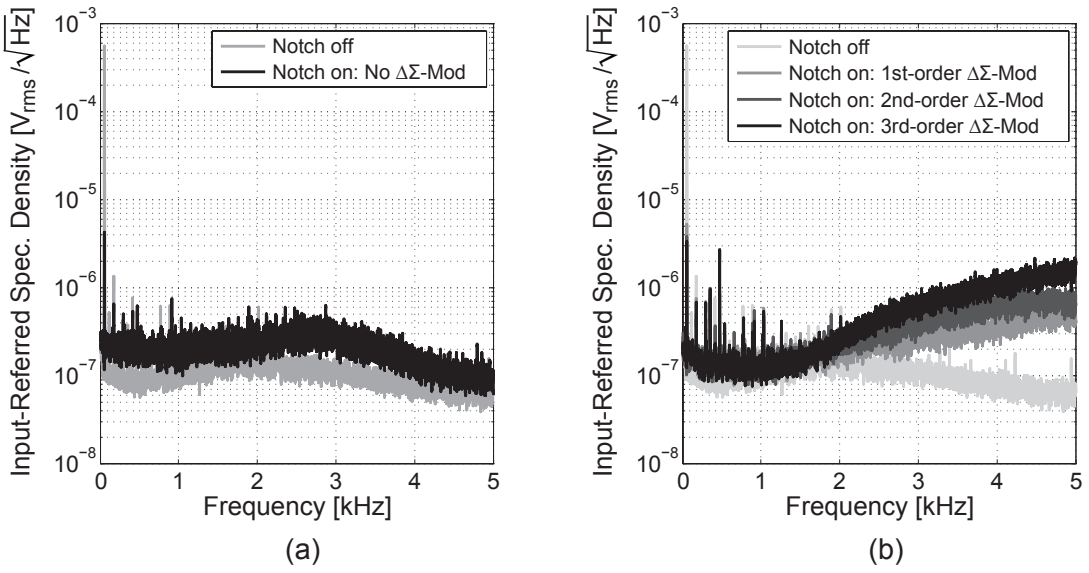

Figure 2-30: Measured input-referred spectral density at the LNA output with an $8mV_{p\text{-}p}$ 60Hz sinusoidal input. (a) Notch is enabled without $\Delta\Sigma$-modulation. (b) Notch is enabled, and 1^{st}-, 2^{nd}-, and 3^{rd}-order $\Delta\Sigma$-modulators are compared.

(0.5Hz to 150Hz) noise floor is significantly increased due to the DAC quantization noise. Figure 2-30(b) shows the spectra when 1^{st}-, 2^{nd}- and 3^{rd}-order modulation are used, and the DAC quantization noise is clearly modulated to higher frequencies.

Finally, Figure 2-31 shows the input-referred spectra (measured at the ADC output) when an input signal of two large tones at 13Hz (desired tone) and 60Hz (interference tone) is applied. In Fig. 2-31(a) with the notch disabled, in addition to the two large tones observed at 13Hz and 60Hz, there are also harmonic distortion spurs at multiples of each of the tones (26Hz, 39Hz, etc., and 120Hz, 180Hz, etc.), as well as intermodulation spurs. These spurs are mainly due to non-linearity introduced by the SAAF, while the non-linearity of the feedback DAC has a lesser effect. Fig. 2-31(b) shows the input-referred spectrum when the 60Hz notch is turned on. In addition to attenuating the 60Hz interference tone by about 40dB while leaving the 13Hz signal tone unaffected, the other spurs due to intermodulation are also removed. The largest remaining distortion is the residual spur at 60Hz, which results from the dithering of the LSB in the DAC. The value of this residual 60Hz spur is determined by DAC reference voltage (0.6V) and the ratio of the DAC unit capacitance to C_{IN} of

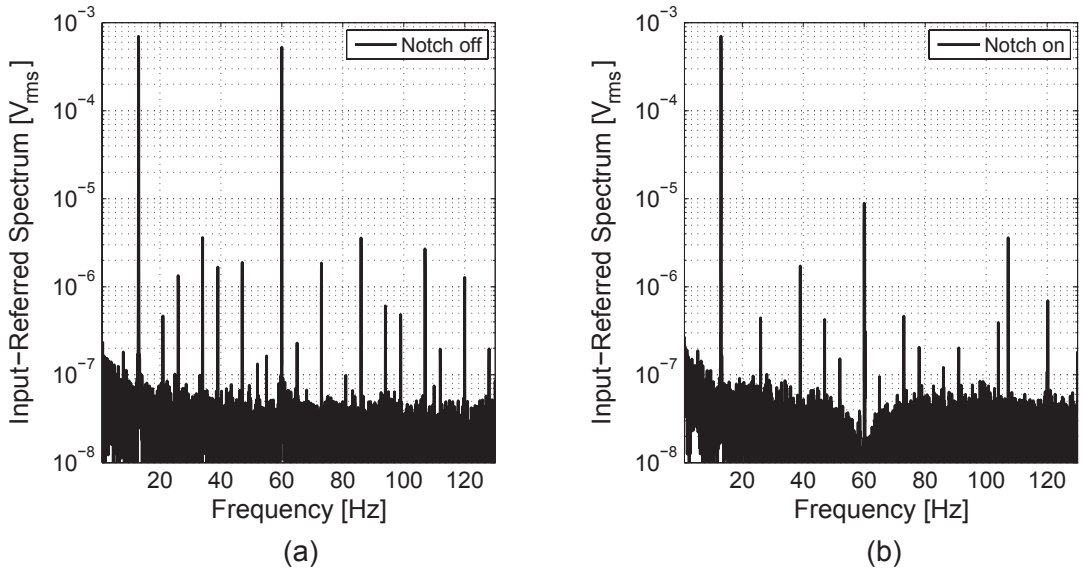

Figure 2-31: Measured input-referred spectra at the output of the ADC with two tones at 13Hz and 60Hz with the PLI notch (a) off and (b) on. $G_{PLI} = 2^{-6}$ and f_{PLI}=60Hz for this measurement, and a Hanning window is applied to the data for plotting purposes.

the LNA. In this system, the residual spur is $26\mu V_{rms}$ when input-referred, or about $1.38mV_{rms}$ at the LNA output which is much smaller than the expected ECG signal.

Linearity measurements of the feedback DAC were also made to verify that it did not limit the system. Figure 2-32 shows the measured INL and DNL of the feedback DAC. Over the 5 die measured, the average INL and DNL were -0.31LSB/+0.77LSB and -0.53LSB/+0.52LSB respectively.

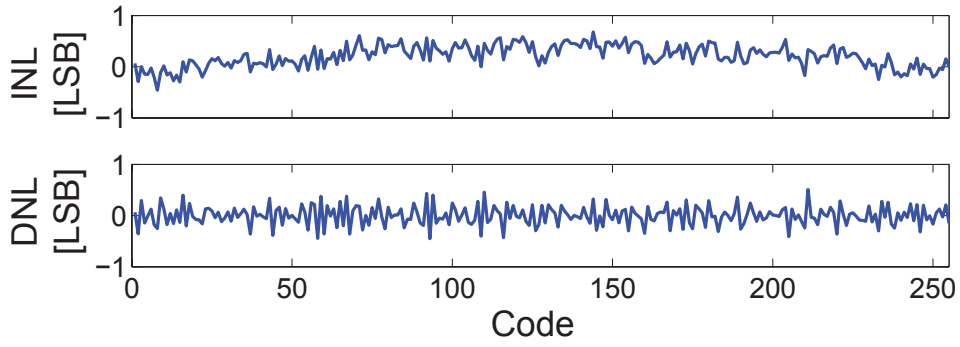

Figure 2-32: Measured static linearity of the charge-redistribution feedback DAC.

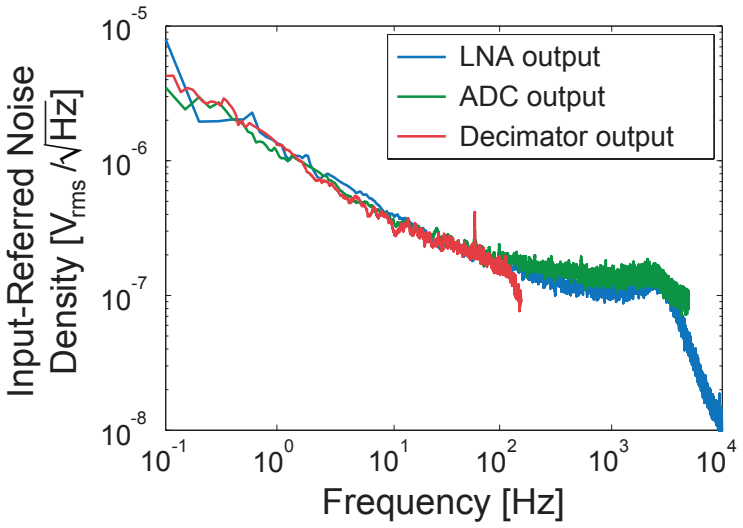

Figure 2-33: Measured input-referred noise density of the system at the output of the LNA, ADC and decimation filter (LNA biased with V_{CM1}=260mV, V_{CM2}=125mV, and a total current of I_{LNA}=650nA).

2.4.3 Noise Performance

In Appendix A.3, it is determined that the ANSI/AAMI specification for noise in ambulatory ECG applications ($30\mu V_{p\text{-}p}$ over 10 seconds in a 150Hz bandwidth) translates to a RMS noise specification of $4.03\mu V_{rms}$ when the output data rate is 500Hz. In this system, the output data rate is 312.5Hz, and therefore the equivalent noise specification is actually $4.17\mu V_{rms}$ ($30\mu V_{p\text{-}p}/7.197\sigma$ to ensure with confidence that the 3,125 samples in a 10 second window do not exceed $30\mu V_{p\text{-}p}$). Figure 2-33 shows the measured input-referred noise density of the system, taken at the outputs of the LNA, ADC, and decimation filter. The noise in this system is dominated by the $1/f$-noise of the input devices in the op-amp of the LNA. At the output of the decimation filter, the total integrated noise is $3.44\mu V_{rms}$ in a 156Hz bandwidth, which is within the ANSI/AAMI specification.

Considering the LNA only, the total input-referred noise integrated from 0.5Hz to 50kHz is $9.26\mu V_{rms}$. Using the common noise efficiency factor (NEF) [81] metric to normalize the total input-referred noise to that of a single BJT with the same current, the LNA on its own achieves a NEF of 5.32 (based on a 2.93kHz bandwidth). Using

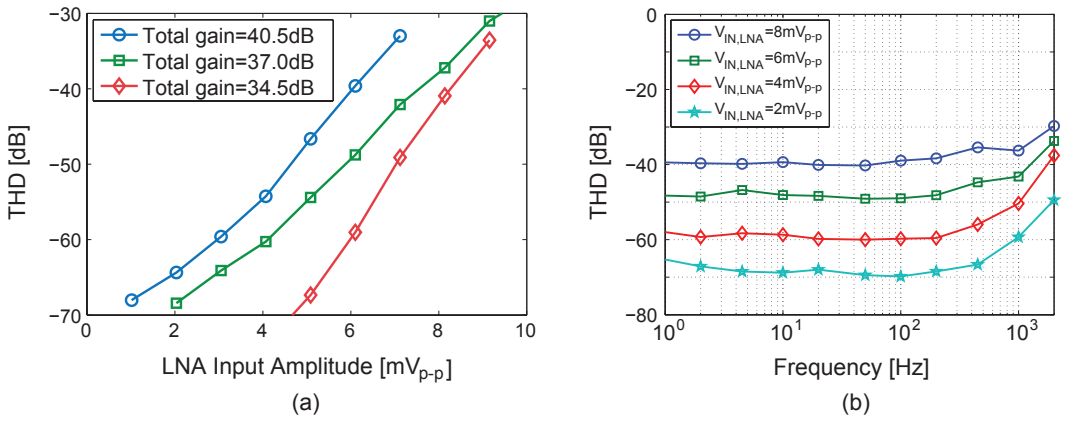

Figure 2-34: Measured THD at the SAAF output versus (a) the LNA input amplitude at a signal frequency of 50Hz, and (b) the signal frequency for various LNA input amplitudes. The measured data in (b) is taken with a total gain of 37dB.

the more recently introduced power efficiency factor (PEF) [26] to account for the impact of supply voltage, the LNA achieves a PEF of 17 which is very competitive with the state-of-the-art (which is typically >20).

2.4.4 Dynamic Range

The dynamic range of a system is the ratio of the largest signal that can be processed without significant distortion, to the minimum detectable signal (usually set by the noise floor). In the previous section, we saw that the measured total input-referred noise of the system was $3.44\mu V_{rms}$. Here, measurements of the total harmonic distortion (THD) are presented in order to determine the largest signal that can be processed by the MSFE. Figure 2-34(a) shows the measured THD versus the system input amplitude for a signal frequency of 50Hz, and Figure 2-34(b) shows that the measured THD is also maintained across the signal frequencies of interest for ECG applications (up to a few hundred Hz). From Fig. 2-34(a), at the lowest gain setting of 34.5dB, the MSFE can accommodate an input of $8mV_{p\text{-}p}$ with better than -40dB (i.e., 1%) THD which is better than the IEC specification of $6mV_{p\text{-}p}$. Combining the maximum allowable input signal of $8mV_{p\text{-}p}$ ($2.83mV_{rms}$) with the minimum detectable signal of $3.44\mu V_{rms}$ (set equal to the input noise), this translates to 58dB of input dynamic range despite just 0.6V supply operation.

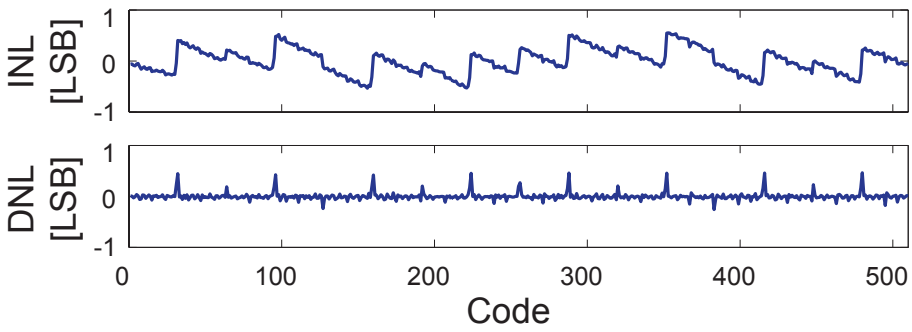

Figure 2-35: Measured static linearity of the dual-DAC SAR ADC.

2.4.5 ADC and Decimation Filter

The static linearity of the dual-DAC SAR ADC was measured across 16 samples, and the average INL and DNL was -0.54LSB/+0.55LSB and -0.23LSB/+0.48LSB respectively. The DAC capacitors were carefully laid out in a common-centroid configuration to ensure good ratiometric matching, and a shield was placed in between the top-plate and bottom-plate routing to minimize non-linearity due to capacitive coupling. Figure 2-35 shows the INL and DNL of a representative part. In the fabricated test chip, the dominant source of non-linearity was the large positive DNL spike every 64 codes, indicating an issue with the MSB/4 capacitor in the DAC (likely due to parasitic routing capacitance). Smaller positive DNL spikes were observed every 32 codes due to parasitic capacitance on the top-plate of the sub-DAC resulting in a systematic mismatch with the main-DAC. Since the non-linearity is dominated by systematic errors (and not random mismatch), the linearity can be improved with better layout and/or adjustment through post-layout extraction, or calibration in the digital domain. Regardless, both INL and DNL are well within ± 1 LSB which is adequate for this system.

The dynamic performance of the ADC at a sampling rate of 10kHz is shown in Figure 2-36(a), where the input frequency was swept from DC up to the Nyquist rate of 5kHz. The SFDR and -THD (in dB) are both maintained above 60dB. The low frequency SNDR is 52.4dB (ENOB of 8.41 bits), dropping to 50dB (ENOB of 8.01 bits) at Nyquist, resulting in an ADC figure-of-merit of 37.3 fJ/conversion-step [82].

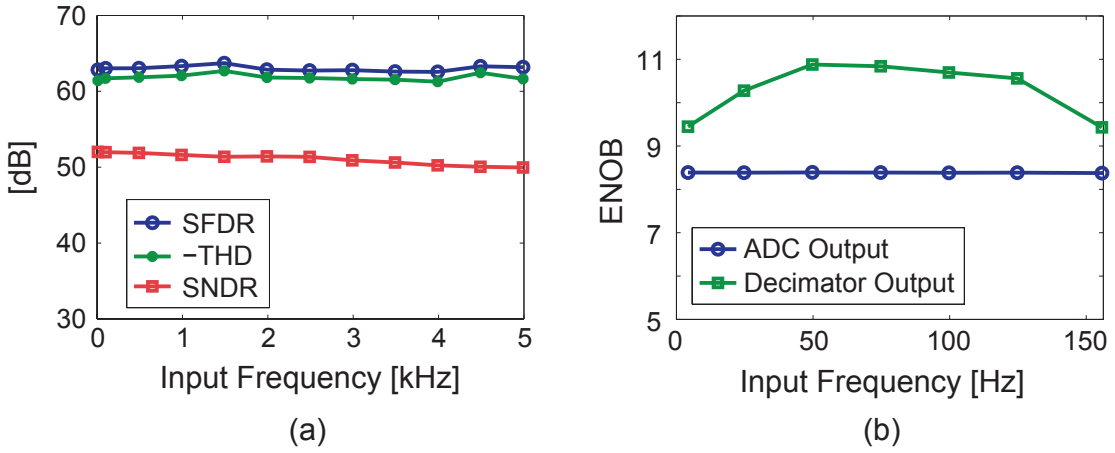

Figure 2-36: (a) Measured dynamic performance of the dual-DAC SAR ADC. (b) Measured effective number of bits at the ADC and decimation filter outputs.

Since the system is oversampled at 10kHz, a decimation filter is necessary for reducing the output data rate (by 32×) to 312.5Hz which is more appropriate for ECG applications. To do so, it removes the out-of-band noise and consequently improves the effective resolution of the system. Figure 2-36(b) shows how the decimation filter improves the ENOB of the system from 8.4 bits at the ADC, to a peak of 10.8 bits at the decimator output. For the decimation filter, the ENOB at low frequency is degraded due to harmonics that lie within the signal band (up to 156Hz), while the high frequency ENOB is degraded because of droop in the CIC filter which attenuates the fundamental tone and limits the SNDR.

2.4.6 Other Front-End Performance Metrics

In Section 2.2.1, specifications for the common-mode rejection, input impedance, electrode offset tolerance, and PLI tolerance were outlined. This section presents the measured results for these front-end metrics.

Common-Mode Rejection

The common-mode gain of the LNA was measured by tying both inputs of the LNA together, and driving it with a $100\text{mV}_{\text{p-p}}$ sinusoid swept from 1Hz to 1kHz. The

CMRR of the LNA was measured across 12 parts from 1Hz to 1kHz, and the average CMRR was 70.4dB which is close to the theoretical value of 73dB. As expected, the CMRR in this system is limited by the mismatch of the LNA input capacitors (C_{IN}) which are off-chip ceramic capacitors.

Input Impedance

Both the common-mode and differential-mode input impedance were measured using the procedure outlined in Section 4.6.4 of [69]. Recall that the common-mode input impedance Z_C must be much larger than the electrode impedance mismatch (typically 10's of kΩ) in order to limit the amount of differential-mode interference converted from a common-mode source (particularly at the line frequency of 50/60Hz). On the other hand, the differential-mode input impedance Z_D must be much larger than the absolute electrode impedance (10's to 100's of kΩ for wet electrodes, a few MΩ for dry electrodes) in order to avoid attenuation of the differential input signal. Z_D is typically measured at the lowest frequency of the signal band because that is where the electrode impedance is usually the largest.

The measured value of Z_C was 179MΩ at 60Hz, which is within the ballpark of the theoretical value of 265MΩ. The discrepancy is likely due to the parasitic capacitance of the bond pads and PCB. The measured value of Z_D was 1.18GΩ at 0.5Hz which is quite close to the theoretical value of 1.27GΩ.

Interference Tolerance

The measured DC electrode offset tolerance was well beyond the IEC specification of $\pm$300mV which is a result of the system being fully *ac*-coupled through the LNA input capacitors. The maximum tolerance for PLI is somewhat arbitrary, but here, it is defined as the amount of PLI at the input that can be tolerated while still keeping the PLI at the output of the system to less than 2% of the magnitude of the desired signal. The measured tolerance for PLI was about 12.6m$V_{p\text{-}p}$ with a feedback BPF gain setting of $G_{PLI} = 2^{-6}$, which allows the system to be able to handle worst case scenarios for PLI coupling.

Table 2.4: Measured breakdown of power consumption in the MSFE at 0.6V.

Single Channel	Low-noise amplifier	390nW
	SINC anti-aliasing filter	290nW
	ADC @ 10kS/s	87nW
	Feedback DAC	2.5nW
	Digital PLI notch filter	180nW
	Digital decimation filter (by 32×)	34nW
	Leakage	166nW
Total power for single channel (PLI notch off)		0.97μW
Total power for single channel (PLI notch on)		1.15μW
Peripheral Circuits	Current reference and bias network	80nW
	400mV supply-indep. reference and reference buffers	1.32μW
	480kHz oscillator and clock generation	310nW
Total peripheral power		1.71μW
Total power (PLI notch off)		2.68μW
Total power (PLI notch on)		2.86μW

2.4.7 Power Consumption

The measured breakdown of power consumption in the MSFE is listed in Table 2.4. All core circuit blocks in the MSFE operate from a 0.6V supply, and the supply current was measured using a Keithley 2400 SourceMeter. A separate 1.8V test supply was used in the I/O ring for test purposes (to interface to the Opal Kelly XEM3010) and the power of the digital level shifters to 1.8V was not included.

At 0.6V, the single-channel MSFE consumes 1.15μW, and the integrated peripheral circuits consume 1.71μW. In contrast, considering just the power of the digital circuits alone, the same implementation at a supply voltage of 1.5V as in [59] would require greater than 3μW of additional digital power, demonstrating the effectiveness of voltage scaling on digital power consumption. The impact of voltage scaling on analog circuits is less straightforward because the supply voltage impacts dynamic range. However, if dynamic range is ignored for now, then analog power scales linearly with supply voltage to the first order. Based on this rather crude assumption, the same implementation at 1.5V would also require 3μW of additional analog power.

2.4.8 Comparison with Recent Work

A comparison of the MSFE with the recent state-of-the-art is provided in Table 2.5. This work achieves lower voltage operation and favorable power consumption

Table 2.5: Comparison of MSFE with recent state-of-the-art. [A]This is the action potential ("spike") recording application in [26]. [B]This is the noise integration bandwidth. [C]This is the maximum input range with an acceptable level of distortion. [D]If available, the ENOB is listed. [E]Not included in the power and area comparison.

	Van Helleputte, ISSCC 2012 [28]	Bohorquez, JSSC 2011 [59]	Muller, JSSC 2012 [26]	This work [27]
Application	ECG	ECG	[A]Spike	ECG
Technology	0.18μm	0.18μm	65nm	0.18μm
Supply voltage	1.2V	1.5V	0.5V	0.6V
Power/channel (blocks included)	17μW (LNA, BPF, ADC@500S/s)	0.86μW (LNA, SAAF)	5.04μW (LNA, BPF, ADC)	1.15μW (LNA, SAAF, ADC/DEC, PLI)
Input-ref. noise ([B]bandwidth)	1.3μV_{rms} (100Hz)	3.4μV_{rms} (100Hz)	4.9μV_{rms} (10kHz)	3.44μV_{rms} (156Hz)
NEF/PEF	n/a	8.3/103	5.99/17.96	5.32/17
[C]Input range	-	10.8mV_{p-p}	3.5mV_{p-p}	8mV_{p-p}
[D]Resolution	12-bit ADC	no ADC	7.15 ENOB	10.8 ENOB
EO tolerance	>1V	>$\pm$300mV	$\pm$50mV	>$\pm$300mV
PLI tolerance	-	5mV_{p-p}	-	>10mV_{p-p}
Motion artifact removal	[E]Yes	No	N/A	No
Additional features	[E]Contact impedance	-	-	-
Die area/channel	Est. >3mm^2	0.225mm^2	0.013mm^2	0.445mm^2

compared to previous ECG front-ends [28, 59][4]. When compared to a 0.5V front-end for neural spike recording [26], this work has favorable input range and effective output resolution. In terms of noise-efficiency, this work also achieves competitive NEF and state-of-the-art PEF because of aggressive voltage scaling down to 0.6V. Lastly, this work achieves good electrode offset and PLI tolerance, in addition to meeting all the target requirements in Table 2.2 for ambulatory ECG systems.

2.5 Physiological Measurements

The MSFE prototype was used to make real ECG measurements on a single male subject. Figure 2-37 shows the test chip being used to measure the 3 standard bipolar limb leads with 3 separate measurements. Lead I, II, and III are defined as LA-RA, LL-RA, and LL-LA respectively (where LA, RA, and LL are left arm, right

[4]The power consumption of the front-end in [59] accounts for only the LNA and SAAF. The MSFE in this work is completely integrated on chip, including the ADC, decimation filter, and PLI cancellation loop.

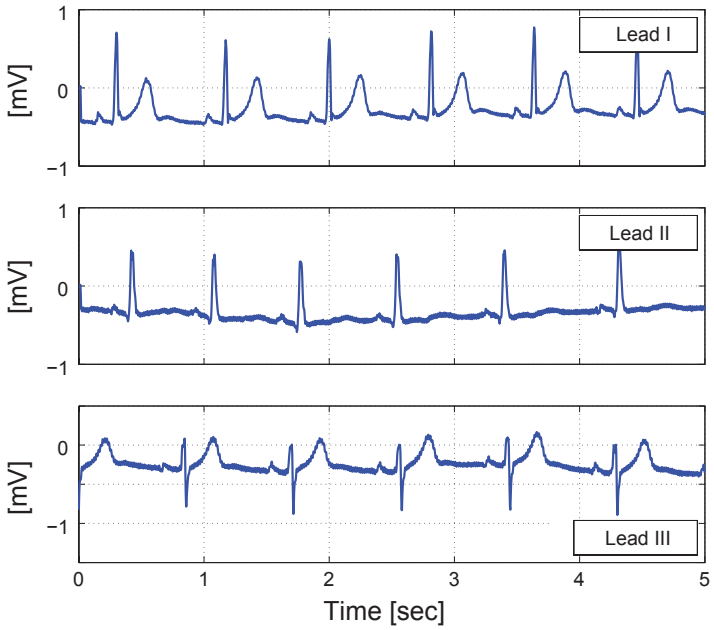

Figure 2-37: Measured ECG (input-referred) using the prototype chip with standard Ag/AgCl electrodes and 2' long unshielded cables in Lead I, II, and III configurations. All measurements were taken with a third patient ground reference electrode to bias the body.

arm, and left leg respectively). All measurements in this section were made with standard Ag/AgCl gel electrodes, 2' long unshielded cables, and a third patient ground reference electrode that biases the body at ground.

In order to emulate a worst case PLI coupling scenario, the input impedance and CMRR of the LNA were intentionally degraded by shunting one of the inputs to ground with a 500kΩ resistor. With the PLI notch off, a large amount of 60Hz interference is coupled onto the ECG, actually causing the QRS peak to clip as shown in Figure 2-38 (top). When the PLI notch is enabled, the PLI is clearly removed and a clean ECG is recovered as shown in Figure 2-38 (bottom).

For wearable ambulatory ECG systems, it is desirable to minimize the size of the device for comfort, wearability, and perhaps concealability. This places constraints on the inter-electrode distance, which has a significant impact on the ECG signal magnitude, as well as the observable features in the ECG. Figure 2-39(a) shows the measured Lead I ECG with the electrodes separated by 21cm, 12cm, 7cm, and 3cm.

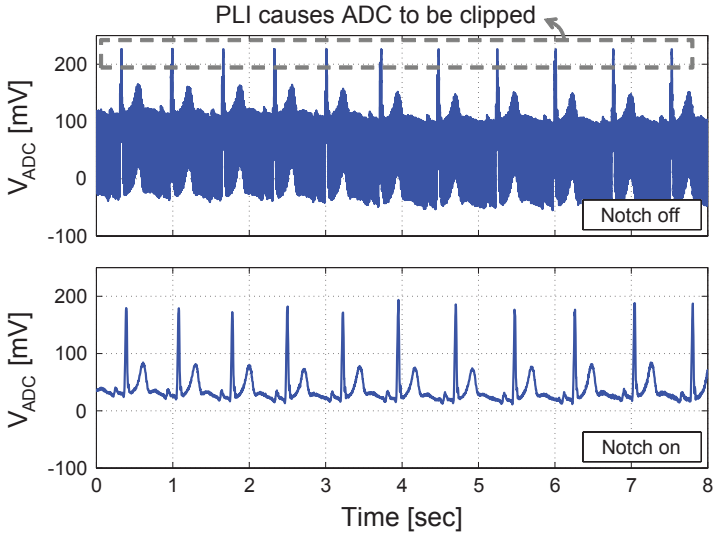

Figure 2-38: Measured Lead I ECG (referred to ADC input) with the notch off (top) and notch on (bottom). For extreme cases where the PLI magnitude is large enough, the PLI can cause clipping in the ADC.

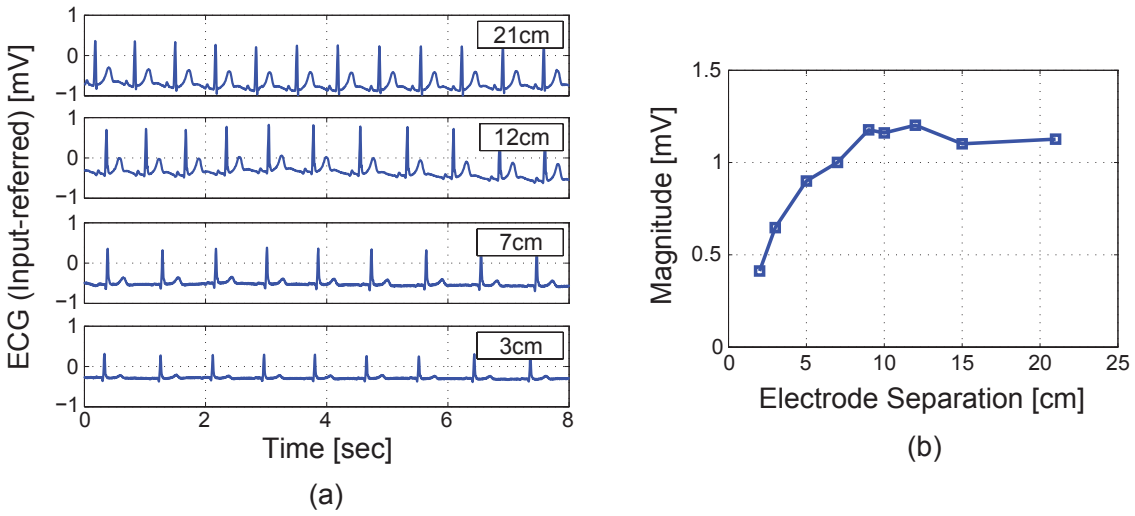

Figure 2-39: (a) Measured Lead I ECG (input-referred) with electrodes separated by 21cm, 12cm, 7cm, and 3cm. (b) Measured magnitude of the QRS peak in Lead I configuration versus the electrode separation.

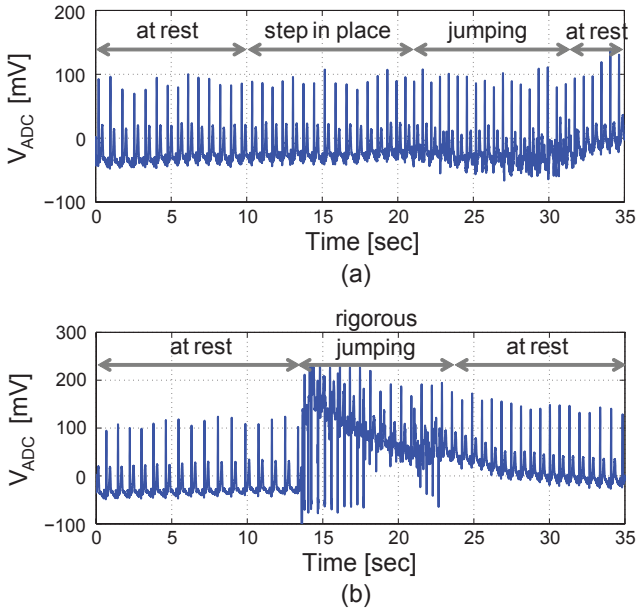

Figure 2-40: Measured ECG (referred to ADC input) with (a) light to moderate activity, and (b) rigorous activity.

At 21cm and 12cm, the P, QRS, and T waves are clearly visible. However, the P wave is lost at 7cm, and only the QRS complex is observable at 3cm. Figure 2-39(b) shows the magnitude of the ECG versus the electrode separation. This data suggests that the electrodes should be at least 10cm apart for good signal quality. However, depending on the type of analysis, the separation may be decreased. For example, rhythm analysis based on the QRS complex may still be possible at 3cm.

Aside from electrode DC offset and 50/60Hz PLI, the other major aggressor plaguing ambulatory ECG systems is the motion artifact that arises when a person is moving. Any physical motion such as walking or running may cause changes in the charge interface of the electrode-skin interface, as well as deformation in the skin. Either effect can inject a significant amount of artifact in the ECG signal. Figure 2-40(a) shows ECG measurements made during light to moderate activity, while Figure 2-40(b) shows an ECG during rigorous activity. Overall, the MSFE is able to perform robustly only for light activity like slow to moderately paced walking. However, since no special efforts were made to mitigate motion artifact in this work, the MSFE is not able to handle rigorous activity.

2.6 Chapter Summary

Symptoms of cardiovascular disease are often very intermittent, necessitating ultra-low-power wearable ECG monitors with long lifetimes. In order to minimize the system power of the entire wearable sensor node, DSPs can accomplish feature extraction and data compression which reduces the power burden of data storage or transmission. Recent biomedical DSPs have leveraged voltage scaling (down to 0.5V) to improve their energy-efficiency. Additional size and power benefits can be obtained by integrating the analog front-end with the DSP back-end. However, current AFEs rely on higher supply voltages in order to perform signal conditioning and accommodate aggressors like electrode offset and PLI. To address these issues, this chapter presented a mixed-signal front-end that leverages a highly-digital architecture and analog circuits optimized for low-voltage in order to operate from a 0.6V supply without sacrificing robustness and dynamic range. Aggressive voltage scaling down to 0.6V also improves power-efficiency, and facilitates integration with low-voltage DSPs for future SoC development.

The highly-digital architecture of the MSFE in this work is based on the first generation of the system designed by Jose Bohorquez [59], but improvements and new contributions are made on several fronts. Compared to [59], the supply voltage is reduced from 1.5V to 0.6V by optimizing the analog circuits for ultra-low voltage by following four general principles for low-voltage analog design. An estimated savings of greater than 5μW is obtained by scaling from 1.5V to 0.6V. $\Delta\Sigma$-modulation of the feedback DAC is added to improve the interference cancellation dynamic range which is crucial at low voltage. An ultra-low-power dual-DAC SAR ADC is designed to provide the oversampling without burning excessive power, and a digital decimation filter is integrated to filter out high frequency noise, reduce the data rate, and improve the effective resolution of the system. The MSFE in this work also provides on-chip low-voltage peripheral circuits like current/voltage references and clock generation, making it a completely integrated front-end solution from a single 0.6V supply.

Table 2.6 provides a summary of all the relevant performance metrics for the

MSFE. The entire MSFE consumes just 2.9μW from 0.6V, which is two orders of magnitude lower than commercially available parts.

Table 2.6: Measured performance summary of the MSFE. [1]The input-referred noise includes all contributions from the LNA, SAAF, ADC, and decimation filter in a 156Hz bandwidth. [2]The NEF and PEF calculations are based on the input-referred noise of the LNA only, integrated from 0.5Hz to 50kHz, using a 2.93kHz 3dB bandwidth.

Technology	0.18μm CMOS
Supply voltage	0.6V
Total power	2.86μW
Bandwidth	0.02-156.25Hz (decimate by 32) 0.02-312.5Hz (decimate by 16)
Gain	LNA: 34.5dB SAAF: 0-30dB ADC: 0, 1.3, 2.9, 4.9dB
[1]Input-referred noise	3.44μV_{rms} (0.5-156Hz)
Maximum input range (1% THD)	8m$V_{p\text{-}p}$
Dynamic range	58.3dB
[2]NEF/PEF (LNA only)	5.32/17
CMRR (1Hz-1kHz)	70.4dB
Input impedance	common-mode: 179MΩ @ 60Hz differential-mode: 1.18GΩ @ 0.5Hz
Electrode offset tolerance	$> \pm$300mV
Maximum PLI tolerance	12.6m$V_{p\text{-}p}$ (input-referred)
ADC linearity	INL: -0.54/0.55LSB DNL: -0.23/0.48LSB
ADC sample rate	10kS/s
ADC ENOB	8.41b @ f_{IN}=100Hz 8.01b @ f_{IN}=4.99kHz
ADC figure-of-merit	37.3 fJ/conversion-step
Area	Single channel (analog): 0.25mm^2 Single channel (digital): 0.20mm^2 Peripherals: 0.18mm^2

Chapter 3

An Invisible Cochlear Implant - Part I: Energy-Efficient Neural Stimulation Waveforms

In the previous chapter, an ultra-low-voltage front-end system for a *wearable* ECG monitor was presented. We now transition to the second focus of this thesis: a fully-*implantable* (i.e., invisible) cochlear implant (CI). Both wearable and implantable devices have many similar functional requirements such as sensing, processing, and communication. However, by virtue of being implanted, implantable devices face stricter volume and energy constraints. Furthermore, many implantable applications require neural stimulation which is a very power-hungry process. Therefore, for implantables, energy-efficient design is not only needed at the architecture and circuit levels, but the sensors and actuators that interface with the human body will also have to be carefully studied.

Our work on demonstrating the feasibility of a fully-implantable cochlear implant (FICI) is broken down into three parts. First, this chapter focuses on developing energy-efficient neural stimulation waveforms aimed at reducing the stimulation power in the CI. Second, in Chapter 4, we present a method to sense sound with a piezoelectric sensor mounted on the middle ear, thereby eliminating the need for an external microphone. Finally, in Chapter 5, we present a system-on-chip that incorporates

the neural stimulation and implantable acoustic sensing contributions, together with reconfigurable sound processing to demonstrate the feasibility of a FICI system.

The focus of this chapter is to reduce the stimulation power (which dominates the overall system power) of neural prostheses like CIs by optimizing the neural stimulation waveform. The first contribution of the chapter is the development of alternate biphasic electrical neural stimulation waveforms that are more energy-efficient than conventional rectangular waveforms using a computational nerve fiber model. The second contribution is the preliminary experimental validation of these simulation results with two sets of *in-vivo* measurements: 1) electrically-evoked compound action potentials (ECAP) from the auditory nerve of two cats, and 2) hearing threshold and loudness perception in two human CI users. Although the target application in this work is for cochlear implants, the results from this chapter are more broadly applicable to other neural prostheses like retinal or vestibular implants, or therapeutic devices such as deep brain stimulators, or any device that delivers electrical stimulation of neural tissue.

This chapter is organized as follows: Section 3.1 provides background on human ear anatomy and gives an overview of a fully-implantable cochlear implant. It also discusses power consumption in a CI and methods to reduce it, and presents an overview of electrical neural stimulation. Section 3.2 presents alternate energy-efficient neural stimulation waveforms from nerve fiber simulations, and Section 3.3 provides experimental validation of the simulation work. A summary is provided in Section 3.4. This work was carried out in collaboration with Rui Jin who ran some of the initial nerve fiber simulations.

3.1 Background

As of December 2010, more than 219,000 people worldwide have received cochlear implants [83], which is more than all other neural prostheses combined. As such, the success of cochlear implants is used as a model for the development of other neural prostheses such as retinal or vestibular implants [13]. Neural prostheses and

implantable therapeutic devices such as deep brain stimulators or vagal nerve stimulators (for treatment of Parkinsons disease, essential tremor, epilepsy) all rely on delivering electrical current into nerve tissue to elicit neural response. The aim of this section is to show why stimulation power in such devices can often dominate the total system power budget, and this is best illustrated by using the cochlear implant as an example.

3.1.1 Anatomy of the Human Ear

This section provides a brief background on the anatomy of the human ear that may be helpful for understanding the discussion on neural stimulation waveforms and implantable acoustic sensing in this chapter and next.

The ear is an intricate organ of the auditory system that performs transduction of acoustical energy to mechanical energy to electrical impulses along the auditory nerve that are interpreted by the brain as sound. The human ear shown in Figure 3-1 is separated into the outer, middle, and inner ear (labeled in blue, green, and orange respectively). The outer ear consists of the pinna (outer ear flap), the ear canal, and the tympanic membrane (ear drum). The umbo is the most depressed part of the ear drum. Sound pressure waves enter the ear canal and vibrate the ear drum back and forth. The motion of the ear drum then couples to three small bones called the malleus, incus, and stapes which comprise the middle ear (also known as the ossicular chain). The ossicular chain provides mechanical filtering and amplification. The motion of the stapes footplate moves the fluid inside the cochlea (inner ear) and causes a wave to travel down the basilar membrane in the cochlea, which is a fluid-filled spiral cavity that is tapered towards the end. A normal hearing person has approximately 30,000 tiny hair cells that line the cochlea, which are excited by the wave in the cochlear fluid. The length of the hair cells are shortest near the oval window (base) of the cochlea and increases toward the round window at the end (apex) of the cochlea. Since the resonant frequency of each hair cell depends on its length, higher frequency waves excite the hair cells near the base, while lower frequency waves excite hair cells at the apical end of the cochlea. This is referred

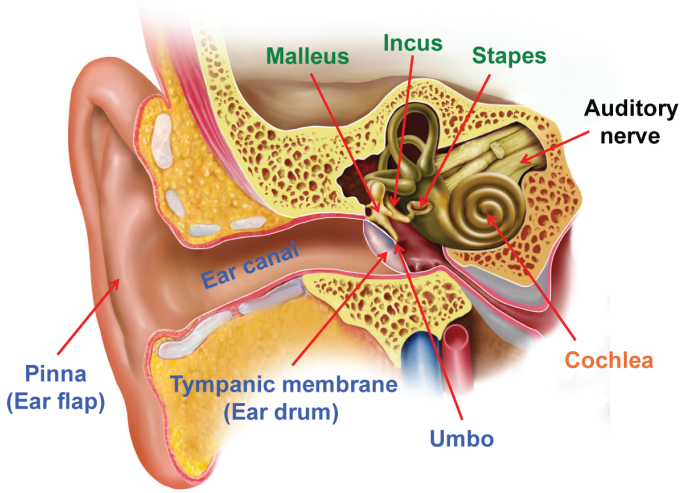

Figure 3-1: Anatomy of the human ear.

to as the tonotopic structure of the cochlea. The motion of the hair cells generate electrical impulses in the auditory nerve, which the brain interprets as sound. Sound pitch is discriminated partly due to to the tonotopic structure of the cochlea, while loudness percepts depend on the amount of hair cells activated in a particular area of the cochlea.

Conductive hearing loss occurs when there is damage to the ossicular chain, while *sensorineural* hearing loss occurs when there is damage to the hair cells in the cochlea. For those who have very few functional hair cells remaining and are profoundly deaf, an implantable cochlear prosthesis can help restore hearing.

3.1.2 A Fully-Implantable Cochlear Implant

The external component of a conventional cochlear implant (CI) that houses the microphone and sound processor raises a number of issues such as usage in the shower or during water sports, as well as concerns with aesthetics and social stigma. These reasons motivate the development of a fully-implantable cochlear implant (FICI) that is totally invisible.

The main challenges with developing a FICI are: 1) developing a wireless power delivery and storage scheme to enable untethered usage for a full day, 2) developing

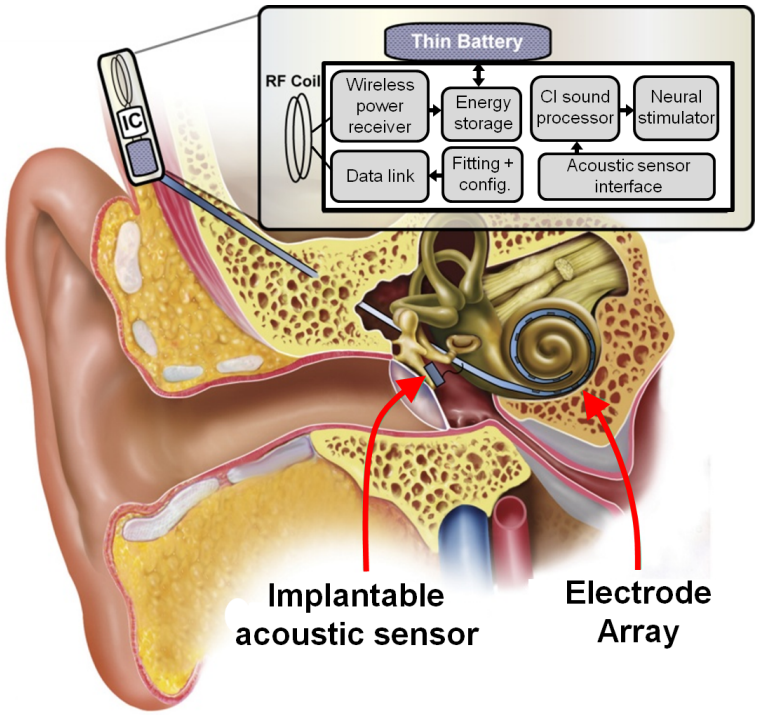

Figure 3-2: Block diagram of a fully-implantable cochlear implant.

an implantable acoustic sensor to detect sound from within the ear, and 3) developing ultra-low-power sound processing and stimulation circuits to ease the burden of energy delivery and storage.

Figure 3-2 shows the conceptual block diagram of a FICI. Compared to a conventional CI, there are a few major differences. First, in this work, an implantable acoustic sensor is mounted at the umbo of the ear drum in order to detect its motion which is a measure of the incoming sound pressure in the ear canal. This replaces the external microphone of a conventional CI. A second difference is the power delivery system. In today's CIs, power must be continuously transferred wirelessly from the external unit to the implanted electronics due to the absence of an implanted battery. In our FICI system, the idea is to rapidly and wirelessly charge up an implanted ultra-capacitor or battery from which the implanted unit can operate for an extended amount of time. This eliminates the need for an external unit, aside from when the FICI recharges while the user sleeps. Lastly, because the FICI relies on implanted energy storage which is limited by the volume, the power consumption of the sound

processing and neural stimulation circuits must be minimized. Therefore, the next section discusses power consumption in a CI which helps identify the power-hungry processes and provides insight on where improvements can be made.

3.1.3 Power Consumption in a CI

Figure 3-3 shows a simplified signal chain in a CI, comprising an acoustic sensor front-end to pick up sound, an analog-to-digital converter, a digital CI sound processor to encode sound, and a neural stimulator to deliver electrical current to the tissue via the electrode array (the wireless link has been omitted for simplicity). The typical power breakdown of the respective components is also shown in Figure 3-3. The acoustic front-end, ADC, and sound processor typically consume a few hundred micro-Watts combined, while the average stimulation power can often be on the order of a few milli-Watts [41, 48, 49]. Therefore, the stimulation power in a CI can represent up to 90% of the total system power budget.

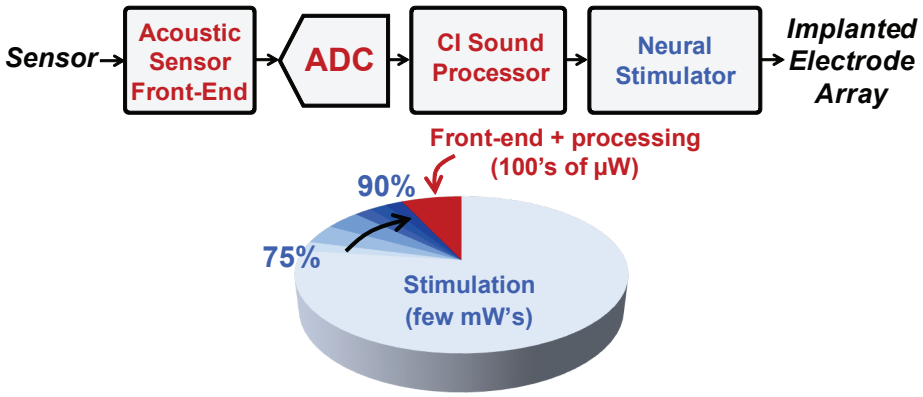

Figure 3-3: Simplified signal chain in a cochlear implant (the wireless link has been omitted), together with its typical power breakdown. The data for the front-end and processing power is based on [41, 84–88]. The stimulation power can represent up to 90% of the total power, dominating the system power budget.

The milli-Watt power consumption in the neural stimulator can be explained by several factors. First, a certain minimum amount of energy must be delivered to reach the threshold for action potential (spike) initiation in a nerve fiber. During this process of delivering charge to the tissue, power is dissipated in the electrical re-

sistance of the electrode-tissue interface. Furthermore, state-of-the-art implants are moving toward more channels (up to 22), higher rates of stimulation, and simultaneous stimulation to achieve "current-steering" for intermediate virtual channels (e.g., in Advanced Bionics HiResolution devices), thereby requiring even more power [13]. Since stimulation power can dominate the total power in a CI, any reduction in stimulation power can translate directly to overall system power reduction, resulting in longer lifetimes or smaller batteries.

3.1.4 Methods for Stimulation Power Reduction

Considering the factors discussed in the previous section, stimulation power in a CI can be reduced at multiple levels of the design. At the physiologic level, the placement of electrodes closer to the target tissue can lower thresholds [89], and optimizing the waveform shape of the stimulus can result in energy savings [90–95]. At the material level, better electrodes with lower impedance can be engineered with new materials. At the algorithm level, alternative sound processing strategies like Asynchronous Interleaved Sampling (AIS) [96] can encode both envelope and phase information and provide higher stimulation rates only for high intensity channels while maintaining a lower average rate of stimulation. As a result, power can be saved while potentially providing higher performance. Finally, at the circuit level, power can be reduced through more efficient stimulation circuits such as adiabatic stimulators [47], or DC-DC converter-based stimulators that use inductors to recycle energy [97].

3.1.5 Neural Stimulation Waveform Efficiency

In this work, our goal is to minimize the energy of the stimulus used in electrical neural stimulation by optimizing its waveform shape. Most neural stimulators today employ charged-balanced biphasic rectangular current pulses shown in Figure 3-4. The first phase is the excitory (cathodic) phase which is responsible for depolarizing a nerve fiber, while the second phase is the recovery (anodic) phase used to ensure that no net charge is delivered to the tissue.

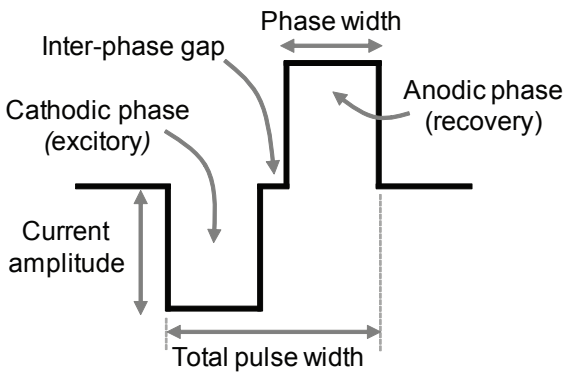

Figure 3-4: Typical charge-balanced biphasic rectangular current waveform.

The rectangular current waveform has been widely adopted in the past because it is simple and fully specified by just the amplitude and duration [91], and it is easily generated using a simple current source. Furthermore, studies have pointed out that the rectangular waveform is the optimal shape for minimizing the energy consumed under the constraint of delivering *equal charge* [98][1]. In the literature, this *equal-charge* constraint has routinely been used to evaluate the efficacy of stimulation waveforms under the assumption that equal charge results in equivalent neural excitation. However, this is a poor assumption because the threshold for neural activation is also a function of the waveform shape [99], and therefore it is possible to achieve neural activation with different amounts of charge injection depending on the waveform shape. As such, the first goal of this work is to investigate alternate waveforms using *neural activation* (instead of charge) as the goal function (i.e., achieving the same neural excitation while minimizing the amount of energy consumed).

The second goal of this work is to provide experimental validation of the efficacy of alternate waveforms with *in-vivo* measurements in the auditory system. This is important because there is a general lack of consensus on the optimal waveform shape, and a lack of experimental data from the auditory system in previous studies [91, 92, 100].

[1]This can be explained intuitively with a simple example. Delivering 1A of current to 1Ω for 1 second requires $E = I^2RT = 1$ Joule. The same total charge can be injected by delivering 2A for 0.5 seconds which would require 2 Joules (2× more energy) because 4× more power is being dissipated over half the time.

3.1.6 Efficiency of Stimulator Circuits

Electrical stimulation of neural tissue can be achieved with current- or voltage-controlled stimulators [101]. Voltage-controlled stimulators are generally more efficient, but ensuring charge-balanced operation is difficult because the current delivered depends on the time-varying electrode impedance. Charge balance in neural stimulators is required to limit the amount of charge build-up at the electrode interface to avoid electrolysis of extracellular fluid that can lead to tissue damage [102,103]. Therefore, current-controlled stimulators are generally preferred because of their safety (known methods of charge balancing) and ease of implementation. However, current-controlled (i.e., current source-based) stimulators can be very inefficient due to the headroom required for the current source transistors. For example, assuming that a stimulation current of $I = 200\mu$A is delivered to an electrode with a resistance of R_{elec} = 5kΩ, the power delivered to the electrode is $P_{elec} = I^2 R_{elec} = 0.2$mW. However, if we assume that the stimulator supply voltage is $V_{DD,stim} = 6$V (in order to provide enough voltage compliance at the maximum stimulation current of 1mA with margin), then the power consumed in the stimulation circuitry is actually $P_{stim} = V_{DD,stim} I$ = 1.2mW. This means that the efficiency of delivering power to the electrode in this example is only $\eta_{stim} = P_{elec}/P_{stim} = 16.7\%$, and most of the power is wasted in the headroom required by the current source.

Recently, an adiabatic stimulator using a bank of capacitors was proposed to increase the efficiency (η_{stim}) of stimulation up to approximately 85% [47]. However, the current delivered was neither constant nor well-controlled. In [97], a DC-DC converter-based stimulator using an inductor to recover energy was designed. It uses feedback to provide the accuracy of current-mode control, with the efficiency of voltage-based systems. However, the loop bandwidth is limited by large passives making this topology not suitable for CI applications that require very short pulses (less than 100μs). Analysis of this approach is provided in Appendix B.

In theory, if the efficiency (η) of the DC-DC converter in an inductor-based approach is 100%, then the efficiency of stimulation (η_{stim}) can approach 100%. In this

ideal case, the power overhead of the stimulation circuitry is zero, and all the power drawn from the stimulator supply $V_{DD,stim}$ is delivered to the electrode (i.e., $P_{stim} = P_{elec}$). Therefore, P_{elec} represents the theoretical minimum power consumption of the stimulation circuits.

Although the inductor-based DC-DC converter approach can be much more efficient than a current-controlled stimulator in theory, practical values of the DC-DC converter efficiency, electrode impedance, stimulation current, and stimulation pulse width for CI applications can limit its benefit over a current-controlled stimulator. Detailed analysis of the efficiency of both the DC-DC converter approach and current-controlled stimulator is provided in Appendix B.

Charge and Energy of Stimulation Waveform

For the rest of this chapter, our analysis will focus on the energy delivered to the electrode (and ignore any wasted overhead power in the stimulation circuitry). This is the most relevant measure of energy because in the limit where η_{stim} approaches 100% (with future improvements in stimulator circuit design), the electrode energy is the theoretical minimum amount of energy that the stimulator has to deliver. Therefore, the energy of the stimulation waveform in the analyses of the rest of this chapter is defined as the energy per unit resistance ($\overline{E}_{elec}$) in the electrode over the phase width (PW), and is given by,

$$\overline{E}_{elec} = \frac{E_{elec}}{R_{elec}} = \int_{<PW>} I^2(t)\mathrm{d}t \tag{3.1}$$

where $I(t)$ is the current of any arbitrary stimulation waveform, E_{elec} is the total waveform (or electrode) energy, and R_{elec} is the electrode resistance. If we account for the stimulation efficiency η_{stim}, then the energy (per Ω) required from the stimulator is given by $\overline{E}_{stim} = \overline{E}_{elec}/\eta_{stim}$.

For the specific case of current source-based stimulators, the more relevant quantity is the average current (or equivalently, total charge) delivered. The total charge

delivered by the waveform is simply

$$Q_{tot} = \int_{<PW>} I(t)\mathrm{d}t, \tag{3.2}$$

and the resulting stimulation energy for a current source-based stimulator is $E_{stim,CS} = V_{DD,stim}Q_{tot}$. Although the analysis in the rest of this chapter focuses on minimizing the waveform energy ($\overline{E}_{elec}$), it can be shown mathematically[2] that any arbitrary (non-rectangular) waveform with energy that is equal to or less than the energy of a rectangular waveform, will require less total charge (Q_{tot}) than the rectangular waveform. Therefore, any waveform that is more energy-efficient (than a rectangular waveform) must also be more charge-efficient, which results in energy savings in current source-based stimulators (assuming the same stimulator supply voltage $V_{DD,stim}$).

3.2 Nerve Fiber Simulations

The first step toward finding energy-optimal waveforms is to use a computational model of a mammalian myelinated nerve fiber to determine the effectiveness of various waveforms on generating action potentials. Figure 3-5 shows a diagram of the widely-used Hodgkin-Huxley type discrete cable model of a single nerve fiber [104]. The model used in this work was provided by Dr. Don Eddington from the Massachusetts Eye and Ear Infirmary, and it is described in detail in Section 2.4.2 of [105]. This model is based on well-accepted models described in [106,107]. The membrane current I_m at the nodes of Ranvier are modeled actively with non-linear conductances (G_{Na+} and G_{K+}) that describe the kinetics of the sodium and potassium ion channels [106, 107], while the internodes are modeled as passive leaky insulators.

The membrane voltage is defined as $V_m = V_i - V_e$, where V_i and V_e are the intracellular and extracellular potentials respectively. C_M is the membrane capacitance, G_L is the membrane leakage conductance, and G_a is the axial conductance. Lastly,

[2]This can be proved using a fundamental mathematical inequality known as Hölder's inequality.

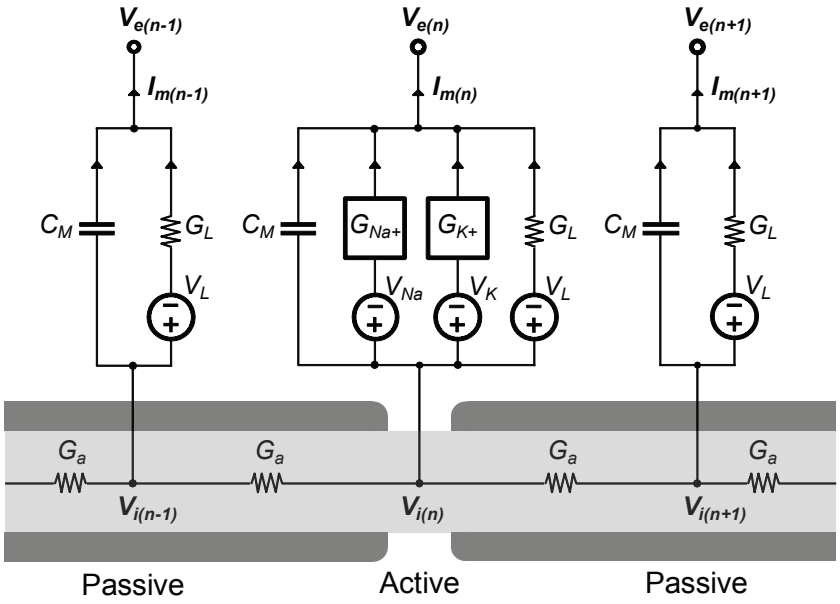

Figure 3-5: Computational model of a single mammalian myelinated nerve fiber [105], where a single node of Ranvier is shown in between two internodes. The membrane current I_m at the nodes of Ranvier are modeled actively with non-linear conductances G_{Na+} and G_{K+} representing the kinetics of the sodium and potassium ion channels respectively, while the internodes are modeled passively as leaky insulators.

V_L is the leak reversal potential, and V_{Na} and V_K are resting potentials.

Finally, it is important to note that these nerve fiber simulations are only useful for determining the energy-efficiency of various waveforms on evoking an action potential in a single nerve fiber which is a binary process (i.e., a particular waveform either triggers a full spike or no spike in the fiber). The influence of waveform shape on auditory perception is addressed later in Section 3.3.2.

3.2.1 Simulation Setup

Figure 3-6 shows the simulation setup used to model the interface comprising the stimulator's n current sources, electrodes, and target neural tissue. Fig. 3-6(a) shows the path of the stimulation current $i(t)$ from the i^{th} current source I_i, through an optional DC blocking capacitor C_{Bi} (usually 100's of nF) to the intracochlear electrode. Fig. 3-6(b) shows a high-frequency model of the intracochlear electrode impedance Z_{Ei}. C_d is the double-layer capacitance at the electrode-solution interface, and R_s is

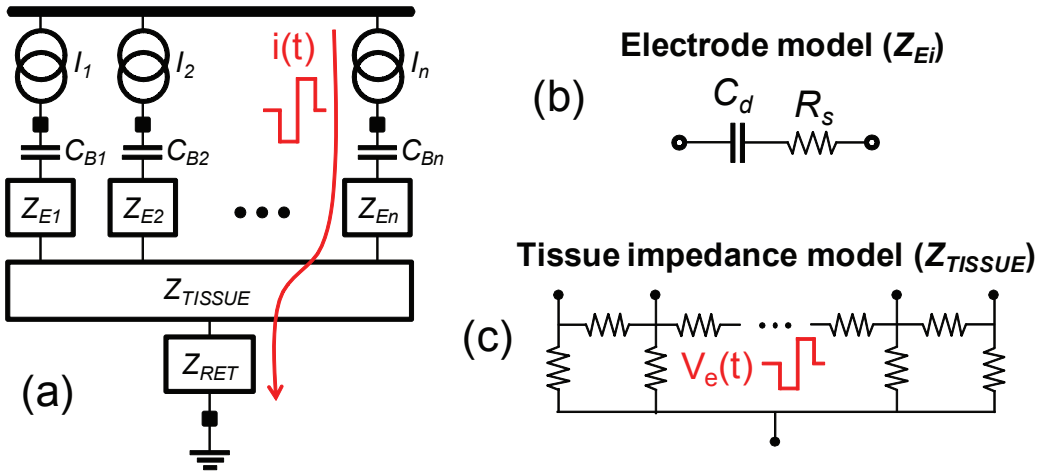

Figure 3-6: Model of the stimulator-electrode-tissue interface used in the nerve fiber simulations.

the solution resistance, and typical values for C_d and R_s in CIs are 5-50nF and 2-10kΩ respectively [108–111]. This high-frequency electrode model is used to calculate the power dissipated in the electrode.

From the intracochlear electrode, the current then spreads through the tissues, cartilage, and bone which can have very different resistivities. There exists very sophisticated 3-D volume conductor models based on histological images that can accurately model the cochlear anatomy [105], but many simpler tissue impedance models use passive networks like the one in Fig. 3-6(c) to model current spread through the tissue [108]. In our simulation setup, for the sake of simplicity, the cochlear tissue is assumed to be homogeneous and isotropic with a resistivity of $\rho_e = 300$ Ω·cm [104]. The voltage generated by the stimulation current $i(t)$ in the tissue at a distance r from the electrode is then calculated as [104]

$$V_e(t) = \frac{\rho_e i(t)}{4\pi r}. \tag{3.3}$$

The potential $V_e(t)$ is then applied to the nerve fiber model as the extracellular potential. Depending on the shape and strength of the stimulus, an action potential may be initiated.

From the tissue, the current is collected by a common extracochlear return elec-

trode which is typically much larger than the intracochlear electrodes. Therefore, its impedance Z_{RET} is typically neglected because it is much smaller than Z_{Ei}.

3.2.2 Energy-Optimal Waveform from a Genetic Algorithm

Because of the complexity and non-linearity of the model shown in Figure 3-5, an analytical solution for the energy-optimal stimulus waveform cannot be determined. However, a heuristic search such as a genetic algorithm (GA) can be used to determine an optimal waveform shape and this was first done by Wongsarnpigoon and Grill in [90]. In general, other global optimization techniques such as simulated annealing can also be used to search for an optimal waveform shape. Inspired by the work in [90], here we focus on the optimization of biphasic stimulation waveforms using stimulation parameters targeted specifically for cochlear implants using the single fiber model described in the previous section.

Genetic algorithms attempt to mimic the process of natural evolution to generate solutions to complex optimization problems. A genetic algorithm works as follows: initially for the first generation of the population, many solutions are generated randomly to expand the space of possible solutions (in order to increase the likelihood of finding a globally optimal solution). A fitness quotient is calculated for each solution of the population (e.g., in this work, the fitness quotient is simply the energy of the waveform, where the lowest energy waveform is the fittest solution), and the fittest solutions are selected as *parent* solutions which are allowed to breed and produce the next generation. To generate a new solution for the next generation, two parent solutions are selected and their characteristics are blended and mutated together. The purpose of blending the parents is to allow their healthy characteristics to be passed on to the next generation, and the purpose of mutation is to allow the possibility for the introduction of a new healthy trait. The fitness quotients are calculated for this new generation, and the process is repeated for a large number of generations until some termination criteria is met.

In this work, the genetic algorithm begins with an initial population of 50 waveforms, each discretized to 20 time steps (i.e., 20 genes). The duty of the waveform is

constrained such that 10 time steps are assigned to each of the cathodic and anodic phases. The choice of 10 steps/phase provides enough degrees of freedom to approximate any arbitrary waveform shape, while being a reasonable level of discretization when considering the limitations of practical stimulation circuits (which have limited bandwidth and settling). In [90], the anodic phase was constrained to be rectangular in shape. In this work, no constraint is placed on the shape of the cathodic or anodic phase to allow both phases to be optimized together. The value of each gene in the initial population is selected randomly between 0 and a full-scale value. For each generation, all waveforms are applied to the nerve fiber model, and the fitness quotient is calculated as the energy of the waveform, plus a severe energy penalty if no action potential is initiated in order to filter out the unfit solutions. At the end of each generation, the fittest waveform is recorded, and the good parent waveforms are allowed to produce offspring for the next generation. At the end of a large number of generations (e.g., 10,000), the output of the simulation is a waveform that is able to trigger an action potential with the lowest energy (up to that point in the simulation).

GA simulation results for a phase width of 25μs

The results of the simulation for a total pulse duration of 50μs (25μs[3] for each of the cathodic and anodic phases) are shown in Figure 3-7, where panels (a) to (e) show the best solution from generations 1, 10, 100, 1,000, and 10,000 respectively, and the evolution of the energy of the GA waveform is shown in Fig. 3-7(f). The energy of a rectangular waveform of equal pulse duration is shown with a dashed line for reference, and it can be seen that the GA waveform is approximately 28% more energy-efficient at the end of 10,000 generations. It is also interesting to note that the optimal GA waveform for a phase width (PW) of 25μs resembles a decreasing exponential cathodic phase, followed by a roughly rectangular anodic recovery phase. This is somewhat different from the truncated Gaussian shapes from [90] and may be explained by the different nerve fiber model used as well as the different constraint on the shape of the anodic phase.

[3]A phase width of 25μs to 50μs (total pulse duration of 50μs to 100μs) is typical in today's CIs.

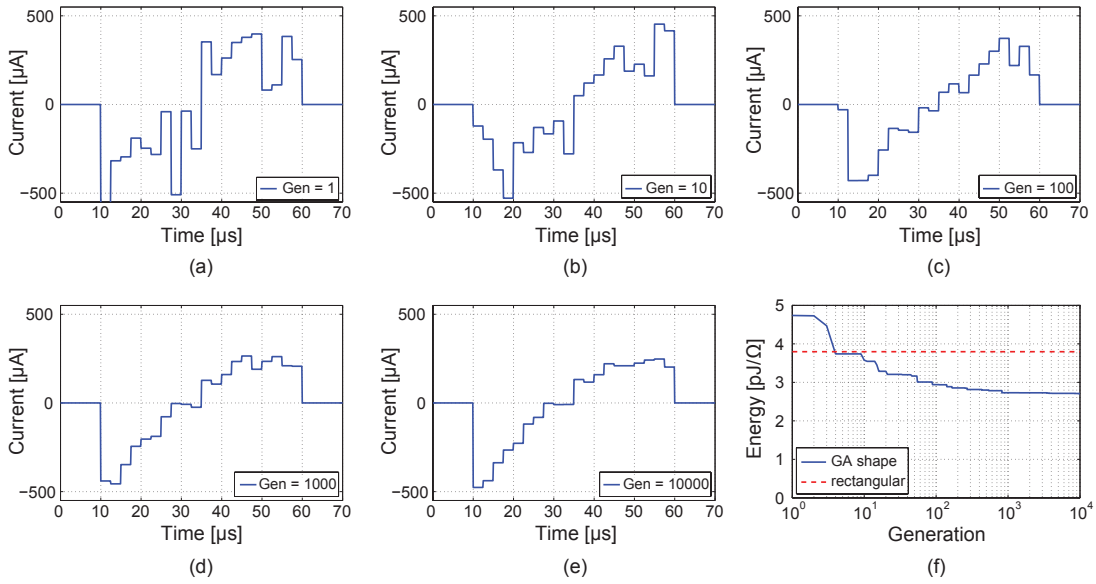

Figure 3-7: Simulation results from the genetic algorithm over 10,000 generations. (a) - (e) show the fittest solution from generations 1, 10, 100, 1,000, and 10,000 respectively. (f) The waveform energy of the best GA shape as it evolves over 10,000 generations. The dashed line shows the energy of the rectangular waveform for reference. The phase width is fixed at 25μs with 10 time steps/phase.

GA simulation results over phase width

Today's CIs may use phase widths ranging from 10μs to 200μs [109, 112]. Therefore, the impact of phase width on the optimal GA shape was also investigated. Figure 3-8 shows the GA shape after 10,000 generations for PWs of 25μs, 50μs, and 100μs, with 10 steps/phase. The corresponding biphasic rectangular waveform at threshold for each of the PWs is overlay for ease of comparison. The GA shape for a PW of 25μs, 50μs, and 100μs requires 2.71pJ/Ω, 1.84pJ/Ω, and 1.49pJ/Ω respectively, which correspond to energy savings of 28%, 35%, and 54% respectively over the corresponding rectangular waveform.

Several observations can be made from the results shown in Figure 3-8. Regardless of the PW, the the cathodic phase is a decreasing waveform with most of the charge being front-loaded in time. Also, the effective duration of the cathodic phase is roughly 30μs to 40μs. For example, for a PW of 25μs (Fig. 3-8(a)), the current is non-zero for the entire cathodic phase. In contrast, for a PW of 50μs or 100μs

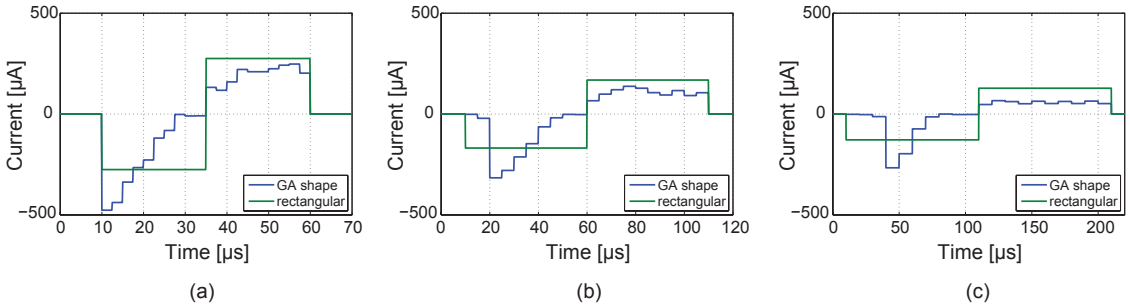

Figure 3-8: GA simulation results after 10,000 generations for a phase width of (a) 25μs, (b) 50μs, and (c) 100μs. A biphasic rectangular waveform is overlay for comparison. Each shape has 10 steps/phase.

(Figs. 3-8(b) and (c)), the genetic algorithm determines that it is not beneficial to deliver charge at the start or the end of the cathodic phase, effectively concentrating most of the current over a 30-40μs window in the middle of the cathodic phase. This explains why the energy savings of the GA shape increases for longer PW, and suggests that there is some critical PW for neural stimulation (beyond which stimulation becomes unnecessarily inefficient). Furthermore, shorter PWs require higher peak currents which is consistent with the strength-duration characteristic for electrical neural stimulation. This could also be attributed to the fact that a longer anodic PW allows the charge in the recovery phase to be more spread out in time, providing less inhibition for the excitory cathodic phase. Finally, the shape of the anodic phase is generally flat for PW=50μs and 100μs, but is somewhat back-loaded in time for PW=25μs.

GA simulation results without a duty cycle constraint

In the last section, the discussion of a critical PW for stimulation suggests that an optimal biphasic waveform may require an uneven split in the duty cycle between the cathodic and anodic phase. In our initial GA simulation, the durations of the cathodic and anodic phases were constrained to be equal (i.e., 50% duty cycle). Therefore, a second GA simulation was performed where that constraint was removed and the durations of the cathodic and anodic phases were allowed to be arbitrary (provided the total pulse duration was fixed). Figure 3-9 compares the results after 10,000

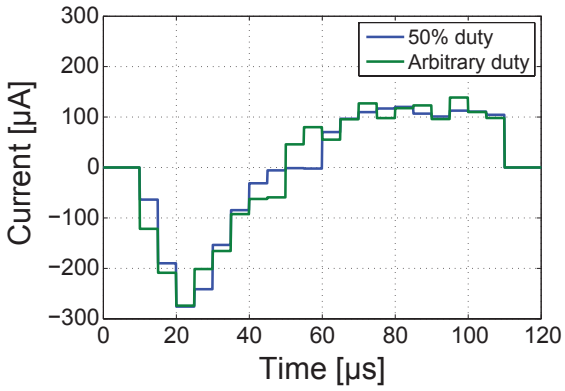

Figure 3-9: GA simulation results after 10,000 generations with the cathodic/anodic duty cycle fixed to 50%, and with no constraint on the duty cycle. Total width of the pulse is 100μs discretized to 20 time steps.

generations from both GA simulations for a total pulse duration of 100μs discretized to 20 time steps. The overall shape from both simulations are similar. However, for the simulation with arbitrary duty cycle, the anodic phase starts slightly earlier and has a 40%/60% split between the cathodic/anodic phases.

GA simulation results over number of genes

The impact of the number of genes (i.e., time steps per phase) on the GA shape was also studied. Figure 3-10 shows the GA shape after 10,000 generations for a PW of 50μs with 5, 10, and 20 steps/phase. In general, a larger number of genes requires more generations for the energy of the GA shape to converge to its steady-state value. However, regardless of the discretization in time, the general shape is still a decreasing front-loaded cathodic phase, followed by a relatively flat anodic phase. The energy savings is roughly independent of the number of genes, and all three cases show an energy savings of approximately 35% when compared to the rectangular shape. In practice, the maximum number of time steps is limited by the stimulator hardware (e.g., bandwidth and settling of the stimulator), and therefore 5 to 10 steps/phase (time step of 5μs to 10μs) appears to be a good tradeoff.

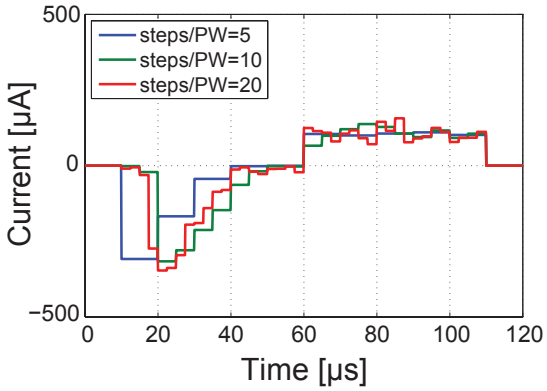

Figure 3-10: GA simulation results after 10,000 generations when varying the number of steps/phase for a phase width of 50μs.

3.2.3 Comparison to Alternative Shapes

The results from the genetic algorithm cannot be guaranteed to be globally optimal, so it is worthwhile to compare it to alternative shapes studied in the literature. A majority of previous studies based on computational models have focused only on monophasic waveforms [91–94] with varying conclusions depending on the nature of the model (e.g., passive linear models versus active non-linear models). Some studies have concluded that rising waveforms like increasing exponentials have reduced thresholds [93–95], while others have refuted those claims by showing that decaying waveforms like decreasing exponentials have shown energy benefits [91, 92]. While monophasic waveforms are conveniently studied using computational models, they are not used in practice because they require a long discharge period following the pulse which limits the pulse rate. In practice, biphasic waveforms are used because charge balancing can be achieved quickly and accurately using the anodic recovery phase.

It is known that the presence of the anodic recovery phase can increase the threshold of the excitory cathodic phase [113, 114], but little work has been done on optimizing the shape of both the cathodic and anodic phases simultaneously. In [95], the authors conclude that rising exponentials for both cathodic and anodic phases is optimal, while in [90] the authors show that the cathodic phase should look like truncated Gaussians with a rectangular anodic phase to be energy-optimal.

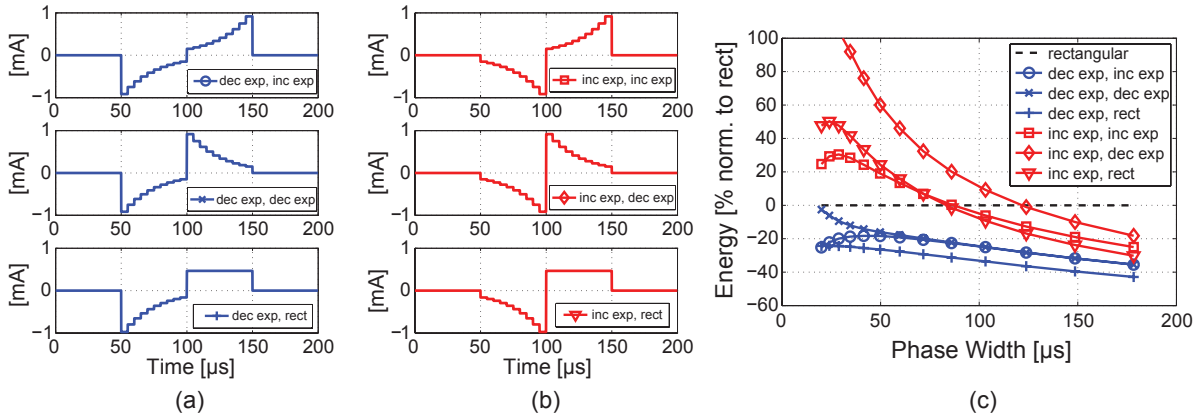

Figure 3-11: Biphasic waveforms with (a) a decreasing exponential cathodic phase, and (b) an increasing exponential cathodic phase. Three different anodic recovery phases are shown for each. (c) Energy of the waveforms in (a) and (b) versus the phase width, expressed as a percentage of the energy of the conventional biphasic rectangular waveform (negative percentages imply energy savings).

Here, the optimal biphasic GA waveforms shown in Figure 3-8 are compared against other common waveforms (discretized in time to 10 steps/phase) found in the literature. Figure 3-11(a) shows alternate biphasic waveforms with a decreasing exponential cathodic phase, and Figure 3-11(b) shows alternate biphasic waveforms with an increasing exponential cathodic phase. For both cases, three different anodic recovery shapes (increasing exponential, decreasing exponential, and rectangular) are shown. Note that the optimal biphasic GA waveform is approximated by the waveform at the bottom of Figure 3-11(a). For all shapes, the waveform energy at the spike threshold is determined for PWs from 20μs to 170μs, and the results are plotted in Figure 3-11(c), where the energy is normalized as a percentage of the energy of a rectangular waveform of equal width (negative percentages imply energy savings with respect to the rectangular waveform).

The results suggest that an increasing exponential cathodic phase is more energy-efficient than a rectangular waveform only for long PWs (PW $>$ 100μs), and becomes very inefficient for short PWs below 100μs. In contrast, decreasing exponential cathodic phases are *always* more energy-efficient than both rectangular and increasing exponential cathodic shapes for all PWs simulated. Among the decreasing exponential cathodic shapes in Figure 3-11(a), there is no difference between a decreasing

exponential and increasing exponential recovery phase for PWs above 50μs, but a decreasing exponential recovery fares worse as the PW falls below 50μs. Finally, it can be seen that the lowest energy waveform among those compared in Figures 3-11(a) and (b) is the decreasing exponential cathodic phase followed by a rectangular recovery anodic phase, which resembles the optimal GA waveform shown in Figure 3-7(e). It achieves an energy savings (with respect to a rectangular waveform) of 25-30% for PWs below 50μs, and the savings increase to approximately 40% at longer PWs.

Impact of Exponential Time Constant

In the previous section, the time constant of the decreasing exponential cathodic phase is fixed at half the phase width. Here, the effect of varying time constant (as a fraction of the phase width) is also considered. Figure 3-12(a) shows waveforms with a decreasing exponential cathodic phase of varying time constants, followed by a rectangular anodic recovery phase of equal charge. Figure 3-12(b) shows the waveform energy normalized as a percentage of the energy of a rectangular waveform. For short PWs below 50μs, a time constant between 0.3$\times$ and 0.5$\times$ the PW appears to be energy-minimal, while for longer phase widths, shorter time constants provide more energy savings.

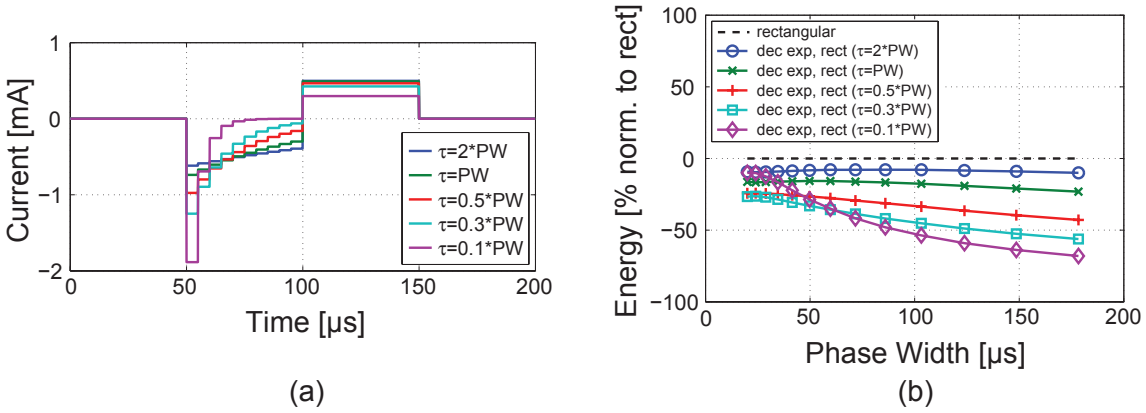

Figure 3-12: (a) Biphasic waveforms with a decreasing exponential cathodic phase of varying time constants, followed by a rectangular anodic phase of equal charge. (b) Energy of the waveforms in (a) versus the phase width, expressed as a percentage of the energy of the conventional biphasic rectangular waveform (negative percentages imply energy savings).

Comparison with Short Rectangular Pulses

From Figure 3-12(a), it is apparent that as the exponential time constant gets shorter, the exponential shape looks more like a short rectangular pulse. Therefore, instead of comparing decreasing exponential waveforms to rectangular waveforms of equal phase width, another perspective is to compare the decreasing exponential waveform to a short rectangular pulse. For example, assuming that the exponential shape of duration t_{PW} is discretized into N time steps, where each time step is a short rectangular pulse of duration $t_{step} = t_{PW}/N$, it is worthwhile to compare the energy of the exponential shape of duration t_{PW} to the energy of a short rectangular pulse of duration t_{step}. Note that in this case, only monophasic stimulation is considered to remove the effect of the anodic recovery phase.

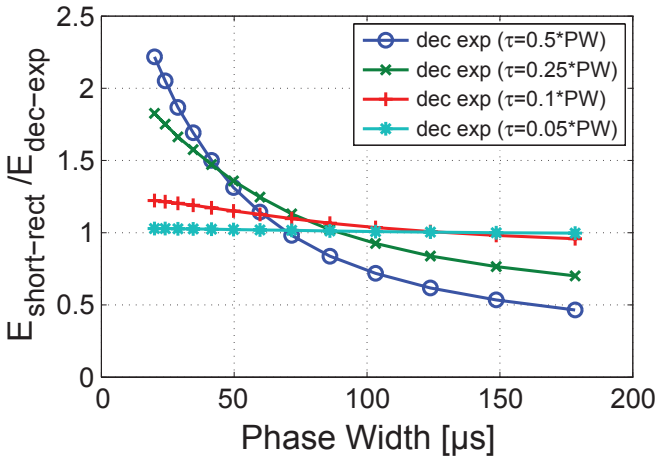

Figure 3-13: Ratio of energy of a short rectangular monophasic pulse ($E_{short-rect}$) with duration equal to 0.2× of the phase width, to the energy of a monophasic decreasing exponential pulse ($E_{dec-exp}$). Four different time constants are plotted.

Figure 3-13 shows the ratio of the energy of a short rectangular pulse of duration $t_{PW}/5$ ($E_{short-rect}$) to the energy of a decreasing exponential waveform of duration t_{PW} ($E_{dec-exp}$), for t_{PW} from 20μs to 170μs. Four different exponential time constants are plotted. As expected, for short time constants like τ=0.05×PW, the ratio is very close to 1. As the time constant increases, the ratio is greater than 1 for short PWs, and less than 1 for longer PWs. In the case of a time constant of 0.5×PW (close to

the value in the optimal GA waveform), the ratio is much greater than 1 for PWs less than 70μs. This implies that the decreasing exponential waveform of duration t_{PW} is more energy-efficient than a short rectangular pulse (of duration $t_{PW}/5$) at short PWs only. This may be explained by the strength-duration characteristic of electrically-evoked neural response which shows that the stimulation duration and spike threshold are inversely proportional. Therefore, as the stimulation duration becomes too small, the threshold for spike generation may increase so much that any benefit from a short PW is negated.

3.3 Experimental Validation with *In-vivo* Measurements

This section presents two sets of *in-vivo* measurement results that complement and validate the nerve fiber simulation results from the previous section. The first set of data consists of ECAP recordings from the auditory nerve of two cats. ECAP data is clinically important because they are commonly used to 1) assist with processor fitting and programming (particularly in children), 2) objectively verify auditory nerve function, and 3) identify redundant or highly interacting channels in CIs [115]. The second set of data comprises the measurement of hearing threshold and loudness perception from two human CI users with Advanced Bionics implants.

3.3.1 ECAP Recordings in Cats

With the help of the Eaton-Peabody Laboratory at the Massachusetts Eye and Ear Infirmary (MEEI), two cats were anesthetized and implanted with cochlear electrode arrays. Stimulation was provided using a wide bipolar configuration, and electrically-evoked compound action potentials (ECAP) in the auditory nerve were recorded. All measurements in this section were made by Dr. Kenneth Hancock at MEEI. Three biphasic waveforms were tested: (1) rectangular, (2) decreasing exponential cathodic, increasing exponential anodic, and (3) increasing exponential cathodic, increasing

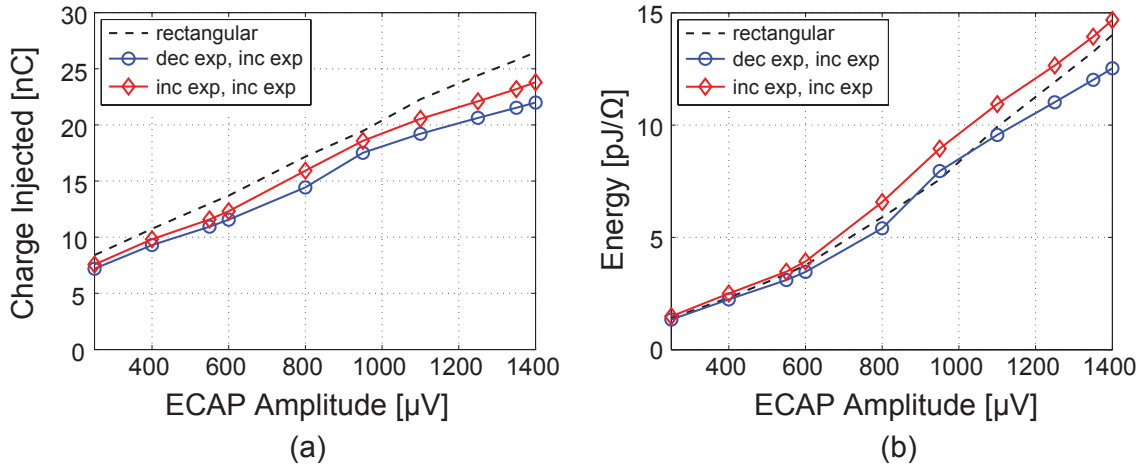

Figure 3-14: (a) Charge injected and (b) energy delivered per phase versus the measured ECAP amplitude for three different biphasic waveforms from one cat subject.

exponential anodic. Figures 3-14(a) and (b) show the charge injected and energy delivered per phase versus the measured ECAP amplitude for the three waveforms for one cat subject. Both alternative waveforms require less charge injected to achieve the same ECAP response. However, in terms of energy, only the decreasing exponential cathodic, increasing exponential anodic shape was more energy efficient than the rectangular shape.

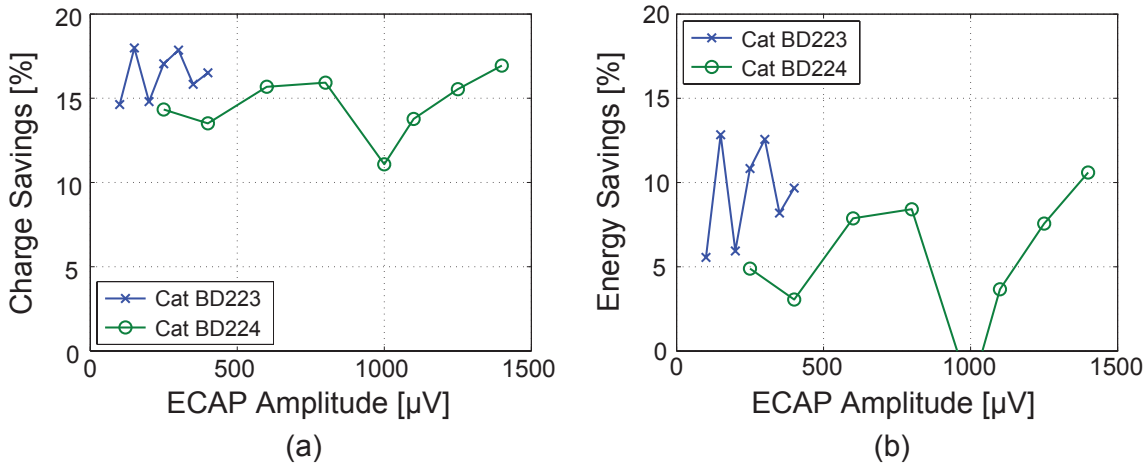

Figure 3-15: (a) Charge savings and (b) energy savings of the biphasic waveform (decreasing exponential cathodic, increasing exponential anodic) for two cat subjects. Data is plotted as percentage savings with respect to the conventional biphasic rectangular waveform.

Figures 3-15(a) and (b) show the charge savings and energy savings of the decreasing exponential cathodic waveform with respect to the rectangular waveform for both cat subjects. For an equal ECAP response, a charge savings of approximately 15% is observed, while the energy savings is approximately 5-10%. Note that the optimal biphasic GA waveform (decreasing exponential cathodic phase, rectangular anodic phase) was not tested with the cat subjects because of limitations with the test setup.

3.3.2 Threshold and Loudness in Human Subjects

While ECAP measurement data from cats help support the nerve fiber simulation results, the data is ultimately limited in terms of what conclusions can be drawn about the loudness perception of a human cochlear implant user. To address this issue, two volunteers with Advanced Bionics devices were recruited, and two waveforms were tested in both subjects by using the test setup at the Cochlear Implant Research Laboratory led by Dr. Don Eddington at MEEI. All measurements in this section were performed by Victor Noel under the MEEI IRB protocol #94-01-003 approved by the MEEI Human Studies Committee and the MIT Committee On the Use of Humans as Experimental Subjects (COUHES). An exponential biphasic waveform (decreasing exponential cathodic, increasing exponential anodic) was tested against the biphasic rectangular waveform using electrode #9 and electrode #7 in subjects 1 and 2 respectively. Both electrodes were in the middle of the electrode array. The first test conducted was a 3-alternative forced choice (3AFC) test with feedback to determine the threshold of hearing for both waveforms, and four trials for each waveform was performed in pseudo-random order.

For the first subject, the average of four trials showed that the required charge injected at threshold for the rectangular and exponential waveforms were 11.4nC/phase and 9.0nC/phase respectively, while the average energy delivered per phase were 2.42pJ/Ω and 2.06pJ/Ω respectively. This corresponds to a charge savings of 21% and energy savings of 15% for the exponential waveform when compared to the rectangular waveform at the threshold of hearing. Similar results were observed with the

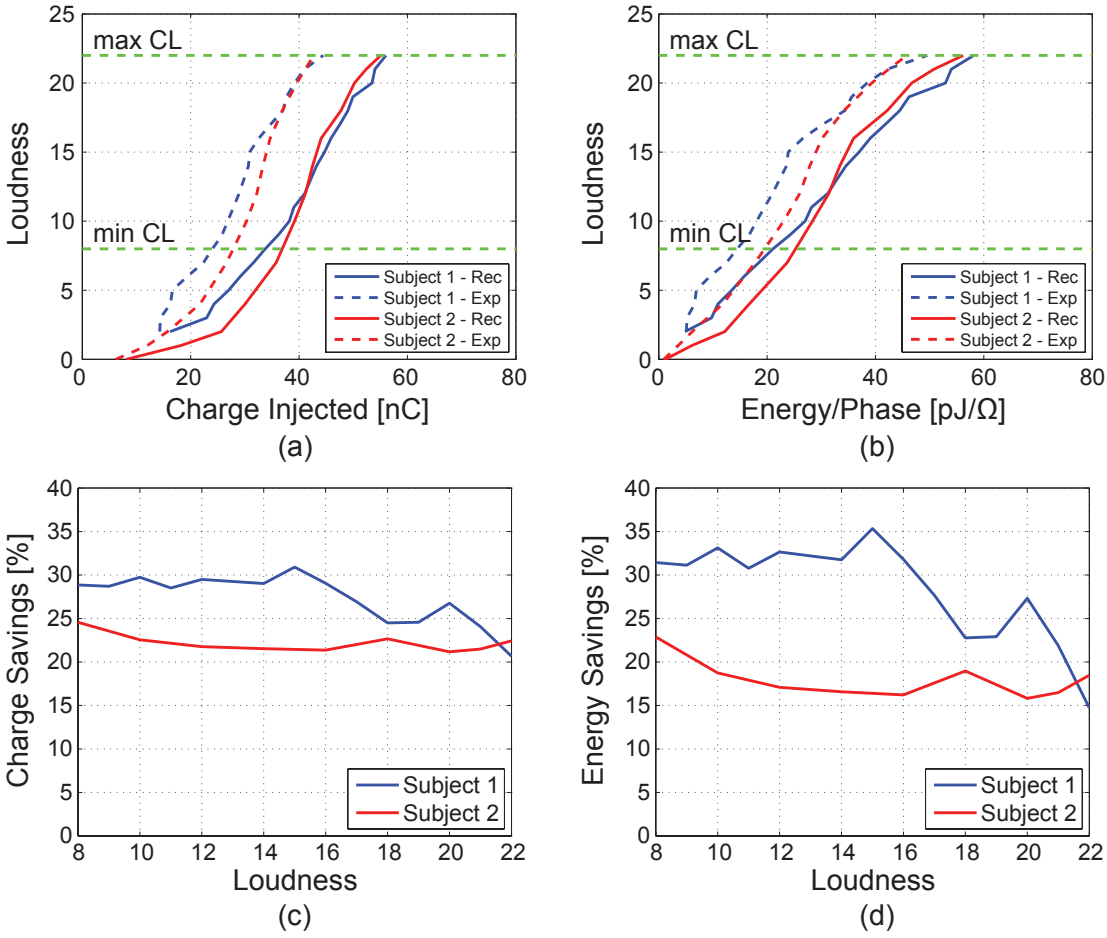

Figure 3-16: Perceived loudness from two human subjects versus the (a) charge injected and (b) energy delivered for a rectangular waveform and exponential waveform (decreasing exponential cathodic, increasing exponential anodic). Average (c) charge savings and (d) energy savings of the biphasic exponential waveform versus the perceived loudness. Data plotted for subject 1 is an average of 2 trials and data plotted for subject 2 is an average of 4 trials.

second subject, where the exponential waveform showed a charge and energy savings of 19% and 11% respectively at the threshold of hearing.

Although threshold data is important for determining the minimum amount of stimulus required, it is actually more important to consider the charge and energy savings over the entire dynamic range of electrical hearing (i.e., from threshold to the most comfortable level). This is because the the stimulus levels provided by the CI during regular usage (e.g., during regular conversation) are well above the threshold of hearing. It is also more relevant from a system perspective to investigate whether savings can be achieved at higher stimulus levels because savings at high power levels are more beneficial than savings at the hearing threshold.

Therefore, the second test conducted on both subjects was a subjective psychophysical loudness perception test using the same two biphasic waveforms. The stimulation amplitude was swept from threshold to just beyond the subject's maximum comfortable level in 50μA steps, and the subjects were asked to rate the loudness on a scale of 0 to 25, with 8 being the minimum comfortable level, and 22 being the maximum comfortable level. The perceived loudness is plotted against the charge injected and energy delivered in Figures 3-16(a) and (b) respectively. The data plotted for subjects 1 and 2 are from an average of 2 and 4 trials respectively. The charge and energy savings of the exponential waveform over the comfortable hearing range are shown in Figures 3-16(c) and (d) respectively. Within the comfortable loudness region (between 8 and 22), charge savings of 20-30% and energy savings of 15-35% across the two subjects were measured.

3.4 Chapter Summary

The power budget of a cochlear implant can be dominated by the stimulation power, and any reduction in the stimulation power transfers directly to overall power savings in the cochlear implant. Therefore, the focus of this work is to investigate alternative non-rectangular stimulation waveforms requiring less charge and energy while providing the same neural response or loudness perception. Most prior work focused on

optimizing the waveform shape of monophasic stimulus. Here, we focus on biphasic waveforms by optimizing the excitory cathodic phase and the recovery anodic phase together without constraining their shape or duty cycle. A computational model of a nerve fiber coupled to a genetic algorithm is used to determine an energy-optimal biphasic waveform that is 28% more energy-efficient than the conventional rectangular waveform at a phase width of 25μs. Independent of the phase width, the energy-optimal GA waveform has a cathodic shape that is decreasing and front-loaded, while the anodic shape is typically flat. Furthermore, the energy-optimal GA waveform appears to have a critical cathodic duration of approximately 30μs to 40μs.

In-vivo measurement results are also presented, where a biphasic exponential waveform (decreasing exponential cathodic, increasing exponential anodic) is compared to a biphasic rectangular waveform (the GA shape could not be tested in cat or human because of limitations with the test setup). ECAP measurements in two cats show a 15% charge reduction and 5-10% energy reduction with the exponential waveform for the same ECAP amplitude response. Data is also presented from two human CI users. The exponential waveform requires 19-21% less charge and 11-15% less energy to reach the threshold of hearing, and it also achieves 20-30% charge savings and 15-35% energy savings within the comfortable hearing range, which is perhaps more relevant than the savings at threshold since that is the region where a CI dissipates the majority of power. Although the experimental effort at this time is still in its infancy and the sample size is quite limited, these initial biological measurements provide encouraging validation of the nerve fiber simulation results. Furthermore, since the threshold and loudness tests in human CI users are subjective, many more subjects will have to be tested in order to gather statistics and provide better confidence of the results.

Chapter 4

An Invisible Cochlear Implant - Part II: Implantable Acoustic Sensing

The previous chapter discussed the investigation of alternate non-rectangular neural stimulation waveforms aimed at reducing the stimulation power, and consequently the power of the entire implant. In this chapter, the focus shifts from the stimulation back-end of the fully-implantable cochlear implant (FICI) system, to the implantable acoustic sensing at the front end. A key enabler of a FICI system is an implantable sensor that is able to sense external acoustic information from within the body.

Recently, totally invisible middle ear implants (MEIs) have been developed to treat conductive hearing loss [116–118]. Most MEIs use an implantable middle ear sensor or a tiny subcutaneous microphone to replace an external microphone. The output transducer (e.g., piezoelectric or floating-mass transducers [118]) of the MEI is usually coupled to the head of the stapes and it provides increased vibration to compensate for hearing loss [116].

Instead of using the sensor readout as an input to the output transducer of a MEI, the readout can be used as an input to the CI sound processor of the FICI system. Inspired by the MEI sensors [116,117], this chapter presents a method for implantable acoustic sensing using a piezoelectric sensor mounted at the umbo of the middle ear.

The main contributions of this chapter are the design of a discrete prototype of a piezoelectric sensor front-end, and the characterization of the sensor on two human cadaveric temporal bones.

This chapter is organized as follows: Section 4.1 provides background on implantable acoustic sensing and presents the circuit model parameters for piezoelectric sensors. Section 4.2 describes the design of a discrete prototype of a piezoelectric sensor front-end, and the test setup for the characterization of the sensor mounted on human cadaveric temporal bones. Characterization results are provided in Section 4.3, and a chapter summary is given in Section 4.4.

4.1 Background

Mild cases of conductive hearing loss can be treated with hearing aids which essentially amplify the incoming sound and produce a louder acoustic output. Middle ear implants (MEIs) on the other hand, may use the natural ear drum (with a sensor attached) as a natural microphone, and an output transducer to provide mechanical vibration coupled onto the stapes which drives the inner ear [118]. Often times, the ossicular chain has to be disarticulated to prevent the output transducer from feeding back to the sensor [116]. In effect, MEIs provide the required impedance matching between external sound waves and the internal ear. MEIs, as their name suggests, are strictly middle ear devices and require the inner ear function to be intact.

To treat individuals with severe sensorineural hearing loss, hearing aids or MEIs are not appropriate, and a cochlear implant is one of the few solutions. In this section, we review recent work on implantable acoustic sensors for MEIs, and show that an umbo-mounted piezoelectric sensor is suitable for a FICI system as well.

4.1.1 Recent Work on Implantable Acoustic Sensing

There have been various approaches to developing an implantable acoustic sensor, and both academic and industry references are summarized in Table 4.1. The most popular approach is to sense the vibration of the umbo or the ossicles induced by sound

Table 4.1: Summary of existing approaches for implantable acoustic sensing.

Reference	Sensor	Description
Academic References		
Park, Bio. Microdev., 2007 [120] Young, TBME, 2012 [119]	MEMS accelerometer	Detect incus or umbo motion
Huang, EMBC, 2007 [121]	MEMS displacement sensor	Detect umbo motion
Maniglia, AJO, 1999 [123]	Magnetic, coupled to electromagnetic coil	Mounted on malleus
Vujanic, ICM, 2002 [124]	Optical vibrometry	Contact-less
Industry References		
Jenkins, OHNS, 2007 [125] (Otologics Carina)	Subcutaneous microphone	Sense sound from under the skin
Perkins, HR, 2010 [126] (EarLens Corp.)	Small microphone in ear canal	Microphone is invisible from outside
Kroll, TIA, 2002 [117] (Envoy Medical Esteem)	Piezoelectric	Detect ossicle motion

pressure waves entering the ear, and interpreting the vibration as sound. MEMS accelerometers [119, 120] and MEMS displacement sensors [121] have been used as vibration sensors. However, these sensors are often too large when compared to the dimensions of the malleus or incus, and mass-loading of the bones can severely dampen the natural frequency response of the middle ear [122]. Furthermore, these types of sensors currently lack the sensitivity and dynamic range required for human hearing [44]. An alternative approach is to sense vibration using a magnetic sensor (mounted on the malleus) coupled to an electromagnetic coil [123]. However this approach is MRI incompatible. To address the issue of mass-loading, contact-less sensing using optical vibrometry has been explored with some success [124]. However, this approach requires high power consumption and has limited reliability due to temporary signal loss.

In the industry, several companies have made significant strides on implantable acoustic sensors using alternative approaches. The Carina device from Otologics uses a subcutaneous microphone to sense sound from under the skin. However, it requires precise placement on soft tissue to minimize pick-up of unwanted body noise arising from actions like talking or chewing [125]. EarLens corporation offers a system that uses a small microphone placed in the ear canal. Although not implantable,

the microphone is invisible from the outside. This system benefits from sensing a wider bandwidth and utilizing the natural filtering of the outer ear to improve sound localization and speech recognition in noisy environments [126]. Finally, the Esteem device from Envoy Medical is a completely implantable middle ear device that uses piezoelectric materials to sense vibration on the malleus, and provide actuation to the stapes [117]. As of 2010, the Otologics Carina and Envoy Esteem are two devices available for clinical use [118].

4.1.2 Feasibility Study with a MEMS Accelerometer

A feasibility study on using MEMS accelerometers as a vibration sensor was performed by Meena Siddiqui at the Massachusetts Eye and Ear Infirmary in October 2010. The goals of that study were to identify 1) the required noise floor of the sensor, and 2) the sensor mass constraints. The umbo was selected as the location for the sensor because previous studies have demonstrated that the umbo experiences the largest amplitude of vibrations in the middle ear [127–129].

In order to determine the sensitivity requirements, pure tones at frequencies from 0.1 to 19kHz were delivered to the middle ear. Laser Doppler vibrometry (LDV) was used to measure the velocity of the umbo vibrations. The sound-to-acceleration transfer curve (normalized to sound pressure level) was derived from the velocity transfer curve. It was determined that for a minimum detectable sound of 40dB SPL, the required integrated noise is $10\mu g_{\text{rms}}$ ($1g = 9.8\text{m}/s^2$) over an 8kHz bandwidth, necessitating a noise floor of approximately $0.1\mu g_{\text{rms}}/\sqrt{\text{Hz}}$. As a point of comparison, a commercially available accelerometer such as the Bosch BMA140 accelerometer has a noise floor of $220\mu g_{\text{rms}}/\sqrt{\text{Hz}}$, a difference of 3 orders of magnitude. Therefore, commercially available accelerometers lack the sensitivity required for an implantable acoustic sensor. Finally, in order to determine the mass-loading limitations, the experiment was repeated with copper disks attached to the umbo ranging in mass from 8mg to 34mg and it was concluded that the sensor should weigh less than 20mg.

As a result of this feasibility study, a decision was made to pursue using piezoelectric materials as the sensing element because of its small size and mass, customizability

(they can be cut to any shape and size), low-power operation, and superior sensitivity required for detecting sound pressures less than 60dB SPL. Furthermore, Envoy's piezoelectric-based Esteem device provides further evidence that piezoelectric sensors are a promising direction.

4.1.3 Piezoelectric Sensor Circuit Model

The piezoelectric material used in this work (PSI-5A4E from Piezo Systems Inc.) is composed of Lead-Zirconate-Titanate (PZT). The sensor is clamped down at one end like a cantilever, while the other end is placed at the umbo of the ear drum. As the umbo vibrates up and down, it exerts a force F on the sensor as depicted in Figure 4-1(a). As the force bends the piezoelectric material, an open circuit voltage is generated across the terminals of the sensor according to [130]

$$V_{OC} = g_{31}\left(\frac{3L}{2Wt}\right)F = g_{31}\left(\frac{3L}{2Wt}\right)mA_U(f)P_{EC}, \tag{4.1}$$

where W, L, and t are the dimensions of the sensor shown in Figure 4-1(a) and g_{31} is the piezoelectric transverse voltage coefficient equal to -11.6$\times$10^{-3} V$\cdot$m/N for PZT [130]. The force from the umbo can be estimated from the ear canal pressure (P_{EC}) according to $F = mA_U(f)P_{EC}$, where m is the mass of the sensor and $A_U = a_{UMBO}/P_{EC}$ is the umbo acceleration normalized to P_{EC}. A typical value of A_U at 1kHz is approximately 1 to 2 (m/s^2)/Pa [119].

For a 3mm $\times$ 3mm sensor that is 20 mils in thickness (PZT has a density of 7800 kg/m^3 which can be used to calculate its mass), V_{OC} ranges from 2.4μV$_{\text{rms}}$ to 0.8mV$_{\text{rms}}$ for sound pressure levels from 40 to 90dB SPL. This expected range of V_{OC} can be used as a guideline when specifying the noise and dynamic range requirements of the sensor front-end.

The charge circuit model of the sensor is shown in Figure 4-1(b), while the voltage circuit model is shown in Figure 4-1(c). These two models are equivalent as long as

$$V_P = \frac{q_P}{C_P}\frac{s}{s+\omega_P}, \tag{4.2}$$

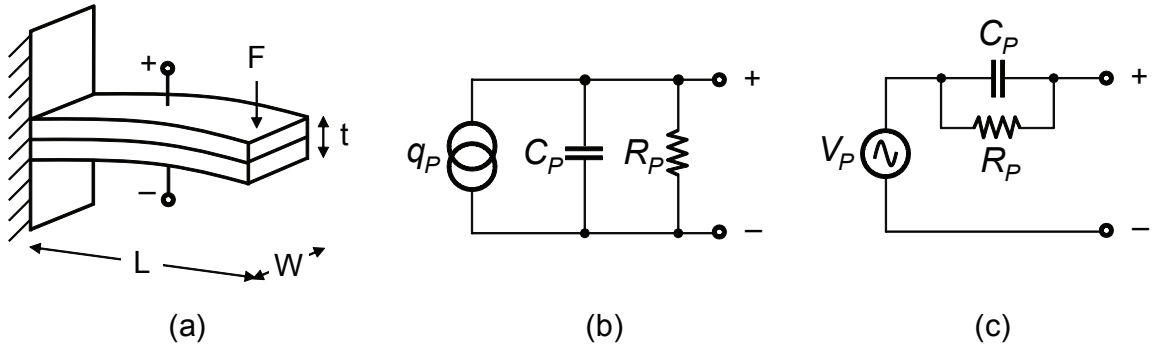

Figure 4-1: (a) Diagram of the piezoelectric sensor configured as a cantilever. The equivalent charge and voltage circuit models are shown in (b) and (c) respectively.

where $\omega_P = \frac{1}{R_P C_P}$, C_P is the capacitance of the sensor, and R_P is the leakage through the insulator between the electrodes. Given that $C_P = \epsilon_r \epsilon_o A/t$ and $R_P = \rho t/A$ where $A = WL$ is the surface area of the sensor, then

$$\omega_P = \frac{1}{R_P C_P} = \frac{1}{\rho \epsilon_r \epsilon_o}, \tag{4.3}$$

where ρ is the resistivity of the insulator, ϵ_r is the relative permittivity of PZT, and ϵ_o is the permittivity of free space. Note that ω_P is determined by the properties of the material alone and is equal to 1 mHz for PZT ($\rho = 10^{10}$ Ω·m and $\epsilon_r = 1800$ [130]). Therefore, for all frequencies of interest, Equation 4.2 simplifies to $V_P = q_P/C_P$ and R_P can be ignored in the circuit models shown in Figures 4-1(b) and (c). For reasonable sizes of the sensor, typical values of C_P range from 0.2nF to 3nF. Lastly, as the circuit model suggests, piezoelectric sensors are high-impedance and therefore can be very low power. However, for the same reason, they can be susceptible to electromagnetic interference [120]. This requires that the wires to the sensor be shielded, or that the front-end be placed as close to the sensor as possible.

4.2 Piezoelectric Sensor Front-End Prototype

In order to investigate the performance of the middle-ear mounted piezoelectric sensor, a discrete prototype of the sensor front-end was designed, and human cadaveric temporal bones were provided from the Eaton-Peabody Laboratory at MEEI. A block

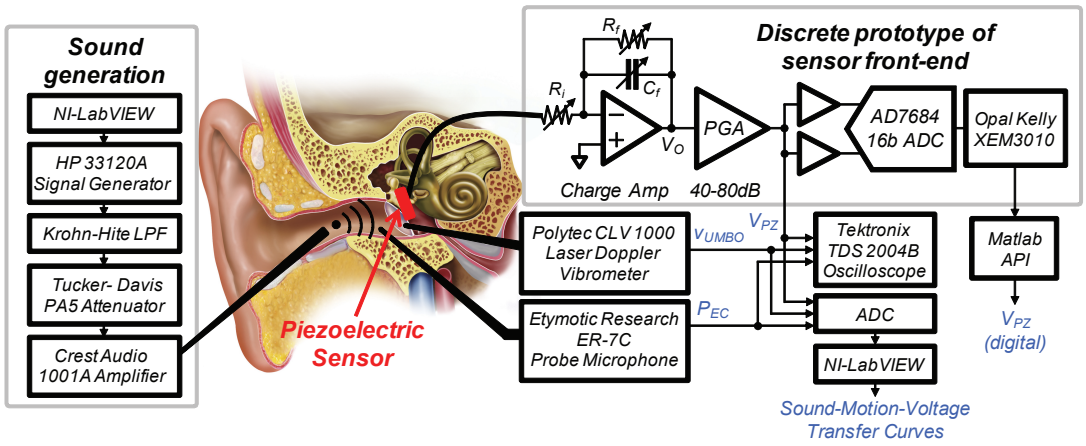

Figure 4-2: Block diagram of the measurement setup and discrete prototype of the sensor front-end for the piezoelectric sensor mounted at the umbo of a human temporal bone. The ear canal pressure (P_{EC}), umbo velocity (v_{UMBO}), and piezoelectric sensor output voltage (V_{PZ}) are measured. The measurements were made possible with the help of Dr. Heidi Nakajima at the Massachusetts Eye and Ear Infirmary.

diagram of the prototype and measurement setup is shown in Figure 4-2.

The entire test setup is controlled by LabVIEW (National Instruments). Input tones ranging from 0.1kHz to 19kHz are generated with a HP33120A signal generator, filtered and amplified by a Crest audio amplifier which drives a speaker connected to a probe tube that funnels the sound into the ear canal of the temporal bone. Ear canal pressure (P_{EC}) is measured by an ER-7C probe microphone (Etymotic Research), and the umbo velocity (v_{UMBO}) is measured by a CLV 1000 laser Doppler vibrometer (Polytec). The discrete prototype is used to amplify, filter, and record the output of the sensor. One terminal of the piezoelectric sensor is biased at the reference voltage (analog ground) from the prototype, while the other terminal is connected to the input of a charge amplifier. Using the charge model of the piezoelectric sensor and ignoring R_P, the band-pass transfer function of the charge amplifier is

$$H_{CA}(s) = \frac{V_O}{q_P}(s) = \frac{-s(\omega_i/C_f)}{(s+\omega_f)(s+\omega_i)}, \tag{4.4}$$

where $\omega_f = \frac{1}{R_f C_f}$ and $\omega_i = \frac{1}{R_i C_P}$ are the high-pass and low-pass corner frequencies respectively, and the mid-band gain is $A_{MB} \approx -1/C_f$. R_i, R_f, and C_f can be tuned on the PCB to provide programmable gain and a pass-band from approximately

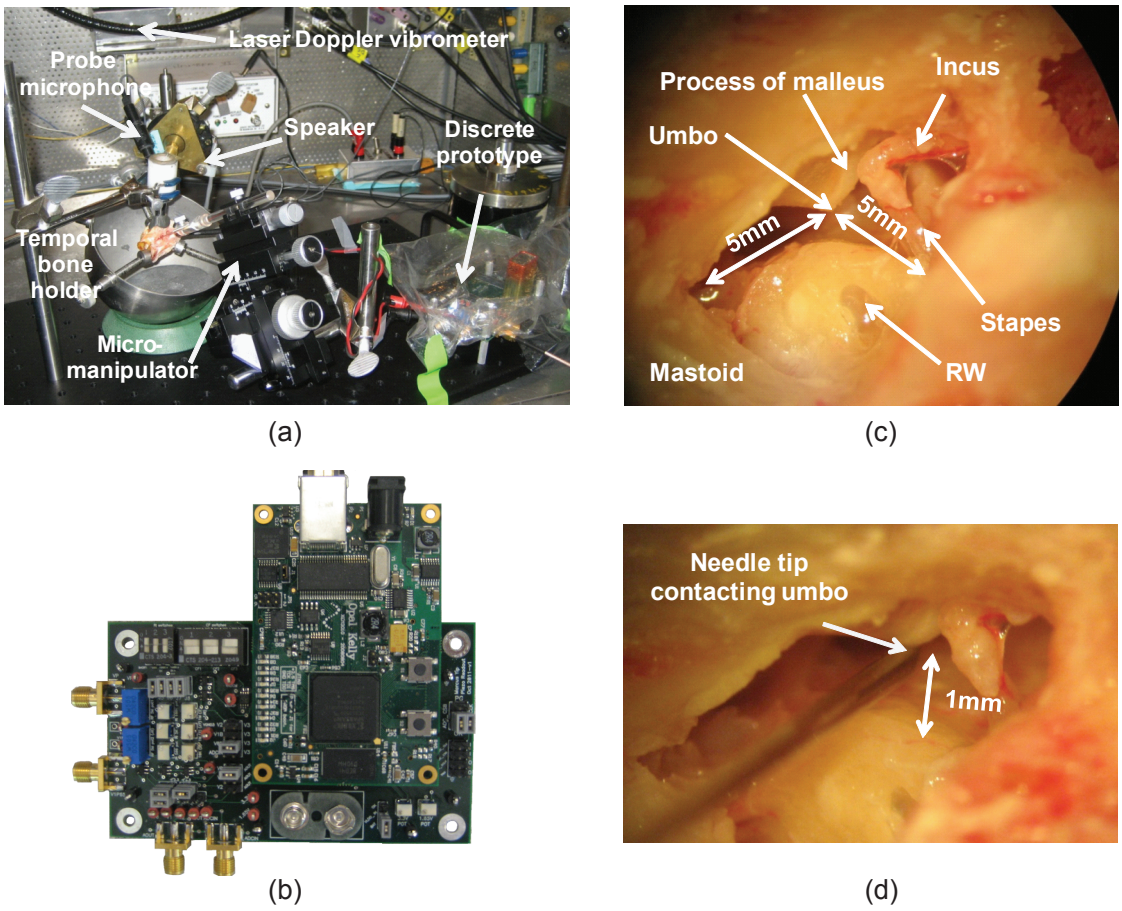

Figure 4-3: Photograph of the (a) actual measurement setup at the Massachusetts Eye and Ear Infirmary, (b) discrete prototype of the piezoelectric sensor front-end, (c) facial recess opening showing the middle ear, and (d) needle tip (which is attached to the piezoelectric sensor) contacting the umbo.

100Hz to 10kHz for a range of sensor sizes (C_P). A second stage programmable-gain amplifier (PGA) provides an additional 40dB to 80dB of gain to generate V_{PZ}. In this system, there is no automatic gain control, and the gain is set (based on the size of the sensor) via a one-time calibration at power-up. All three outputs (P_{EC}, v_{UMBO}, and V_{PZ}) are recorded by LabVIEW and the transfer characteristics from the ear canal pressure (P_{EC}), to the umbo velocity (v_{UMBO}), to the sensor output voltage (V_{PZ}) can be calculated. Furthermore, the output voltage from the prototype is also digitized by a 16-bit ADC (AD7684) and recorded by Matlab on a laptop computer through the USB interface of the Opal Kelly XEM3010 development board for post-processing. The charge amplifier, PGA, and ADC drivers all use the AD8603 precision op-amp from Analog Devices.

Figures 4-3(a) and (b) show photographs of the measurement setup at MEEI and the discrete prototype. The temporal bone is held in place by a temporal bone holder, and the piezoelectric sensor is positioned by a micro-manipulator external to the temporal bone. Fig. 4-3(c) shows the middle ear cavity of the temporal bone with a wide open facial recess which was drilled open by Dr. Heidi Nakajima. In order to make contact with the umbo within the limited space, a 22 gauge hypodermic needle (1.5cm in length) was epoxied to the piezoelectric material and extended toward the umbo as shown in Fig. 4-3(d). The vibration from the umbo is transferred through the stiff needle to the piezoelectric sensor which is clamped at the micro-manipulator. Improved methods to mount the sensor on the mastoid closer to the umbo will have to be developed in the future.

4.3 Characterization Results from Human Cadaveric Temporal Bones

This section presents the characterization results of the umbo-mounted piezoelectric sensor using two different human cadaveric temporal bones, labeled bone096 and bone098. Measurements were taken between October 2011 and July 2013 at the

Massachusetts Eye and Ear Infirmary with the generous help of Dr. Heidi Nakajima.

4.3.1 Transfer Characteristics

The transfer characteristics of bone096 is shown in Figure 4-4. Fig. 4-4(a) shows the umbo velocity (normalized to P_{EC}) where both the shape and magnitude of the curve closely match those found in literature [131, 132]. The umbo velocity increases with frequency at a slope of +1 (i.e., +6dB/octave) up to 1kHz, peaking at 0.2 (mm/s)/Pa which is within the typical range. Fig. 4-4(b) shows the transfer characteristic from v_{UMBO} to the output of sensor front-end V_{PZ}. A peak is observed around 1.5kHz likely due to the resonant frequency of the sensor and attached needle. A similar observation was found in [133] where the authors explore using fins attached to the sensor to dampen out the resonance. Combining the plots from Figures 4-4(a) and (b), the overall transfer characteristic from P_{EC} to V_{PZ} is shown in Fig. 4-4(c). V_{PZ} shows a band-pass characteristic which peaks at 1kHz and falls off steeply beyond 7kHz which is high enough to capture the characteristics of speech.

An important observation of the V_{PZ}/P_{EC} response is that it increases with a slope of +6dB/octave up to around 1kHz. This is important because most CI sound processing strategies require a pre-emphasis high-pass filter with a slope of +6dB/octave up to 1.2kHz [134]. The pre-emphasis filter compensates for the -6dB/octave roll-off which occurs in speech spectrum that is radiated from the lips. Therefore, in any system using this mounted sensor, the pre-emphasis is already embedded.

From the umbo velocity, the displacement and acceleration can be calculated by integrating and differentiating the velocity curve. Figure 4-5 shows the corresponding umbo displacement and acceleration calculated from the velocity. The umbo displacement curve (Fig. 4-5(a)) is flat at low frequencies and rolls off after 1kHz which matches the results from [119, 135]. Furthermore, the low frequency displacement of 30 nm/Pa is also within the range of 20-40 nm/Pa found in the literature. Similarly, the umbo acceleration curve (Fig. 4-5(c)) increases at a rate of +12dB/octave up to around 1kHz which matches measurements found in [119]. Therefore, it is safe to conclude that the measurements are from a representative temporal bone.

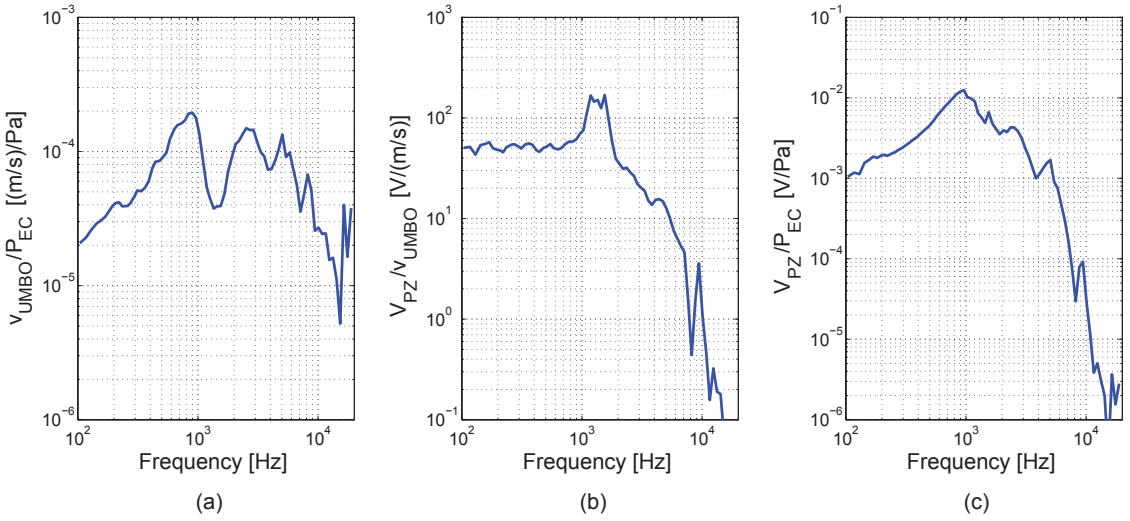

Figure 4-4: Transfer characteristics (a) from ear canal pressure (P_{EC}) to umbo velocity (v_{UMBO}), (b) from v_{UMBO} to the charge amplifier output voltage (V_{PZ}), and (c) from P_{EC} to V_{PZ}.

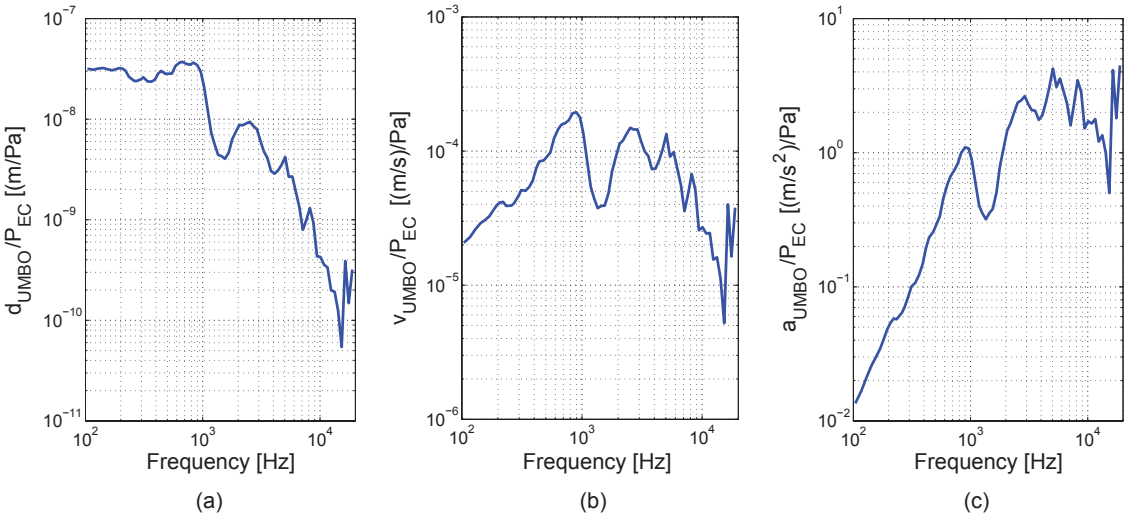

Figure 4-5: Spectra of the (a) umbo displacement, (b) umbo velocity, and (c) umbo acceleration normalized to the input sound pressure. The umbo velocity is measured with laser Doppler vibrometry.

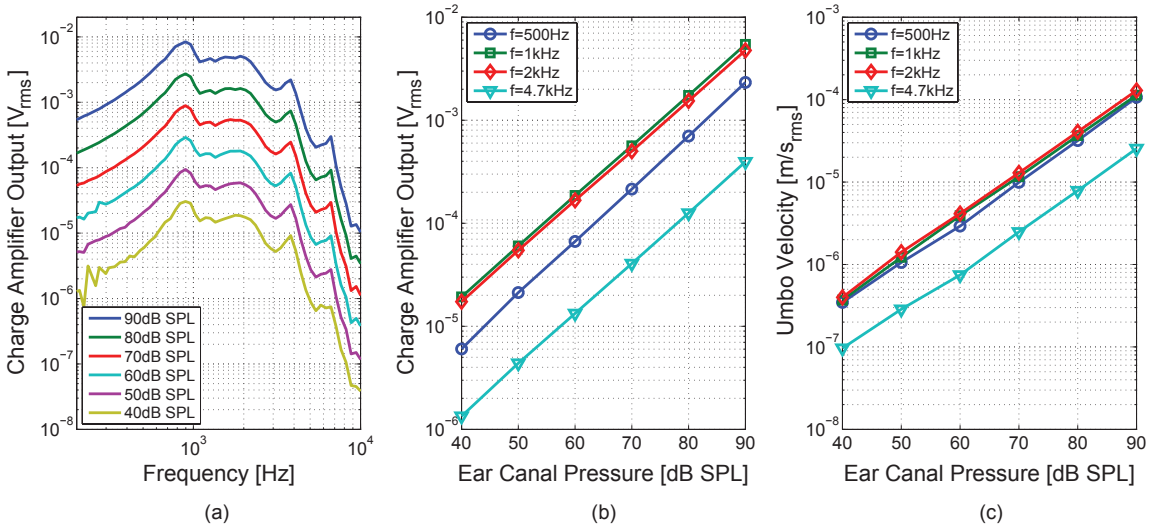

Figure 4-6: (a) Spectrum of the charge amplifier output for sound pressure levels from 40 to 90dB SPL. (b) Charge amplifier output voltage and (c) umbo velocity versus ear canal pressure at 500Hz, 1000Hz, 2000Hz, and 4700Hz.

4.3.2 Channel Linearity

Figure 4-6(a) shows the output of the charge amplifier for sound pressure levels from 40 to 90dB SPL in the ear canal of bone098. Considering that conversational speech ranges from 45 to 75dB SPL and that the dynamic range of speech is at most 50dB [136], the piezoelectric sensor has adequate performance in terms of sensitivity and dynamic range. In comparison to accelerometer-based systems [119] which has a minimum detectable signal of 60dB SPL at 500Hz and 35dB SPL at 2kHz, the piezoelectric sensor in this work has better low-frequency sensitivity.

Figures 4-6(b) and (c) show the linearity of the charge amplifier output and umbo velocity versus the ear canal pressure respectively. Both the piezoelectric sensor and middle ear mechanics show excellent linearity with respect to the ear canal pressure.

4.3.3 Repeatability and Umbo Loading

This section presents measurement results that address the repeatability of the sensor readout over time and across temporal bone samples, as well as the effect of umbo loading. Fig. 4-7(a) shows the umbo velocity of bone096 tested twice over a period

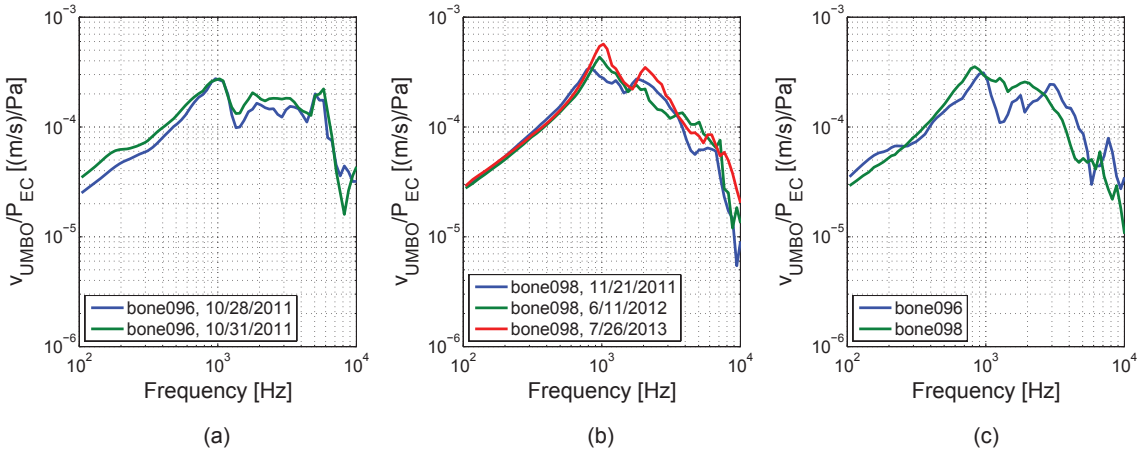

Figure 4-7: Transfer characteristic from ear canal pressure (P_{EC}) to umbo velocity (v_{UMBO}) for (a) bone096 measured twice over 4 days, (b) bone098 measured three times over 20 months, and (c) a comparison between two bones.

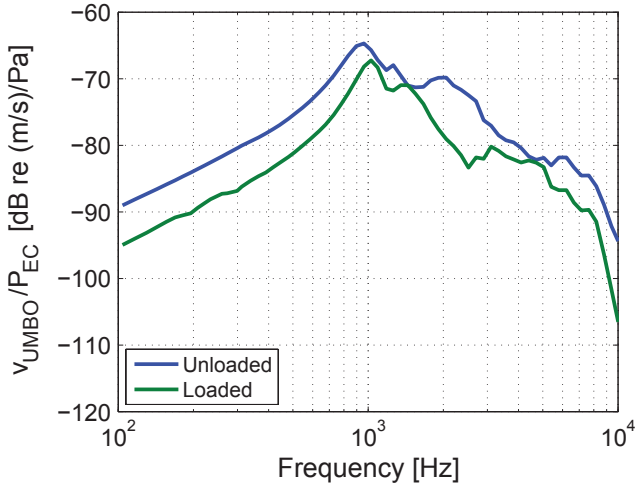

Figure 4-8: The effect of loading the umbo with the piezoelectric sensor on the transfer characteristic from ear canal pressure (P_{EC}) to umbo velocity (v_{UMBO}). Data is measured from bone098.

of 4 days, and Fig. 4-7(b) shows the umbo velocity of bone098 tested three times over the course of 20 months. In general, the temporal bones show good repeatability over both short and long term. The low-frequency behavior of bone096 varied by just a few dB, while the peak of the velocity for bone098 shifted slightly over time. Each bone was frozen in between measurements, and thawed adequately before each measurement. Fig. 4-7(c) compares the umbo velocity of both bones. The shape of the umbo velocity curves is remarkably similar despite being measured from two different temporal bones.

Finally, the effect of loading the umbo with the sensor is shown in Figure 4-8. On average, umbo loading decreases the umbo velocity by about 5dB.

4.4 Chapter Summary

A method for implantable acoustic sensing using a piezoelectric sensor mounted at the umbo of the middle ear is presented in this chapter. The sensing method is inspired by middle ear implants used to treat conductive hearing loss. However, when used in a CI system for sensorineural hearing loss, disarticulation of the ossicular chain is not necessary because there is no mechanical feedback from an output transducer as in a MEI.

A discrete prototype of a piezoelectric sensor front-end is used to help characterize the sensor which is mounted on two human cadaveric temporal bones. The measured umbo motion (displacement, velocity, and acceleration) using laser Doppler vibrometry shows good agreement with results found in the literature, ensuring that the characterization results are from representative temporal bones. Both the middle ear mechanics and piezoelectric sensor readout show excellent linearity, and the sensor is able to detect sounds from 300Hz to 10kHz over a 50dB dynamic range from 40 to 90dB SPL. Furthermore, a pre-emphasis of +6dB/octave is embedded in the sensor output which is useful for CI sound processing. Finally, the sensor readout is shown to be repeatable over time and across temporal bone samples.

Chapter 5

An Invisible Cochlear Implant - Part III: System-on-Chip

In the previous chapter, an umbo-mounted piezoelectric sensor was used to demonstrate implantable acoustic sensing with human cadaveric temporal bones. That work was inspired by the sensors used in middle ear implants, but is instead applied to a fully-implantable cochlear implant (FICI) in this work. A necessary step toward achieving the goal of a FICI is the development of electronic hardware that enables many of the functions that a FICI requires. Therefore, this chapter focuses on the design and implementation of a proof-of-concept system-on-chip (SoC) for a fully-implantable cochlear implant.

The goal of the work described in this chapter is to develop an ultra-low-power chip to incorporate many of the earlier results, together with sound processing capabilities that is required in cochlear implant systems. The first contribution of this chapter is the design of an integrated piezoelectric sensor front-end optimized for low noise and ultra-low power consumption when compared to the discrete prototype with off-the-shelf components described in Chapter 4. The second contribution is the design of a highly-reconfigurable multi-rate sound processor that leverages the energy-efficiency of ultra-low-voltage digital design. Finally, the third contribution is the design of an arbitrary waveform stimulator to allow the SoC to deliver energy-efficient stimulation waveforms as discussed in Chapter 3. A fabricated prototype SoC in a

high-voltage 0.18μm CMOS process is used to demonstrate the feasibility of future fully-implantable CIs.

This chapter is organized as follows: Section 5.1 presents a brief background on sound processing in today's CIs. Section 5.2 provides an overview of the architecture of the CI SoC. Sections 5.3, 5.4, and 5.5 describe the detailed analysis and design of the three subsystems respectively. Sections 5.6 and 5.7 present detailed measurement and system demonstration results of the SoC prototype, and Section 5.8 provides a summary of the chapter. The design of the stimulator subsystem was carried out in collaboration with Rui Jin who implemented the core analog circuits of the stimulator.

5.1 Background

The sound processing and stimulation strategies of a cochlear implant play a crucial role in the speech recognition performance of a cochlear implant user [13, 137, 138]. This section provides a brief review of current techniques.

5.1.1 Cochlear Implant Sound Processing Strategies

Currently, all commercially available cochlear implants employ a multi-electrode array to stimulate the auditory nerve at different locations in order to emulate the tonotopic structure of the cochlea. However, the performance of the cochlear implant depends on the sound processing strategy (i.e., how information is extracted from the acoustic signal and what information is delivered to each electrode). Multi-channel sound processing strategies can be loosely classified into three categories: (1) envelope-based strategies, (2) n-of-m strategies, and (3) feature-based strategies [137]. A fourth category of fine structure processing techniques which involves delivering phase and timing information is a current area of research. Complete surveys on sound processing strategies can be found in [13, 137–139], but some common strategies are highlighted here:

- **Continuous Interleaved Sampling**: An example of an envelope-based ap-

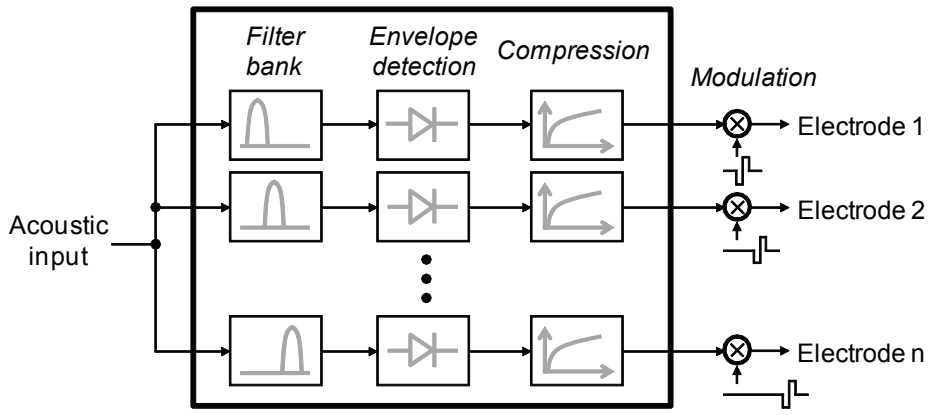

Figure 5-1: Block diagram of the Continuous Interleaved Sampling (CIS) sound processing strategy.

proach is the Continuous Interleaved Sampling (CIS) strategy. It is the default and most ubiquitous strategy used by all cochlear implant manufacturers. Figure 5-1 shows a block diagram of the CIS strategy [140]. Each electrode delivers biphasic current pulses in a time-interleaved (non-overlapping) manner to minimize channel-to-channel interaction in the cochlea. The amplitude of the pulses is determined by the amount of energy in a particular frequency band. This is accomplished by first passing the sound through a bank of logarithmically spaced band-pass filters. The envelope of each channel is then extracted with rectification and low-pass filtering. Following that, the envelope is compressed by a non-linear function (usually logarithmic) in order to scale the dynamic range of the acoustic signal to the dynamic range of electrical hearing. Finally, the compressed envelope is used to modulate the amplitude of a train of interleaved current pulses. The advantage of this approach is that a reasonably high rate of stimulation (several 100's to 1000's of pulses per second) can be achieved with minimal simultaneous channel interaction. A high rate of stimulation is important for preserving temporal information in speech [141].

- n**-of-**m **Strategies**: The Spectral Peak (SPEAK) strategy is an example of a n-of-m approach used by the Nucleus processor from Cochlear Limited [142]. This strategy is similar to CIS, except that the sound is filtered by a bank of 20 (i.e.,

$m = 20$) band-pass filters with center frequencies spanning 250Hz to 10kHz. The envelope of each channel is extracted as in the CIS approach, but only the 5 to 10 (i.e., $n = 5$ to 10) largest amplitudes are selected depending on the frequency content of the sound. The stimulation rate of each electrode varies from 180 to 300 pulses/sec depending on the number of maxima selected, as well as the patient's preference. For broadband sound, more channels are selected and the stimulation rate is reduced. For narrowband sound, less maxima are selected and the stimulation rate is increased to provide more temporal information [137]. Note that if $n = m$, this strategy is essentially the CIS strategy.

- **Feature-based Strategies**: The F0/F2 and F0/F1/F2 strategies are feature-based approaches which involve delivering stimulation based on information about the fundamental frequency (F0) and the first two formants of speech (F1 and F2) [137]. In the F0/F2 strategy, the fundamental frequency and the second formant are extracted using band-pass filters and zero-crossing detectors. The amplitude of F2 is also estimated through envelope detection. During voiced speech, the appropriate electrode determined by the value of F2 is stimulated at a rate of F0 pulses/sec, with the current strength set by the amplitude of F2. During un-voiced speech, the electrode is stimulated at quasi-random intervals at an average rate of 100 pulses/sec. In the F0/F1/F2 strategy, the first formant frequency (F1) and its amplitude are also extracted, and an extra electrode is stimulated according to the F1 frequency and energy.

There have been numerous studies on the effectiveness of these strategies in recognition of vowels, consonants, and sentences in both noisy and quiet environments [137, 140]. Although the CIS strategy is consistently the best performer, the performance of each strategy is still largely patient-specific and therefore modern sound processors often offer more than one strategy [13].

5.1.2 Number of Spectral Channels

Aside from the sound processing strategy, the number of spectral channels in a processor is also an important parameter. In [141], using acoustic simulations with normal hearing listeners in a quiet environment, Shannon et al. showed that as few as 3 to 4 channels of spectral information can result in good speech recognition performance. This implied that dynamic temporal information in a few broad spectral channels may be sufficient for speech recognition. A similar study with CI users in quiet using the SPEAK processor was presented in [143]. There, Fishman and Shannon found that the average performance improved as the number of stimulation electrodes was increased from 1 to 4, but no differences were found between 7-, 10-, or 20-electrode processors.

Although good performance can be attained with a small number of channels in quiet, the introduction of background noise can severely degrade performance in CI users. In [144], Fu et al. showed that the performance of normal hearing listeners using acoustic simulations improved up to approximately 16 to 20 channels in noise. However, for actual CI users in noise, the best CI listeners improved their performance only up to 7 electrodes. CI users with low levels of speech recognition could not benefit from more than 4 electrodes. These studies raise questions about whether or not CI users can actually make use of all the electrodes (up to 22) that are available in today's CIs.

5.2 Architecture Overview

A detailed block diagram of the fully-implantable cochlear implant system-on-chip (SoC) is shown in Figure 5-2. The system can be separated into three main subsystems: 1) the piezoelectric sensor front-end (PZFE), 2) the low-voltage reconfigurable sound processor, and 3) the arbitrary waveform stimulator and high-voltage electrode switch matrix.

The PZFE operates from a 1.5V analog power supply and comprises three stages: a charge amplifier (CA) to interface with the umbo-mounted piezoelectric sensor

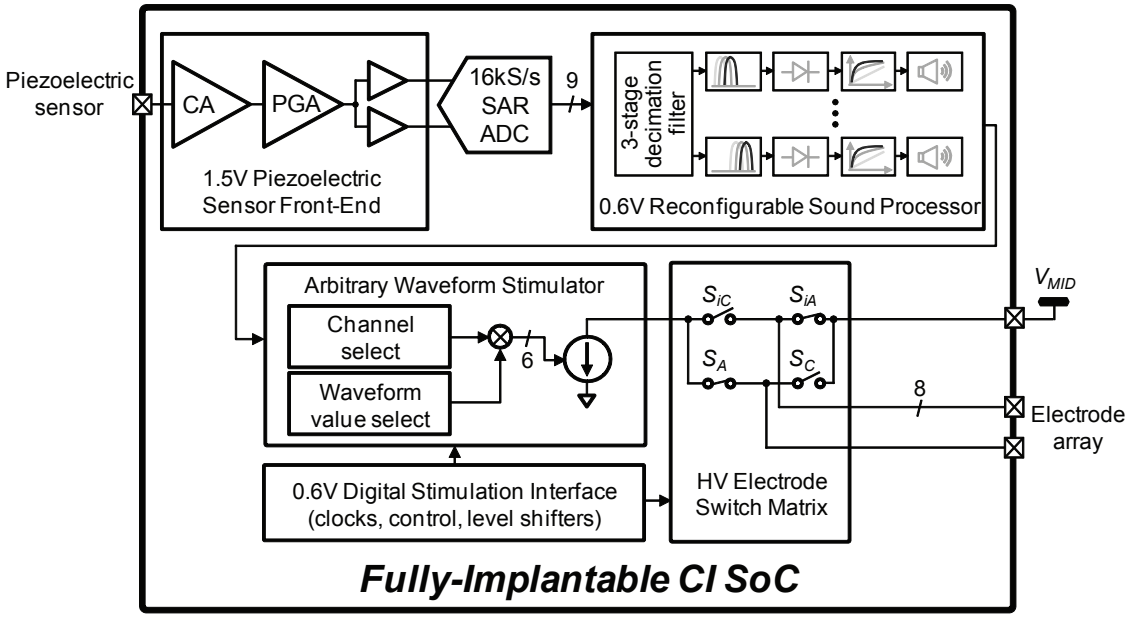

Figure 5-2: Block diagram of the fully-implantable cochlear implant system-on-chip.

(described in Chapter 4), a programmable-gain amplifier (PGA), and a single-ended to differential ADC driver. The PZFE provides a mid-rail reference voltage $V_{ref,PZ}$ to bias one terminal of the sensor, while the other terminal is connected to the input of the CA. The ADC driver stage also provides analog level conversion from $V_{ref,PZ}$ = 750mV down to $V_{adc,cm}$ = 300mV which is the input common-mode for the ADC. The ADC is a differential 16kS/s 9-bit SAR ADC operating from a supply of 0.6V.

The output of the ADC is processed by a 0.6V reconfigurable digital sound processor that implements the CIS sound processing strategy. The processor in this work has many programmable features. First, the number of channels can be configured between 4-, 6-, or 8-channel modes to enable a power-performance tradeoff. The filter bank has reconfigurable coefficients to adjust the filter bandwidths for the three modes of operation, and multi-rate signal processing is leveraged to reduce power and area. The channels are logarithmically spaced to emulate the tonotopic structure of hearing, and the bandwidth of the processor covers 300Hz to 5.5kHz. Furthermore, processor settings like global channel gain, rectification type, and amount of compression are all programmable. The processor also has the capability to adjust the volume level (i.e., threshold and most comfortable level) for each channel individually

Table 5.1: Summary of required supply voltages in the CI SoC.

Block(s)	Supply Voltage [V]
Piezoelectric sensor front-end	1.5
SAR ADC Digital CIS sound processor Digital stimulation interface	0.6
Intermediate level shifters	1.8
Stimulator current source circuits	3.3
V_{MID}	5 to 10
V_{DDG}	7 to 12

to provide patient fitting capability.

The processor outputs data at an analysis rate of 1kHz. The output of each channel is a 6-bit value that represents the logarithmically compressed energy in each frequency band. This value is used to modulate a train of current pulses (1,000 pulses/sec per electrode) that is delivered to the electrode corresponding to the channel (high-frequency channels map to basal electrodes, low-frequency channels map to apical electrodes).

The interleaved operation of the CIS sound processing strategy conveniently allows for a single current source to be interleaved between all channels. This is accomplished by using a high-voltage electrode switch matrix to select the active electrode and control the direction of current flow. The SoC is designed to be used with monopolar electrode arrays, where a common return electrode is used as the return path for all electrodes. Furthermore, an on-chip digital stimulation interface controller allows the waveform of the stimulation pulses to be programmed to any arbitrary shape. The shape parameters are loaded onto the SoC through a serial programming interface. The controller also provides the control signals for the switch matrix.

Since the stimulator drives the electrode-tissue interface which may have impedances in the 10's of kΩ range, it must have high voltage compliance in order to deliver 100's of μA of stimulation current. Therefore, the stimulation subsystem does require a number of supply voltages. The stimulation current is drawn from a high voltage supply (V_{MID}) which can be 5V to 10V. The control signals to the high-voltage

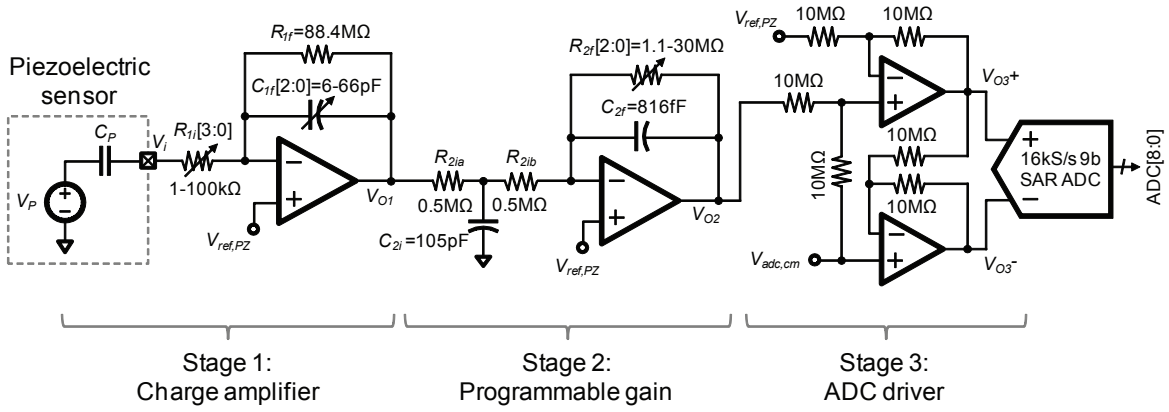

Figure 5-3: Schematic of the 3-stage piezoelectric sensor front-end with a 16kS/s, 9-bit SAR ADC.

electrode switch matrix operate from V_{DDG} which can range from 7V to 12V. The digital stimulation interface controller performs its computations at 0.6V, and 1.8V level shifters are required between the 0.6V and V_{DDG} domains. Lastly, the current source circuits operate from a supply voltage of 3.3V. A summary of the required supply voltages in the SoC is provided in Table 5.1.

5.3 Piezoelectric Sensor Front-End

The design of a front-end prototype for a piezoelectric sensor was explored in Section 4.2, but an integrated circuit version is optimized here for low-power and low-noise performance. Figure 5-3 provides an overview of the 3-stage piezoelectric sensor front-end (PZFE) and ADC. Details of each stage are described in this section.

5.3.1 Stage 1: Charge Amplifier

The architecture of the integrated charge amplifier is identical to the one used in the discrete prototype of Chapter 4, but noise analysis and rationale for component sizing is provided here. Figure 5-4(a) shows stage 1 of the PZFE with noise sources, and Figure 5-4(b) shows the equivalent block diagram, where $v_{n,i}$, $v_{n,f}$, and $v_{n,a}$ are the noise from R_i, R_f, and the op-amp respectively, and $A(s)$ is the open-loop transfer

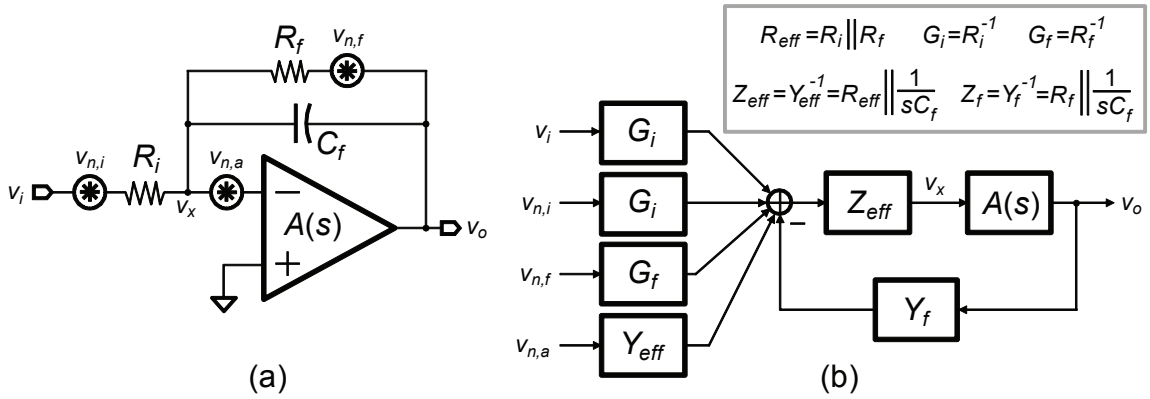

Figure 5-4: (a) Equivalent circuit for the charge amplifier in Figure 5-3 including noise sources, and (b) the corresponding block diagram.

function of the op-amp. The block diagram can be used to determine the transfer functions from the input and noise sources to the output. Since we are interested in determining the transfer function from the piezoelectric sensor voltage V_P to the output, the input voltage v_i can be referred back to V_P through the following transfer function

$$H_P(s) = \frac{v_i}{V_P} = \frac{s}{s+\omega_i}, \tag{5.1}$$

where $\omega_i = \frac{1}{R_i C_P}$. For the frequencies of interest, the loop gain of the feedback loop is large (because of large $A(s)$) and so the closed-loop transfer function of the feedback loop is the inverse of the feedback factor Y_f,

$$H_{CL}(s) \approx \frac{1}{Y_f} = \frac{1}{C_f(s+\omega_f)} \tag{5.2}$$

where $\omega_f = \frac{1}{R_f C_f}$. Therefore, the signal transfer function of stage 1 from V_P to v_o is given by $H_{STG1}(s) = \frac{v_o}{V_P} = \frac{v_o}{v_i}\frac{v_i}{V_P}$ which evaluates to

$$H_{STG1}(s) = G_i H_{CL}(s) H_P(s) = \frac{C_P}{C_f}\frac{s\omega_i}{(s+\omega_f)(s+\omega_i)} \tag{5.3}$$

which is equivalent to the result in Equation 4.4 but with V_P in the denominator (i.e., using the voltage model of the sensor). Note that the negative polarity of the charge amplifier has been ignored here for simplicity. The low-frequency high-pass

corner is set by ω_f, while the high-frequency low-pass corner is set by ω_i. Since the sound processor bandwidth is 300Hz to 5.5kHz, the constraints for ω_f and ω_i are $\omega_f < 2\pi(300\text{Hz})$ and $\omega_i > 2\pi(5.5\text{kHz})$. Within the band of interest ($\omega_f < \omega < \omega_i$), the mid-band voltage gain is simply C_P/C_f.

Since typical values of C_P range from 0.2nF to 3nF (Section 4.1.3), C_f is a tunable capacitor that is sized small enough to provide adequate gain for small values of C_P, and large enough to limit the gain for large values of C_P so as to not saturate the front-end at high sound pressure levels. The implemented feedback capacitor (called C_{1f}[2:0] in Fig. 5-3) is a 3-bit switched-capacitor that is tunable from 6pF to 66pF and is non-uniformly spaced to provide programmable mid-band gain in 3dB steps. Given that the high-pass corner of the charge amplifier is set by $f_f = \frac{1}{2\pi R_f C_f}$ which must be <300Hz in this system, R_f is constrained by the minimum value of C_f and is set to 88.4MΩ (called R_{1f} in Fig. 5-3). Finally, R_i is a tunable resistor which ensures $f_i >$ 5.5kHz for all expected values of C_P. It is implemented as a 4-bit switched-resistor (called R_{1i}[3:0] in Fig. 5-3) which is logarithmically spaced from 1kΩ to 100kΩ.

Noise Analysis

The expected signal magnitude from a 3mm × 3mm × 0.5mm sensor was analyzed back in Section 4.1.3. At 40dB SPL, the minimum expected signal is approximately $3\mu\text{V}_{\text{rms}}$ which sets an upper bound on the noise of the PZFE. Here, we analyze the noise of stage 1 which dominates the noise performance of the entire PZFE.

Noise from R_i: The noise transfer functions referred to V_P can be determined by calculating the noise transfer function to the output v_o, and then dividing by $H_{STG1}(s)$. For R_i, the noise transfer function referred to V_P is $H_{np,i}(s) = \frac{V_P}{v_{n,i}} = \frac{v_o}{v_{n,i}}\frac{V_P}{v_o}$ which evaluates to

$$H_{np,i}(s) = G_i H_{CL}(s) \cdot \frac{1}{H_{STG1}(s)} = \frac{s+\omega_i}{s}. \tag{5.4}$$

For the frequencies of interest ($\omega \ll \omega_i$), this reduces to $H_{np,i}(s) \approx \omega_i/s$. Therefore,

the noise spectral density of R_i referred to V_P ($V^2_{np,i}(f)$) can be determined by multiplying the noise spectral density of R_i ($V^2_{n,i}(f) = 4kTR_i$) by the square of $H_{np,i}(f)$,

$$V^2_{np,i}(f) = |H_{np,i}(f)|^2 V^2_{n,i}(f) = \frac{4kT}{R_i}\left(\frac{1}{2\pi f C_P}\right)^2. \tag{5.5}$$

Noise from R_f: Following the same analysis as above, the noise transfer function of R_f (referred to V_P) is

$$H_{np,f}(s) = G_f H_{CL}(s) \cdot \frac{1}{H_{STG1}(s)} = \frac{R_i}{R_f}\frac{s+\omega_i}{s} \approx \frac{R_i}{R_f}\frac{\omega_i}{s} \tag{5.6}$$

in the bandwidth of interest ($\omega \ll \omega_i$). Therefore, the noise spectral density of R_f referred to V_P is

$$V^2_{np,f}(f) = |H_{np,f}(f)|^2 V^2_{n,f}(f) = \frac{4kT}{R_f}\left(\frac{1}{2\pi f C_P}\right)^2. \tag{5.7}$$

Noise from the op-amp: Finally, the noise transfer function of the op-amp referred to V_P is

$$H_{np,a}(s) = Y_{eff} H_{CL}(s) \cdot \frac{1}{H_{STG1}(s)} = \frac{R_i C_f (s+\omega_{eff})(s+\omega_i)}{s} \tag{5.8}$$

where $\omega_{eff} = \frac{1}{R_{eff}C_f} \approx \frac{1}{R_i C_f}$ since $R_f \gg R_i$. For typical component values, $\omega \ll \omega_i \ll \omega_{eff}$, so the transfer function reduces to $H_{np,a}(s) \approx \omega_i/s$. Therefore, the noise spectral density of the op-amp referred to V_P is

$$V^2_{np,a}(f) = |H_{np,a}(f)|^2 V^2_{n,a}(f) = \left(\frac{1}{2\pi f R_i C_P}\right)^2 \left[\frac{4kT}{\kappa g_{m1}}\left(\frac{\text{NEF}}{2.02}\right) + \frac{K_f}{WLC_{ox}f}\alpha_f\right] \tag{5.9}$$

where the first and second terms in the square bracket are the total thermal noise and $1/f$-noise of the op-amp respectively (similar to the analysis in Section 2.3.3). The details of the thermal noise term are provided later in this section in Equation 5.11, and the $1/f$-noise term can be made negligible by sizing W and L large enough.

Total noise in stage 1: The total noise spectral density of stage 1 referred to V_P is calculated by summing up the noise from R_i, R_f, and the op-amp as follows,

$$\begin{aligned} V_{np,tot}^2(f) &= V_{np,i}^2(f) + V_{np,f}^2(f) + V_{np,a}^2(f) \\ &= \underbrace{\frac{4kT}{R_i}\left(\frac{1}{2\pi f C_P}\right)^2}_{\text{noise from } R_i} + \underbrace{\frac{4kT}{R_f}\left(\frac{1}{2\pi f C_P}\right)^2}_{\text{noise from } R_f} + \\ &\quad + \underbrace{\left[\frac{4kT}{\kappa g_{m1}}\left(\frac{\text{NEF}}{2.02}\right) + \frac{K_f}{WLC_{ox}f}\alpha_f\right]\left(\frac{1}{2\pi f R_i C_P}\right)^2}_{\text{noise from op-amp}}. \end{aligned} \quad (5.10)$$

This expression is valid for $\omega \ll \omega_i$ within the bandwidth of interest. From Equation 5.10, we see that the noise spectral density of R_i, R_f, and the op-amp thermal noise all have a $1/f^2$ shape, while the op-amp flicker noise has a $1/f^3$ characteristic. The input-referred noise due to R_i and R_f is reduced for larger values of C_P. However, the op-amp noise depends on the product $R_iC_P = \frac{1}{\omega_i}$ which is generally a fixed parameter. Therefore, the op-amp noise is independent of C_P (since R_i is varied to keep R_iC_P constant) and does not benefit from larger values of C_P. Furthermore, in general, the low-frequency input-referred noise is independent of C_f.

From noise simulations with typical component values, the noise from R_f is negligible because of its large value ($R_f = 88.4\text{M}\Omega$). The relative contributions of noise from R_i and the op-amp vary depending on the value of C_P. For $C_P = 0.5\text{nF}$, the simulated total integrated noise from 300Hz to 5.5kHz is $2.5\mu\text{V}_{\text{rms}}$, where R_i accounts for approximately 60% of the total noise power (40% from the op-amp). For $C_P =$ 3nF, the total integrated noise is reduced to $1.7\mu\text{V}_{\text{rms}}$, where R_i accounts for only 20% of the total noise power (80% from the op-amp). Therefore, the op-amp noise dominates for large values of C_P, and its schematic design is considered next.

Stage 1 Op-Amp Schematic Design

Figure 5-5 shows the schematic of the op-amp used in stage 1. The input stage is a folded-cascode op-amp with source-degenerated bias transistors ($\text{M}_{5,6}$ and R_S) to

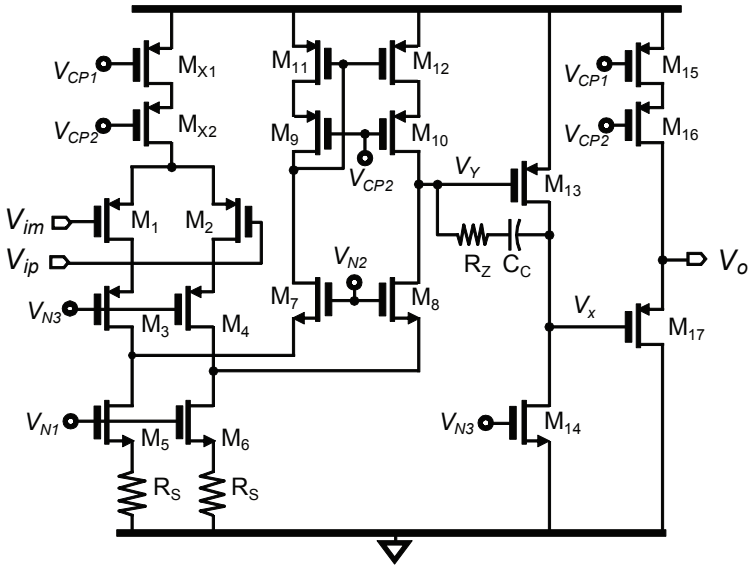

Figure 5-5: Schematic of the op-amp used in the charge amplifier.

improve the noise performance [145]. The input devices are PMOS transistors with very large W and L to limit the $1/f$-noise, and so the op-amp noise is dominated by thermal noise. The second stage is a common-source stage to increase its open loop gain, and the output stage is a PMOS source-follower with low output impedance which is needed to drive the resistive load presented by stage 2 of the PZFE.

The noise from the common-source and source-follower stages can be neglected because of the high gain provided by the folded-cascode stage. Also, any noise from the current source devices M_{X1} and M_{X2} is common-mode and therefore can be ignored. Furthermore, the noise of the cascode devices (M_3-M_{10}) can be neglected because their noise current circulates locally and does not add to the total noise [146]. Therefore, the only devices that contribute noise are $M_{1,2}$, $M_{11,12}$, and the source-degeneration resistors R_S. Detailed noise analysis can be found in [145], and the results are summarized here. Ignoring $1/f$-noise, the total input-referred thermal noise of the op-amp is

$$V_{n,a-th}^2(f) = \frac{4kT}{\kappa g_{m1}} \underbrace{\left[1 + \left(\frac{\alpha_2}{\alpha_1}\right)^2 \frac{2\kappa}{g_{m1}R_S} + \left(\frac{\alpha_3}{\alpha_1}\right)^2 \frac{4\kappa g_{m11}}{3g_{m1}}\right]}_{=\frac{\mathrm{NEF}}{2.02}}, \tag{5.11}$$

where g_{m1} and g_{m11} are the transconductances of $M_{1,2}$ and $M_{11,12}$ respectively, k is

the Boltzmann constant, T is the temperature in Kelvin, κ is a process constant (approximately 0.7), and α_1, α_2, and α_3 are the current transfer functions from the noise current of $M_{1,2}$, R_S, and $M_{11,12}$ to the output current respectively (note that $\alpha_{1,2,3} \leq 1$ and are typically close to 1 if the op-amp is designed properly). $V^2_{n,a-th}(f)$ is equal to the thermal noise term of Equation 5.9.

Equation 5.11 highlights the relative contributors to the total noise. The input pair contributes an input-referred noise of $\frac{4kT}{\kappa g_{m1}}$ which is the factored term in front of the square bracket. The source-degeneration resistors contribute an additional factor of $\chi_2 = \left(\frac{\alpha_2}{\alpha_1}\right)^2 \frac{2\kappa}{g_{m1}R_S}$, while devices $M_{11,12}$ (biased in strong-inversion) contribute a factor of $\chi_3 = \left(\frac{\alpha_3}{\alpha_1}\right)^2 \frac{4\kappa g_{m11}}{3g_{m1}}$. Therefore, the sum of the terms in the square bracket $(1 + \chi_2 + \chi_3)$ represents the amount of noise beyond the noise of the input pair, and is equal to $\frac{\text{NEF}}{2.02}$ where 2.02 is the noise efficiency factor (NEF) for a MOS differential pair in sub-threshold [145].

The bias current of the input devices (I_{D1}) is chosen to provide the required g_{m1} for the desired noise floor. The input devices $M_{1,2}$ are biased in sub-threshold to maximize their g_m/I_D ratio, thus maximizing g_{m1} for a given I_{D1}. The source-degeneration resistors R_S effectively replace the noise of $M_{5,6}$ with their own noise, which can be made much lower as long as $R_S \gg 1/g_{m5,6}$. Another benefit of resistive source-degeneration is that resistors contribute only white noise, while $M_{5,6}$ would contribute both thermal and $1/f$-noise [145]. In order to minimize χ_2, $R_S \gg 1/g_{m1}$ is also satisfied. Finally, χ_3 is minimized by ensuring $g_{m11} \ll g_{m1}$ which is accomplished by setting $I_{D11} = I_{D1}/10$ and biasing $M_{11,12}$ in strong-inversion to minimize g_{m11}.

Considering the techniques just described, the resulting values for χ_2 and χ_3 from simulation are 0.18 and 0.073 respectively, resulting in a simulated NEF = 2.02×(1 + 0.18 + 0.073) = 2.53 for the op-amp, which is just slightly above the theoretical minimum of 2.02.

5.3.2 Stage 2: Programmable-Gain Amplifier

The available gain from the CA in stage 1 is determined by the size of the sensor (C_P), and the available values of C_{1f}[2:0] from 6pF to 66pF. Therefore, stage 2 of the PZFE

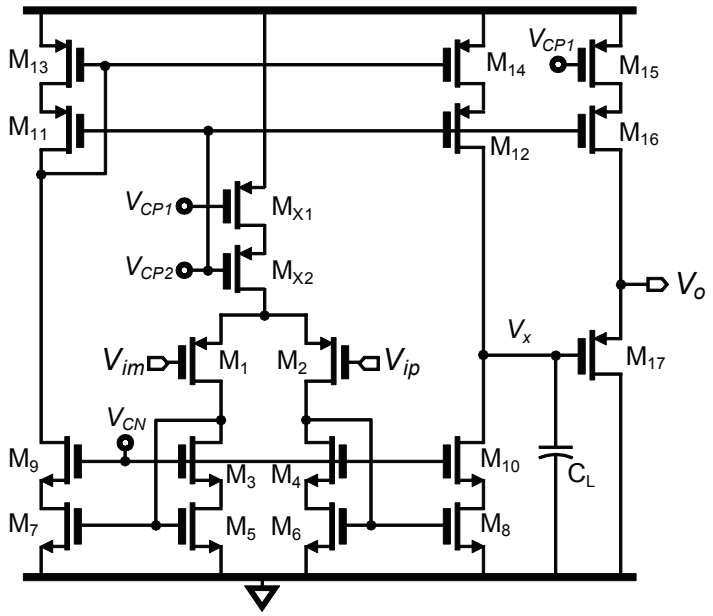

Figure 5-6: Schematic of the op-amp used in the programmable-gain amplifier.

is a 2-pole programmable-gain amplifier (PGA) to provide additional programmable gain if necessary. The transfer function of the PGA shown in Figure 5-3 is

$$H_{STG2}(s) = \frac{-R_{2f}}{R_{2i}} \cdot \frac{\omega_{2i}\omega_{2f}}{(s+\omega_{2i})(s+\omega_{2f})} \tag{5.12}$$

where $R_{2i} = R_{2ia} + R_{2ib}$, and $\omega_{2i} = \frac{1}{(R_{2ia}||R_{2ib})C_{2i}}$ and $\omega_{2f} = \frac{1}{R_{2f}C_{2f}}$ are the two poles of the PGA that help provide anti-aliasing in the PZFE. The DC gain is $A_{STG2} = -R_{2f}/R_{2i}$, where $R_{2i} = 1\text{M}\Omega$ and R_{2f} is a switched-resistor that is logarithmically spaced between 1.1MΩ and 30MΩ to provide programmable gain in 6dB steps (from 0.83dB to 29.5dB). Given that $R_{2ia}||R_{2ib} = 250\text{k}\Omega$ and $C_{2i} = 105\text{pF}$, the first pole ω_{2i} is set to 6kHz. The second pole ω_{2f} is set by R_{2f} and $C_{2f} = 816\text{fF}$, and is typically a higher frequency pole for small values of R_{2f}. However, at the maximum value of $R_{2f} = 30\text{M}\Omega$, ω_{2f} has a minimum value of 6.5kHz which provides additional filtering.

The schematic of the op-amp used in stage 2 is shown in Figure 5-6. The input stage is a cascoded current mirror op-amp to achieve high gain, and good input and output range. The output stage is a PMOS source-follower to provide low output impedance to drive the resistive load presented by stage 3 of the PZFE.

The tradeoff with a cascoded current mirror op-amp is that it has higher noise (when compared to an alternative like a folded-cascode stage). However, low-noise design is not as crucial in stage 2 because of the gain provided by stage 1. The noise of stage 2 decreases with larger values of R_{2f}, and is in general negligible when input-referred through stage 1.

5.3.3 Stage 3: ADC Driver and Low-Voltage SAR ADC

The third and last stage of the PZFE shown in Figure 5-3 is a single-ended to differential amplifier that must be able to drive the input capacitance of the SAR ADC which is approximately 480fF (differentially). This is accomplished with a series connection of a non-inverting amplifier (gain = 2) and an inverting amplifier (gain = -1). Therefore, the single-ended to differential conversion provides an additional gain of 12dB (4V/V).

Since the ADC operates from a low supply voltage of 0.6V, stage 3 also provides analog level conversion from $V_{ref,PZ}$ = 750mV to the ADC input common-mode of $V_{adc,cm}$ = 300mV. This is accomplished with appropriate biasing of the feedback network of 10MΩ resistors. Finally, the op-amps used in stage 3[1] are two-stage op-amps that leverage the cascoded current mirror stage in Figure 5-6, with a high-power common-source output stage to drive the ADC capacitance.

The ADC in this work reuses the design from the MSFE for ECG monitoring described in Chapter 2, with the exception of the dual-DACs. The comparator and DAC circuits are ported from that work, and the SAR logic is simplified to control a single DAC in this work. The design details of a similar ADC can be found in [1].

[1]Note that a wide input and output range is not strictly required in stage 2 (since the input and output are both biased at $V_{ref,PZ}$ = 750mV), but the same cascoded current mirror op-amp stage is used in the ADC driver of stage 3 to reduce design time. In stage 3, the input and output DC operating point can be between 300mV and 525mV, and therefore the cascoded current mirror op-amp is suitable there.

5.4 Low-Voltage Reconfigurable CIS Sound Processor

The sound processor in this work implements the Continuous Interleaved Sampling (CIS) strategy [140] because it is the most ubiquitous strategy in cochlear implants today, and it is used by all cochlear implant manufacturers. The main goals for the design of the sound processor in this work are 1) ultra-low-power operation and 2) highly reconfigurable features to enable a power-performance tradeoff, and patient-specific fitting capabilities. The first goal is accomplished by leveraging ultra-low-voltage digital processing at 0.6V to operate the processor close to its minimum energy point to maximize energy-efficiency [147]. The second goal is addressed with a custom architecture featuring a multi-rate reconfigurable filter bank and highly-programmable processor parameters which are set via a serial programming interface. The processor was coded in Verilog and synthesized and placed using the Synopsys digital tool flow. Verification was performed with Synopsys NanoSim to ensure operation at 0.6V. The processor architecture is presented next, followed by details on power scalability, and the reconfigurable filter bank.

5.4.1 Processor Architecture

The block diagram of the reconfigurable CIS sound processor is shown in Figure 5-7, where reconfigurable features are labeled in blue. The processor spectrally decomposes the signal with a logarithmically spaced filter bank and then extracts the envelope from each channel (channels A to H represent low to high frequency channels). The envelope is then compressed to fit the patient's electric hearing dynamic range. The output of each channel is used to modulate a train of interleaved current pulses delivered to the corresponding electrode. The architecture in this work uses several techniques to reduce the overall power and area and provide highly reconfigurable features, and they are discussed next.

The number of channels can be reconfigured between 8-, 6-, and 4-channel modes

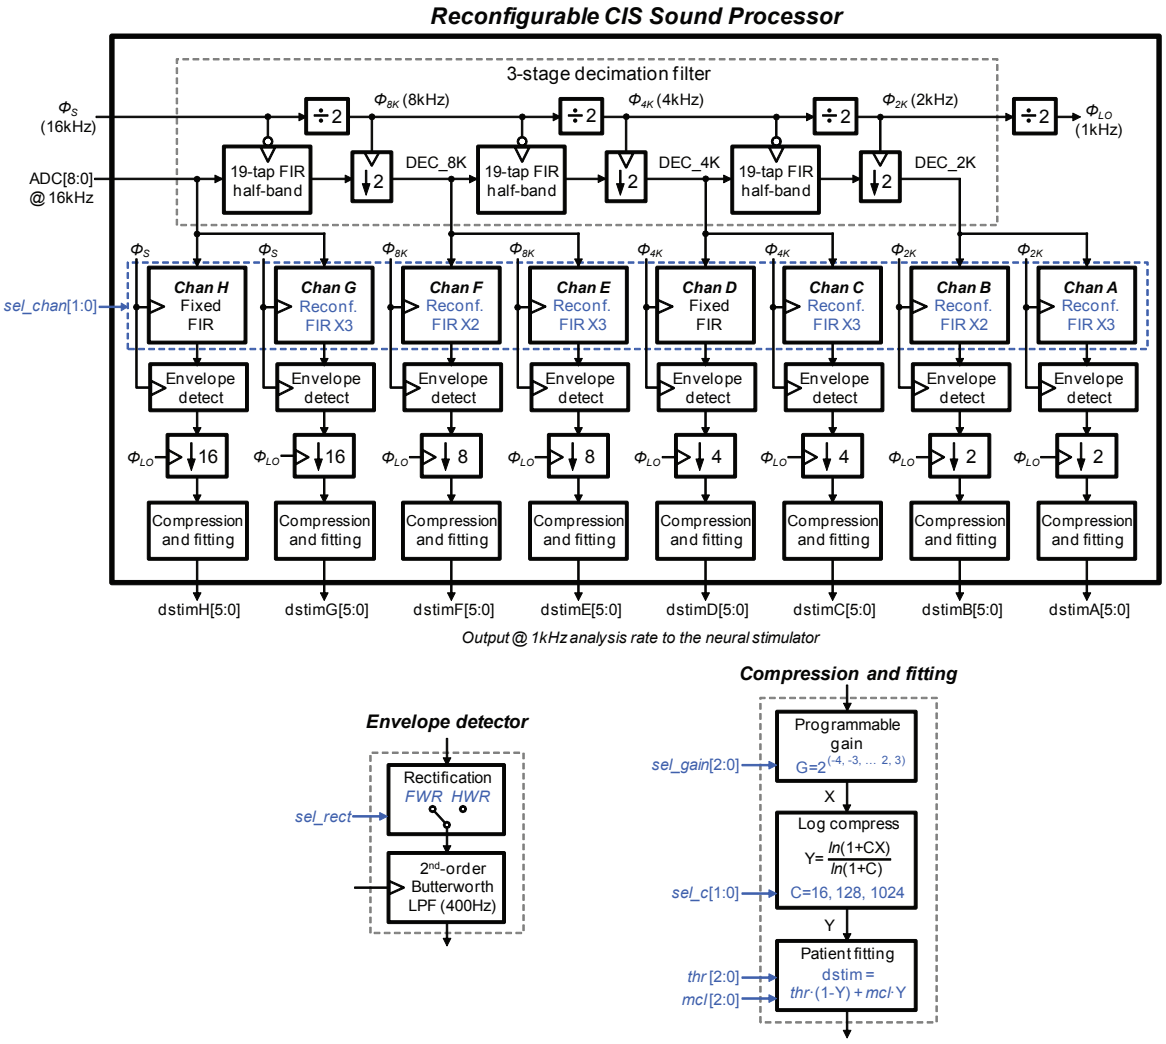

Figure 5-7: Block diagram of the reconfigurable multi-rate CIS sound processor, with the reconfigurable features highlighted in blue. The processor can be configured in 4-, 6-, or 8-channel mode by setting *sel_chan*[1:0]. Clock-gating is applied to channels that are turned off. Details for the *envelope detector* and *compression and fitting* blocks are shown at the bottom.

to enable a power-performance tradeoff. This is motivated by the findings in [141,143, 148,149] which show that good speech recognition scores in CI users can be achieved with as few as 4 channels, and patient performance improves with the number of channels up to around 7 or 8 channels, at which point the performance begins to saturate. As the number of channels changes, a different subset of channels are selected from channels A to H. Therefore the bandwidth and center frequency of each filter in the filter bank may need to change. This is accomplished with FIR filters with 3 levels of reconfigurability in both the coefficients and number of taps, and this will be discussed in Section 5.4.3. By properly clock-gating the processor, reducing the number of channels can reduce the processor power consumption. But more importantly, it reduces the neural stimulation power which dominates the entire SoC. The channel mode can be set with 2 bits (sel_chan[1:0]) in the serial programming interface.

A logarithmically spaced filter bank requires that higher frequency channels are wider in bandwidth, while low frequency channels need to be narrow and more selective. At a fixed sampling rate, this implies that the low frequency channels require higher order filters (e.g., more taps in a FIR filter) which increases both power and area[2]. In this work, this issue is addressed by using multi-rate signal processing, where the the low frequency channels operate at low sample rates, while high frequency channels operate at higher sample rates. This allows each filter to provide the necessary selectivity with a reasonable filter order. In the architecture shown in Figure 5-7, the incoming data from the ADC is decimated in 3 stages, resulting in data rates of 2kHz, 4kHz, 8kHz, and 16kHz. Since channels A to H represent low to high frequency channels, channels A/B, C/D, E/F, and G/H operate at a data rate of 2kHz, 4kHz, 8kHz, and 16kHz respectively. The decimation filter uses 19-tap half-band FIR filters (previously described in Section 2.3.8) to perform anti-aliasing.

Following the filter bank, the envelope of each channel is extracted by rectifying

[2]For example, a 2kHz wide band-pass FIR filter with 1kHz transition bands and 40dB stop band attenuation requires just 23 taps at a sampling rate of 16kHz. In contrast, a 200Hz wide band-pass FIR filter with 100Hz transition bands and 40dB stop band attenuation requires over 200 taps at the same sampling rate of 16kHz.

the signal and then passing it through a low-pass filter as shown in the bottom left of Figure 5-7. The type of rectification has an effect on speech recognition scores and sound quality [134], and therefore both full-wave (the default) and half-wave rectification are available in the envelope detector by setting the *sel_rect* bit. The low-pass cut-off of the envelope detector is not a crucial parameter as long as it is between 50Hz to 500Hz [141], and a standard 2^{nd}-order Butterworth filter with a 400Hz cut-off is used to limit the bandwidth of the envelope before it is downsampled to the output data rate of 1kHz.

The last step performed by the processor is dynamic range compression and patient fitting shown in the bottom right of Figure 5-7. The channel gain (applied globally to all channels) can be set with 3 bits (*sel_gain*[2:0]) from 2^{-4} to 2^{3}. Then, the signal is logarithmically compressed in amplitude according to the following function,

$$Y = \frac{\ln(1 + CX)}{\ln(1 + C)} \tag{5.13}$$

where C is the compression factor. Logarithmic compression is needed because of the well-known loudness growth function of electrical hearing that describes the linear relationship between acoustic sound intensity in dB SPL and electrical stimulation intensity in amperes [150]. A typical value for C in clinical processor fittings is 1024 [151]. However, there is evidence that different amounts of compression can be beneficial [136], and therefore three settings are available in this processor: $C = 1024$, 128, and 16, selectable using 2 bits (*sel_c*[1:0]). The compression functions are stored in three small look-up tables, each with 512 6-bit wide values, requiring approximately 200 gates each. Finally, each channel in the processor has individual threshold (THR) and most-comfortable-level (MCL) settings that can be used during an audiology fitting session to adjust the dynamic range between the minimum and maximum current level for each electrode to ensure a patient specific fit [151]. Each channel has 3-bit programmability in both THR and MCL. The outputs of the processor are dstimA[5:0] to dstimH[5:0] corresponding to channels A to H which are updated at an analysis rate of 1kHz.

5.4.2 Power Scalability

The reconfigurable number of channels and programmable processor parameters can influence the processor power, but more importantly, it has a greater impact on the stimulation power which dominates the overall SoC power. The scalability of the stimulation power with respect to the processor settings is summarized here.

- **Number of channels**: The sound processor power (from 0.6V), current source power (from 3.3V), and stimulation power (from 1.8V, V_{MID}, and V_{DDG}) all scale linearly with the number of channels. This is a consequence of clock-gating the digital circuits, and power-gating the analog circuits (e.g., the current source) when they are not in use.
- **Channel gain and compression factor**: In general, the stimulation power increases monotonically with the channel gain because more current is delivered at higher gains. For a small value of the compression factor (e.g., $C = 16$, which provides very little compression), the stimulation power is roughly linear with the channel gain. For a large value of the compression factor (e.g., $C = 1024$), the stimulator power becomes linear with the logarithm of the channel gain because the channel outputs are heavily compressed. Lastly, the stimulation power increases with the compression factor because lower magnitude envelopes are emphasized more with greater compression.
- **Rectification type**: Using full-wave rectification requires slightly more stimulation power than half-wave rectification because the average envelope value is larger.
- **THR and MCL settings**: Stimulation power increases with higher THR and MCL settings.
- **Stimulation parameters**: Stimulation power increases with the stimulation phase width because the current source is active for a longer amount of time. However, the increase in power is not strictly linear with phase width because less current is required at longer phase widths due to the strength-duration characteristic of nerve fibers.

5.4.3 Reconfigurable Filter Bank

Filter banks can be implemented with IIR or FIR filters. IIR filters are typically more hardware efficient, but they suffer from instability and non-linear phase (i.e., non-constant group delay) which can lead to frequency dispersion. Therefore, FIR filters were chosen for their linear phase (i.e., constant group delay) which can have a positive effect on sound quality and speech intelligibility [152]. Other benefits include their unconditional stability and regular structure which make them easier to design.

When the processor (Fig. 5-7) is configured in 8-channel mode, all channels are active and channels A to H map to channels 1 to 8. In 6-channel mode, channels D and H are clock-gated (i.e., turned-off) and channels A-C and E-G map to channels 1 to 6 respectively. Finally, in 4-channel mode, channels B, D, F, and H are clock-gated, and channels A, C, E, and G map to channels 1 to 4 respectively. This channel mapping is summarized in Table 5.2. Reconfiguring the channel mode requires that the cut-off frequencies of the individual filters vary with the channel mode and they are also summarized in Table 5.2. The channel bandwidths for the 4- and 6-channel modes are based on [153], and bandwidths for the 8-channel mode is based on [134]. Regardless of the channel mode, the processor bandwidth covers 300Hz to 5.5kHz.

In order to achieve the required reconfigurability in the filter responses, the filter bank leverages three types of FIR filters with different levels of reconfigurability as shown in Figure 5-8. Because the filters have symmetric coefficients, the filter is folded to reduce the number of coefficient multiplications by half [154]. The coefficients are quantized to 8-bit precision which is the minimum possible without significantly affecting the frequency response, and the filter word lengths are optimized for the given filter coefficients. Furthermore, the filters are implemented in transposed form so that there is at most one adder in the critical path.

Fig. 5-8(a) shows the most reconfigurable filter (called FIR X3) which can be programmed to three different filter lengths: 14-, 16-, or 20-taps used in the 4-, 6-, or 8-channel mode using the control signals ($md8$, $md68$, $md4$) defined in Table 5.3 in the next section. This filter is used for channels A, C, E, and G which are active in

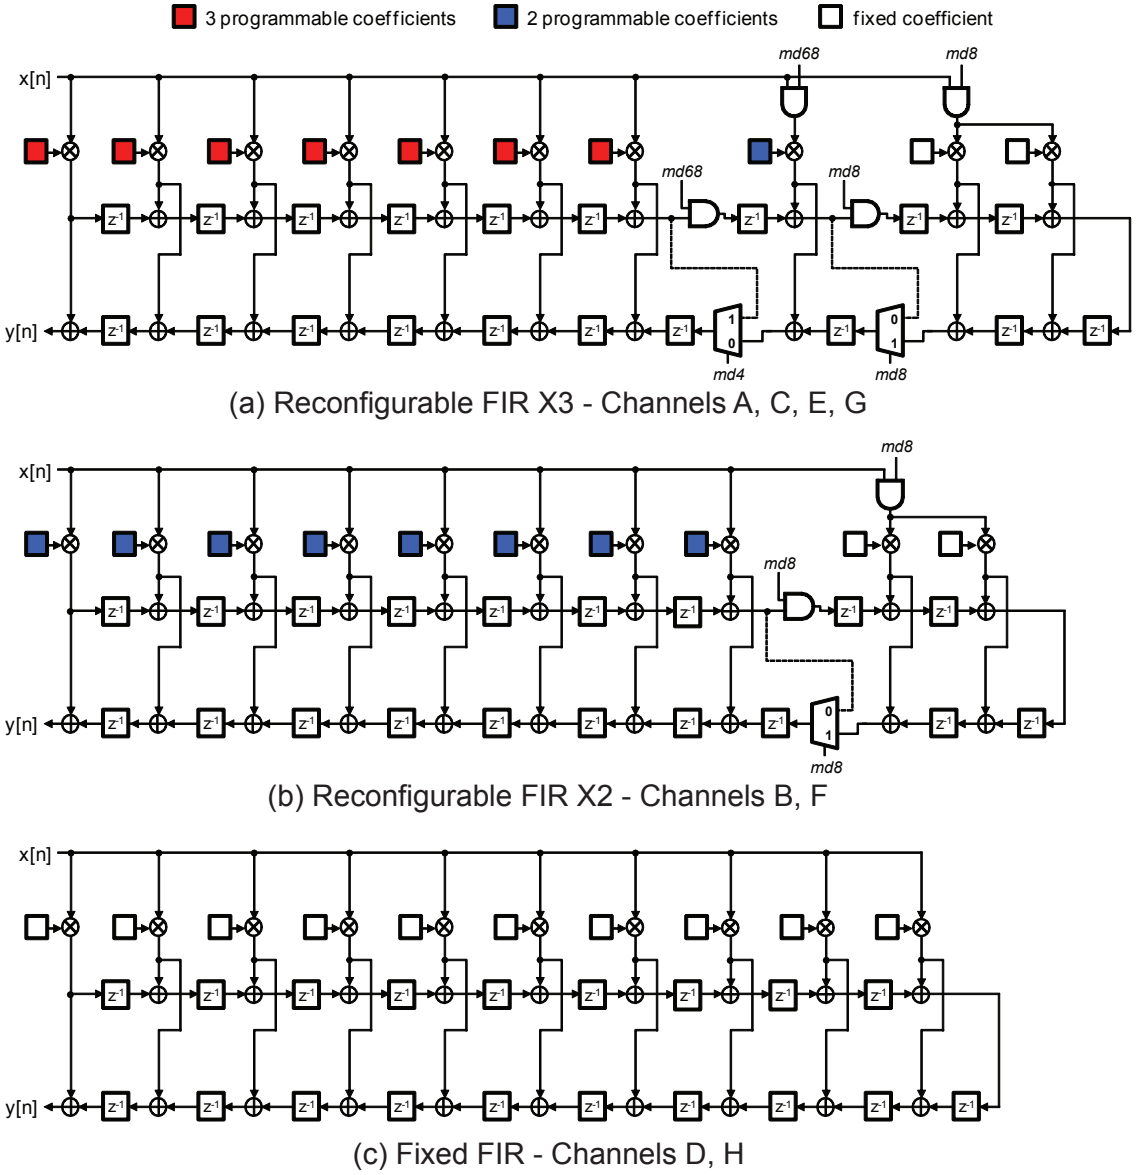

Figure 5-8: Three different FIR filters used in the reconfigurable filter bank shown in Figure 5-7. (a) Reconfigurable FIR filter that can be reconfigured into 3 modes: 14-, 16-, and 20-tap used in 4-, 6-, and 8-channel modes. (b) Reconfigurable FIR filter that can be reconfigured into 2 modes: 16- and 20-tap used in 6- and 8-channel modes. The filter is clock-gated when in 4-channel mode. (c) Fixed 20-tap FIR filter used in 8-channel mode only. The filter is clock-gated when in 6- and 4-channel modes. The reconfiguration logic signals are defined in Table 5.3.

Table 5.2: Reconfiguration of the multi-rate filter bank (shown in Figure 5-7) comprising channels A - H into 4-, 6- and 8-channel modes, where the filters are reconfigured into 14-, 16- and 20-tap FIR filters respectively. The filter cut-off frequencies are shown in parentheses, and the clock frequency for each filter is indicated by f_{clk} at the top of each column.

Mode	*sel_chan* [1:0]	Channel mapping							
		f_{clk}=2 kHz		f_{clk}=4 kHz		f_{clk}=8 kHz		f_{clk}=16 kHz	
		Ch A	Ch B	Ch C	Ch D	Ch E	Ch F	Ch G	Ch H
4-channel (14-tap FIR filter bank)	00	**1** (300-600Hz)	off	**2** (600-1300Hz)	off	**3** (1300-2600Hz)	off	**4** (2600-5500Hz)	off
6-channel (16-tap FIR filter bank)	01	**1** (300-500Hz)	**2** (500-800Hz)	**3** (800-1300Hz)	off	**4** (1300-2000Hz)	**5** (2000-3400Hz)	**6** (3400-5500Hz)	off
8-channel (20-tap FIR filter bank)	1x	**1** (300-425Hz)	**2** (425-625Hz)	**3** (625-900Hz)	**4** (900-1300Hz)	**5** (1300-1850Hz)	**6** (1850-2650Hz)	**7** (2650-3825Hz)	**8** (3825-5500Hz)
Filter structure		Reconf. FIR X3	Reconf. FIR X2	Reconf. FIR X3	Fixed	Reconf. FIR X3	Reconf. FIR X2	Reconf. FIR X3	Fixed

all three modes. Furthermore, the coefficients shown in red can be chosen from three values, the coefficients in blue can be chosen from 2 values, and fixed coefficients are shown in white. The number of taps in the filter increases with the number of channels because more selectivity is required. Similarly, Fig. 5-8(b) shows the reconfigurable FIR X2 filter used for channels B and F. It can be programmed to two different filter lengths: 16- or 20-taps used in the 6- or 8-channel modes only. FIR X2 is clock-gated when the processor is in the 4-channel mode. Finally, Fig. 5-8(c) shows the 20-tap FIR filter used for channels D and H which is fixed because it is used in the 8-channel mode only, and clock-gated during the 4- or 6-channel modes.

Effective Frequency Response

A simplified block diagram of the multi-rate filter bank is shown in Figure 5-9(a) where x[n] is the input from the ADC at a sampling rate of $f_S = 16$kHz. In order to determine the effective frequency response of the filter bank at $f_S = 16$kHz, the downsampling identity shown in Figure 5-10 can be used [80]. Based on the equivalence of the block

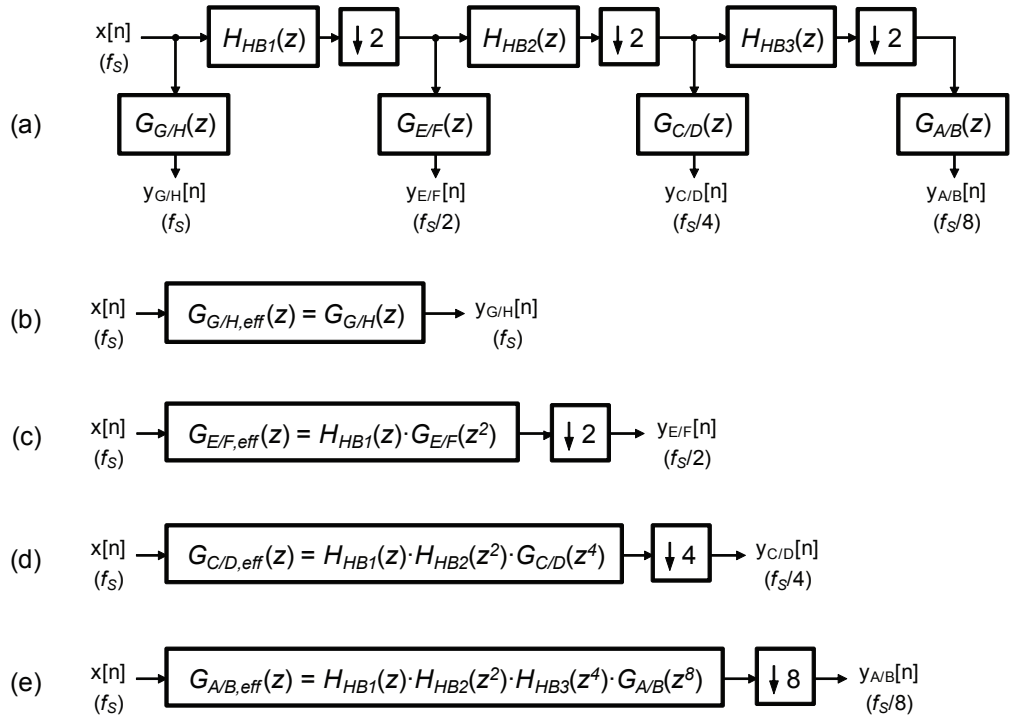

Figure 5-9: (a) Simplified block diagram of the multi-rate filter bank shown in Figure 5-7. The effective filters based on the downsampling identity for channels G/H, E/F, C/D, and A/B are shown in (b) - (e) respectively.

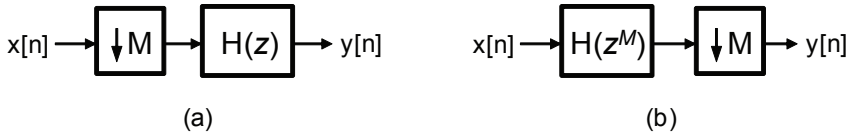

Figure 5-10: The systems in (a) and (b) can be shown to be equivalent [80]. This is known as the downsampling identity.

diagrams shown in Figures 5-10(a) and (b), the effective frequency response of each channel can be determined and they are shown in Figures 5-9(b)-(e) respectively.

Looking at channel A for example, its effective filter response is given by

$$G_{A,eff}(z) = H_{HB1}(z)H_{HB2}(z^2)H_{HB3}(z^4)G_A(z^8) \tag{5.14}$$

where $G_A(z)$ is the response of channel A at $f_S/8 = 2$kHz, and $H_{HB1}(z)$, $H_{HB2}(z)$, and $H_{HB3}(z)$ are the half-band filter responses of the 1st, 2nd, and 3rd stage of the decimation filter. Therefore, although $G_A(z)$ is only a 20-tap filter (in 8-channel

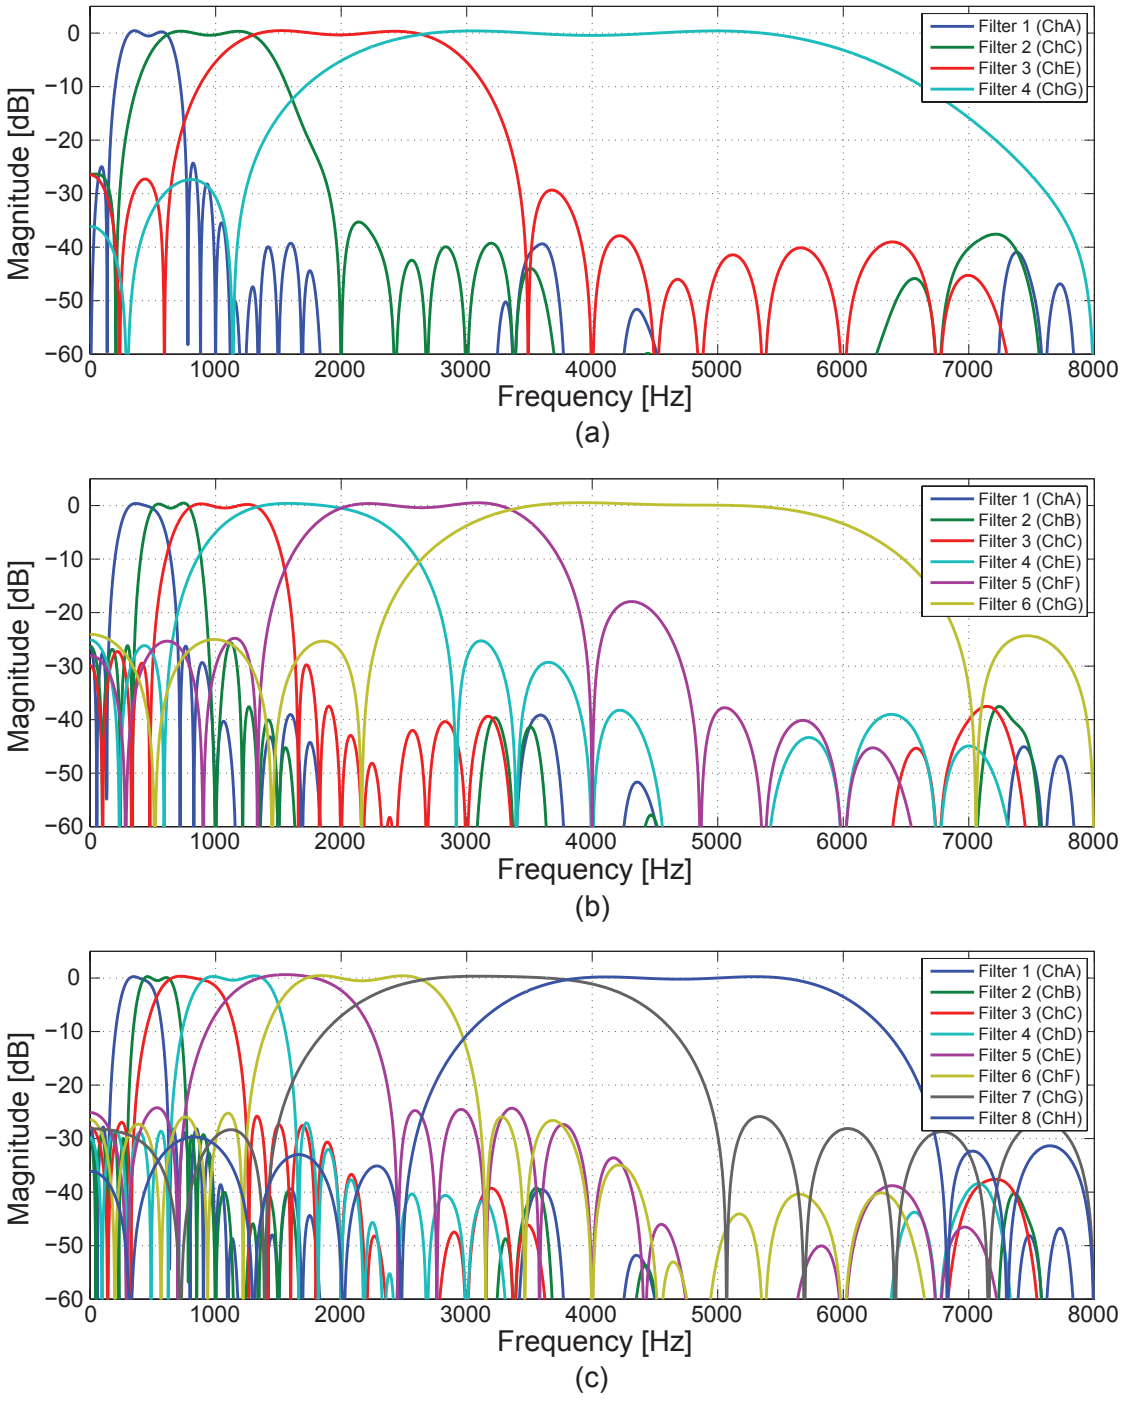

Figure 5-11: Effective frequency response of the multi-rate filter bank at 16kHz reconfigured in (a) 4-channel, (b) 6-channel, and (c) 8-channel modes.

mode), the effective filter is a much higher order filter because of the decimation filtering and multi-rate signal processing.

The effective frequency responses of the logarithmically spaced filter bank (at 16kHz) reconfigured in 4-, 6-, and 8-channel modes is shown in Figure 5-11(a)-(c) respectively. In this work, the stop band is limited to a moderate attenuation of approximately 30dB.

5.5 Arbitrary Waveform Stimulator

The output from each channel of the sound processor is the compressed temporal envelope of the signal in that particular frequency band, and it is used to determine the strength of the stimulus delivered to the corresponding electrode. This section presents details on the back-end of the SoC which is the arbitrary waveform neural stimulator. It is broken down into three sub-blocks: the high-voltage electrode switch matrix, the current source, and the low-voltage digital waveform interface. Many of the analog circuits presented in this section were implemented and simulated by Rui Jin, and their description is included here to facilitate understanding of the overall stimulator subsystem.

5.5.1 High-Voltage Electrode Switch Matrix

Since the CIS sound processing strategy relies on interleaved stimulation, a single current source is interleaved between all electrodes using the high-voltage switch matrix shown in Figure 5-12. E_i (for $i = 1$ to 8) are the active intracochlear electrodes and E_{com} is the common return electrode of a monopolar electrode array. A high-frequency R_sC_d electrode model between E_i and E_{com} models the impedance of the electrode-tissue-electrode interface. The switches S_C, S_A, S_{iC}, and S_{iA} are used to select the active electrode (E_i) and control the direction of current flow. Table 5.3 shows the channel mode selection logic used to determine which electrodes are active based on *sel_chan*[1:0]. Electrode E_i is active when S_i is asserted.

Fig. 5-12(a) shows the cathodic phase of a stimulation pulse through E_4 when S_C

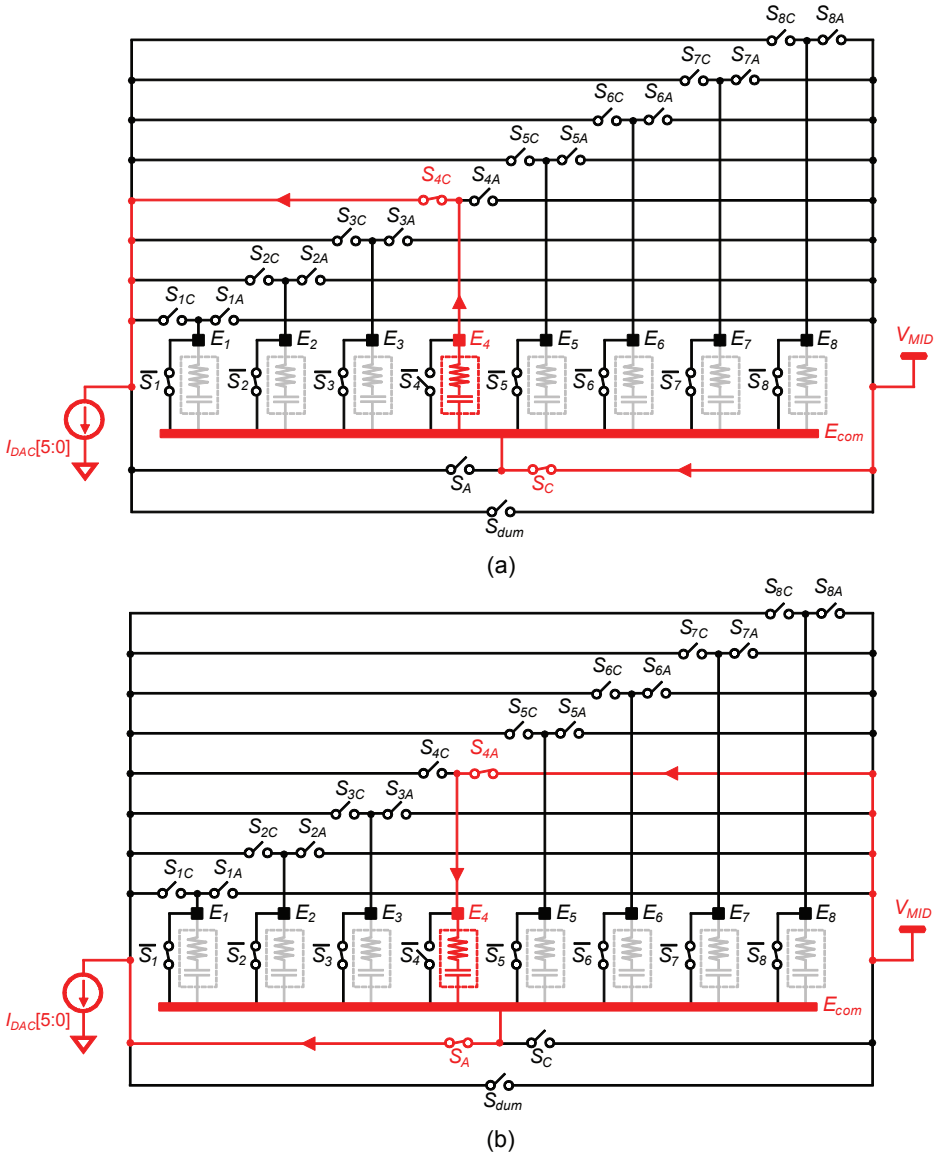

Figure 5-12: Schematic of the 8-channel high-voltage electrode switch matrix during (a) the cathodic phase, and (b) the anodic phase of electrode 4 (E_4). The path of current flow is highlighted in red, and the direction is indicated by the arrows.

Table 5.3: Channel mode selection logic used in the digital control and electrode switch matrix. *sel_chan*[1:0] is programmed through the serial programming interface.

Mode	*sel_chan*[1:0]	*md*4	*md*6	*md*8	*md*68	Active switches, S_i
4-channel	00	1	0	0	0	$i = 1, 3, 5, 7$
6-channel	01	0	1	0	1	$i = 1, 2, 3, 5, 6, 7$
8-channel	1x	0	0	1	1	$i = 1, 2, 3, 4, 5, 6, 7, 8$

and S_{4C} turn on. The current flows negatively through E_4, and the path and direction of the current flow is highlighted in red. The value of the current is determined by the 6-bit current source (I_{DAC}[5:0]) which will be described in Section 5.5.2. Fig. 5-12(b) shows the anodic phase following the cathodic phase, where S_C and S_{4C} turn off, and S_A and S_{4A} turn on which reverses the direction of current flow in the electrode. During each of the cathodic and anodic phases, an on-chip digital waveform controller controls the current source to provide any arbitrary waveform shape.

Since the switch matrix works like an H-bridge, current always flows from V_{MID}, which is a high-voltage supply (5V to 10V), to ground. In between the cathodic and anodic phases, an optional switch S_{dum} can be used to insert an inter-phase gap. After the completion of each pulse, S_i is de-asserted and therefore switch $\bar{S}_i$ is turned on to short the electrode to ensure that any residual charge is removed. Although not pictured, a large DC blocking capacitor (220nF) is also placed in series with the electrodes to ensure that there is no DC current flowing to the tissue for safety reasons.

Switch Matrix Timing Diagram

Figure 5-13 shows the timing diagram of the digital control for the electrode switch matrix in 8-channel mode over a single analysis period (i.e., the period over which each electrode gets stimulated once). The circuit block diagrams generating all of the signals in the timing diagram is provided later in Figure 5-16. The start of each stimulation cycle begins on the rising edge of ϕ_{LO} (1kHz) which generates a *stim_start* pulse and also asserts *en*33 which is used to enable the 3.3V supply from which the current source circuits operate.

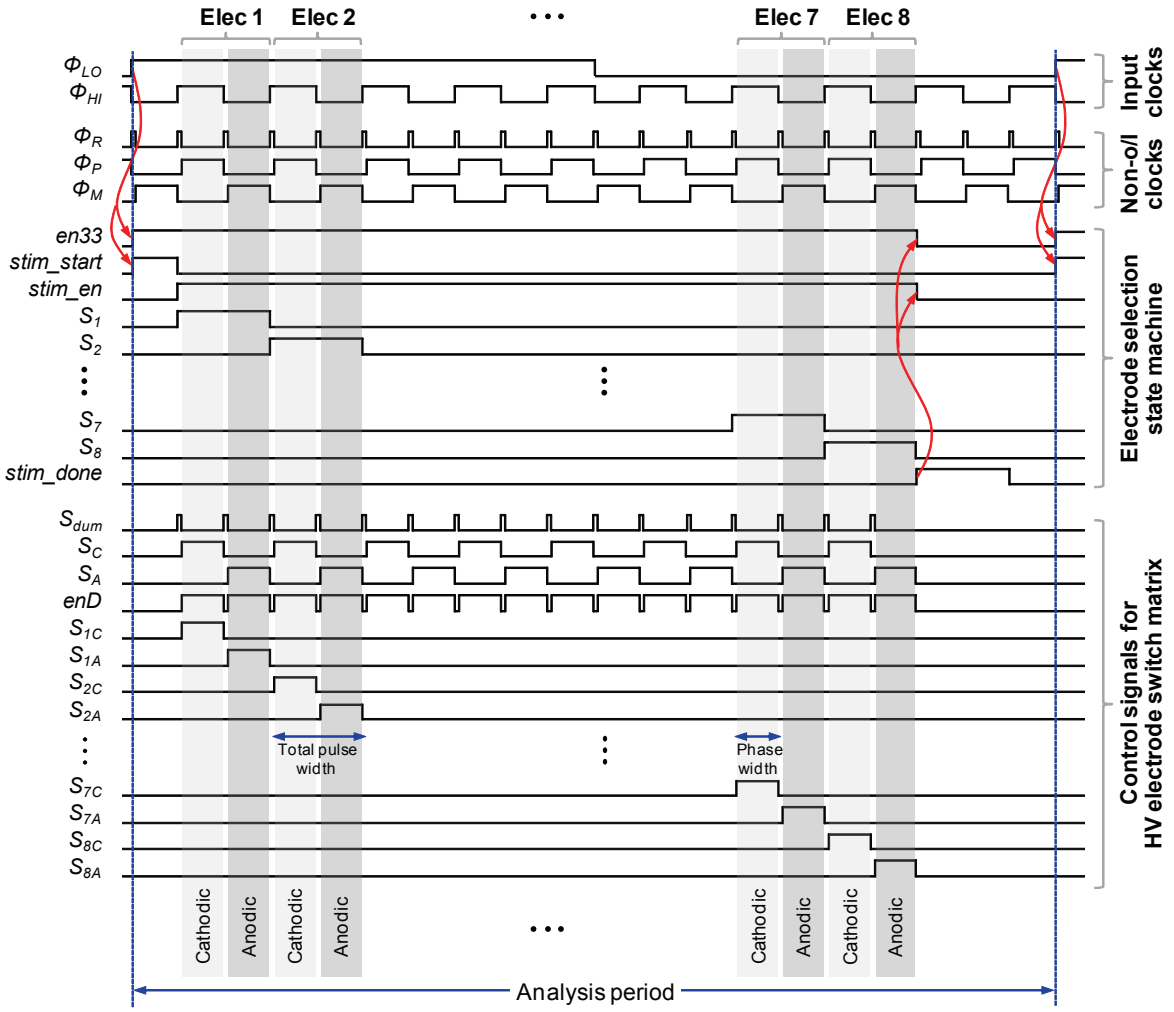

Figure 5-13: Timing diagram of the digital control of the electrode switch matrix in 8-channel mode over a single analysis period.

Stimulation is enabled when *stim_en* is asserted, and the electrode selection signals S_i are generated by an electrode selection state machine that is clocked by ϕ_{HI}. Note that *en*33 rises half a cycle before *stim_en* is asserted in order to provide adequate time for the current source circuits to settle after powering up. Non-overlapping clocks ϕ_R, ϕ_P, and ϕ_M are generated from ϕ_{HI}. The switch signals S_C, S_A, S_{iC}, S_{iA}, and S_{dum} are derived from gating ϕ_R, ϕ_P, and ϕ_M with the appropriate control signals (*stim_en* and S_i). Stimulation is complete on the positive edge of *stim_done*, which de-asserts *stim_en* and *en*33. Note that the stimulation phase width is governed by the frequency of ϕ_{HI} which is a variable frequency clock. The phase width is equal to $\frac{1}{2f_{\phi_{HI}}}$ and the total pulse width is twice that value.

Finally, *enD* is used to gate the digital input code to the current source to ensure that it is in a low-current state when disabled which helps with settling transients.

5.5.2 Current Source

For a comparison between a current source-based stimulator and a DC-DC converter-based stimulator, the reader can refer to Appendix B. The current source used in this work is based on the voltage-controlled resistor (VCR) topology from [155] which is chosen for its large output impedance and high voltage compliance. However, in this work, a current-steering DAC is designed to achieve fast settling in the current source. High output impedance is needed to ensure that the current delivered does not depend on the electrode impedance. Voltage compliance is defined as the range of the stimulator supply voltage (V_{MID}) for which the current remains within 1% of its nominal value, and therefore a high voltage compliance is needed for efficiency and also to prevent saturation of the current source when the electrode impedance and the stimulation current are high.

The schematic of the current source operating from a 3.3V supply is shown in Figure 5-14. It provides 6 bits of resolution which is typically sufficient for electrical hearing using cochlear implants [84]. The current source uses feedback to force a reference voltage V_{DSREF} across a resistor, shown in Fig. 5-14 as a voltage-controlled

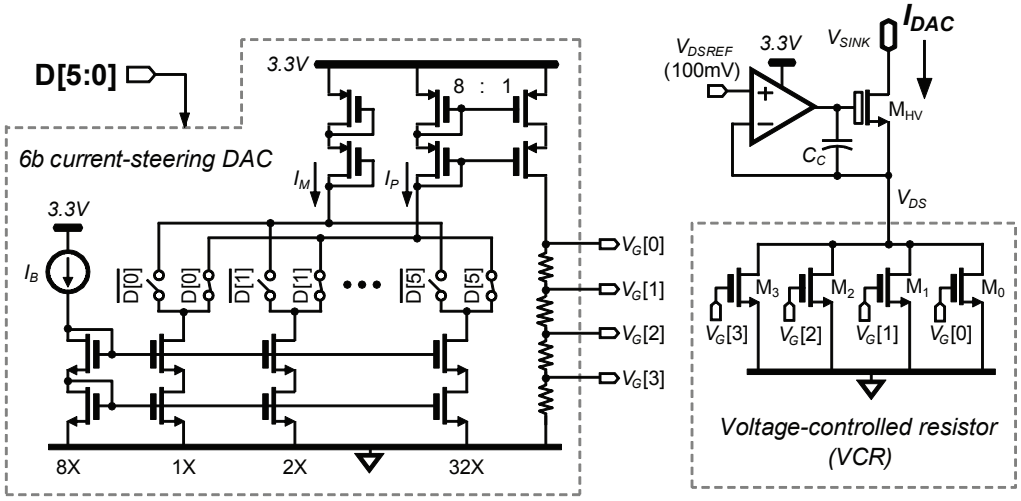

Figure 5-14: Schematic of the 6-bit current source with a fast-settling current-steering DAC and large voltage compliance based on [155]. All devices are 3.3V transistors except for M_{HV} which is a high-voltage NMOS transistor.

resistor with value R_{VCR}. Therefore, the output current is simply

$$I_{DAC} = \frac{V_{DS}}{R_{VCR}} = \frac{V_{DSREF}}{R_{VCR}} \tag{5.15}$$

where $V_{DS} = V_{DSREF}$ through the virtual short (as long as the loop gain is high). Either the voltage V_{DSREF} or resistance R_{VCR} can be varied to change I_{DAC}. However, varying V_{DSREF} would change the voltage compliance of the circuit and therefore R_{VCR} is used as the variable element. The VCR comprises a bank of NMOS transistors that act together to provide a linear resistance. As the name suggests, the value of R_{VCR} is a function of the gate voltages V_G[3:0] which are generated with the current-steering DAC shown on the left of the figure. M_0 is the main triode device that carries the majority of I_{DAC}, and it is controlled by $V_G[0]$. Considering M_0 only, its resistance in the linear region is given by

$$R_{VCR,M0} = \left[\frac{\mathrm{d}i_{ds}}{\mathrm{d}v_{gs}}\right]^{-1} \approx \frac{1}{\mu C_{ox}\frac{W}{L}(V_{gs} - V_t)} = \frac{1}{K(V_G[0] - V_t)} \tag{5.16}$$

where $K = \mu C_{ox}\frac{W}{L}$ and V_t is the threshold voltage. Therefore, by combining Equations 5.15 and 5.16, we see that $I_{DAC} = V_{DSREF}K(V_G[0] - Vt)$. Therefore, as long

as $V_G[0]$ is linear with the DAC code, I_{DAC} will be linear as well. M_1 to M_3 are auxiliary devices controlled by $V_G[1:3]$ (which are a fraction of $V_G[0]$) that carry a small fraction of the total current to help linearize M_0. V_{DSREF} is set to 100mV to ensure that M_0 is in triode, and to limit the headroom of the current source to achieve high compliance. The output device (M_{HV}) is the only high-voltage device in the current source because it connects to the high-voltage switch matrix that drives the electrodes. It must be kept in saturation to keep the loop gain of the feedback loop high, and requires approximately 100mV of drain-source voltage. Therefore, the output node V_{SINK} can be as low as 200mV which maximizes the voltage compliance of the circuit (e.g., 6.8V of compliance is achieved when $V_{MID} = 7$V).

The DAC in Figure 5-14 leverages current steering to achieve very fast settling which is needed to generate arbitrary waveforms with a short time step (e.g., a phase width of 25μs with 10 steps/phase requires a time step of 2.5μs). The input code D[5:0] from a digital arbitrary waveform controller steers binary-weighted currents to I_P which is then mirrored to the output branch. A resistor string converts that code-dependent current into the control voltages $V_G[3:0]$ which are linear with D[5:0] as desired. The bias current I_B is set such that the full scale of I_{DAC} is 1mA which is needed for cochlear implant applications.

5.5.3 Digital Waveform Interface

The previous two sections described the core analog circuits of the stimulator. That is, the high-voltage electrode switch matrix used to select the active electrode and control the direction of the stimulation current, which is generated by a fast-settling 6-bit current source with high output impedance and voltage compliance. In this section, the digital control for the stimulator is discussed. In order to minimize power consumption, the digital controller operates at 0.6V. Therefore, a two-stage level shifter shown in Figure 5-15 is used between the ultra-low-voltage digital domain and the high-voltage domain of the switch matrix. An intermediate voltage of 1.8V is used in the level shifter between 0.6V and V_{DDG} which can be 7V to 12V.

Figure 5-16 shows the block diagrams of the digital control circuits for the stimu-

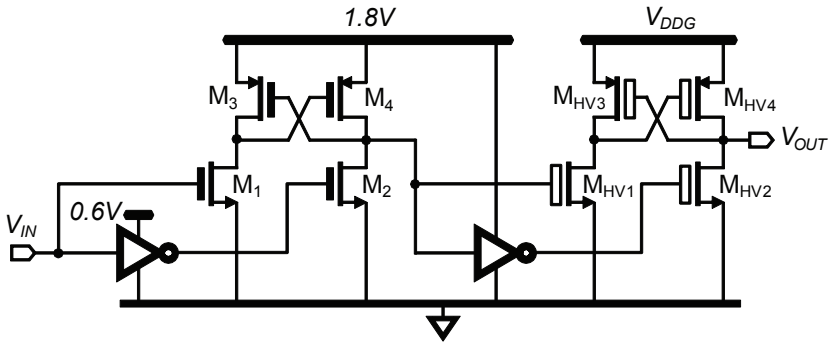

Figure 5-15: Schematic of the two-stage level shifter from 0.6V to 1.8V to V_{DDG} used to interface between the ultra-low-voltage digital control and the high-voltage electrode switch matrix. High-voltage devices are indicated with a hollow gate symbol.

lator. Fig. 5-16(a) shows the custom-designed electrode selection state machine that generates S_1 to S_8 which are asserted in sequence to signify the active electrode (E_1 to E_8). Control signals $md68$ and $md8$ (see Table 5.3) are used to reconfigure the state machine between 4-, 6-, and 8-channel modes. The state machine is triggered with a positive edge of ϕ_{LO}, and the signals S_i are shifted out serially with $\phi_{HI,D}$ which is a delayed version of ϕ_{HI}. This is to ensure that S_i transitions just after the the negative edges of S_A and S_{iA} to avoid glitching in the switches. Other control signals (*stim_en*, *stim_start*, $en33$, and *stim_done*) that govern the stimulation cycle are also generated by this state machine (see Figure 5-13).

Fig. 5-16(b) shows the generation of the control signals for the switch matrix. A non-overlapping clock generator generates ϕ_R, ϕ_P, and ϕ_M which are shown in Figure 5-13. The control signals for the switch matrix are generated by gating S_i and *stim_en* with the non-overlapping clocks. The width of the inter-phase gap (i.e., the width of ϕ_R) is set by *sel_ipg*[1:0] which control the timing of one-shot circuits.

Finally, Fig. 5-16(c) shows the digital arbitrary waveform interface that controls the shape of the pulses delivered to the electrodes. The corresponding timing diagram over two stimulation pulses is shown in Figure 5-17. An example using electrode 1 (E_1) will help illustrate the operation of the control. In the timing diagram shown in Figure 5-17, the waveforms for E_1 are shown in red. While S_1 is asserted, the

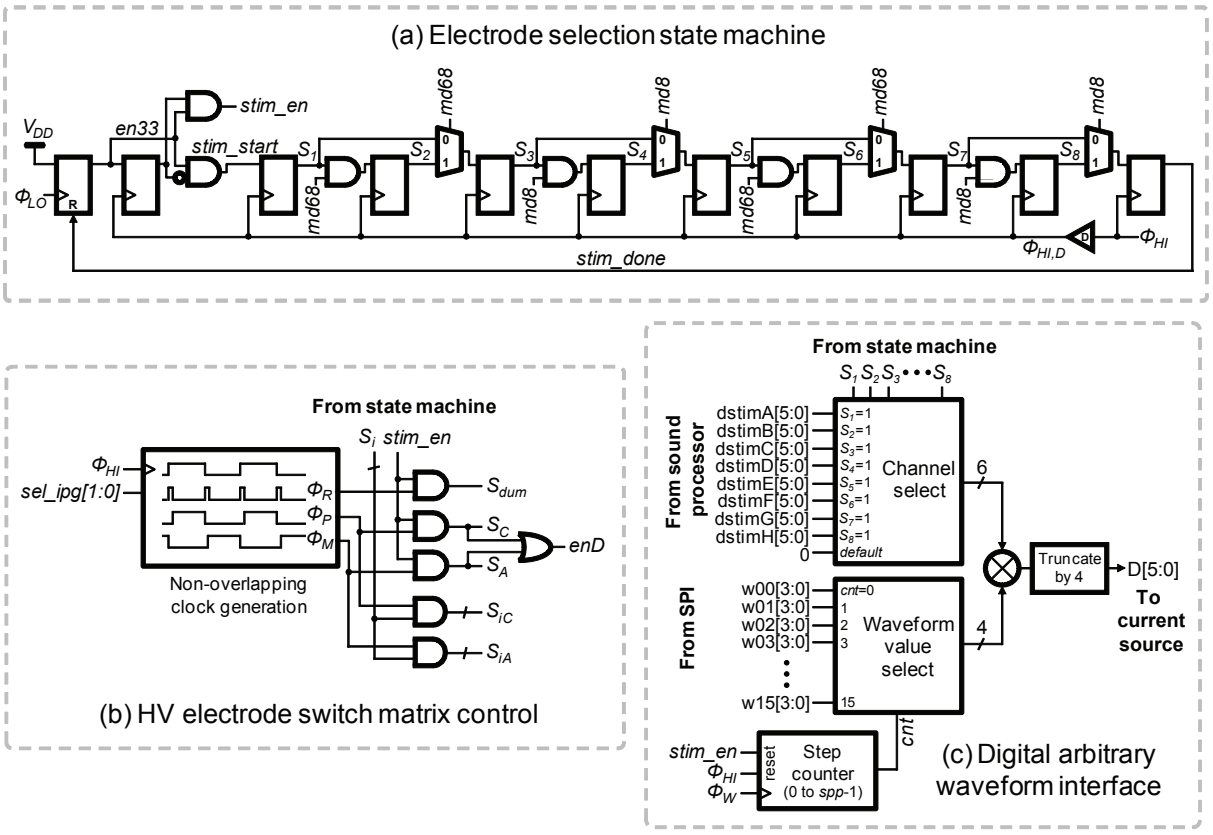

Figure 5-16: Ultra-low-voltage (0.6V) digital control of the stimulator. (a) Electrode selection state machine, where $S_1 - S_8$ determine the active electrode in the array. Control signals $md68$ and $md8$ are defined in Table 5.3, where $md68$ is asserted for 6- and 8-channel modes, and $md8$ is asserted for 8-channel mode only. (b) High-voltage switch matrix control generation. S_i is determined by the channel modes listed in Table 5.3, and sel_ipg[1:0] sets the duration of the inter-phase gap. (c) Digital arbitrary waveform interface for the stimulator. Buses w00[3:0] to w15[3:0] determine the waveform shape, and dstimA[5:0] to dstimH[5:0] are the outputs of each channel from the sound processor used to modulate the strength of the waveform.

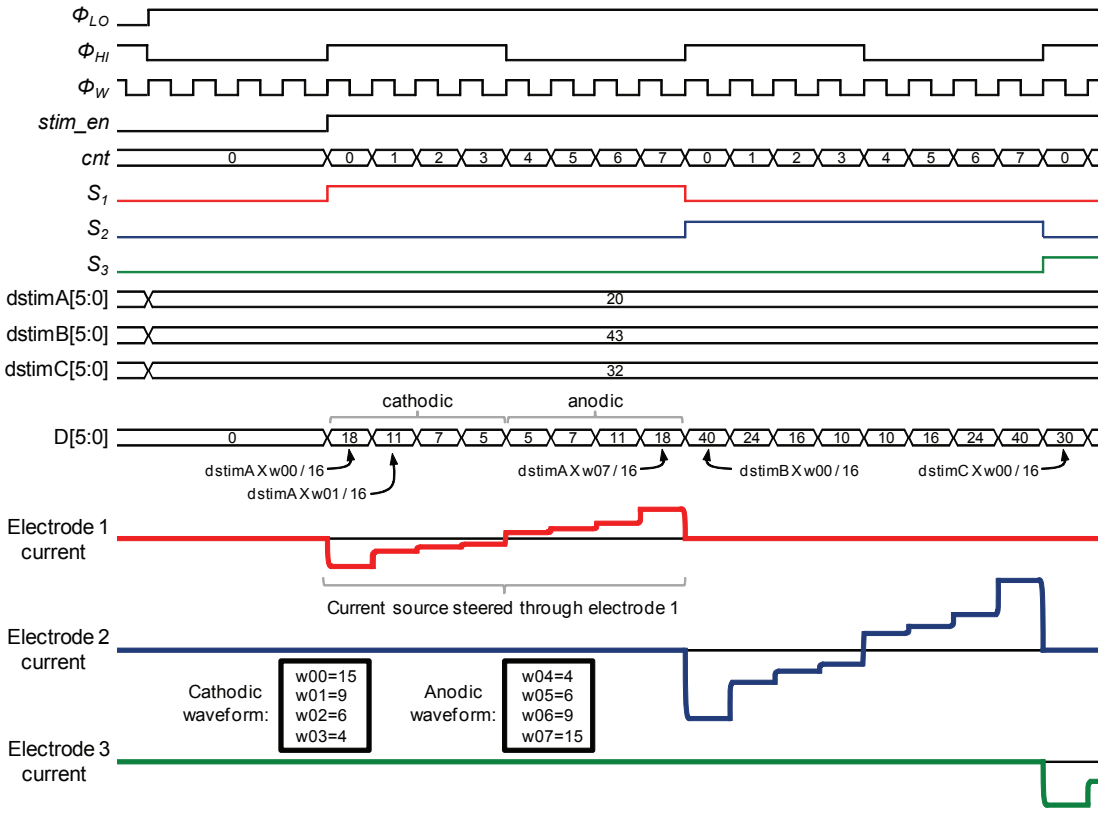

Figure 5-17: Timing diagram of the digital arbitrary waveform interface shown in Figure 5-16(c). This example shows 8 steps per pulse (4 per phase), but the interface supports up to 16 steps per pulse (8 per phase). Electrode 2 receives more current than electrode 1 because channel B (dstimB[5:0]) from the sound processor has a larger value than channel A (dstimA[5:0]). All values shown are in decimal format.

channel select block selects dstimA[5:0] from the sound processor. At the same time, a step counter clocked on ϕ_W (a high-frequency waveform clock) counts from *cnt* = 0 to *spp* − 1, were *spp* is the number of steps per pulse. In this example, *spp* = 8 for simplicity, but the actual implementation allows for up to 16 steps/pulse (i.e., 8 steps/phase). The value of *cnt* determines which waveform value (w00[3:0] to w15[3:0]) is multiplied with dstimA[5:0] to generate the output D[5:0] which is the input to the current source. The waveform values (w00 to w15) are programmed onto the SoC through the serial programming interface. In this way, the output of the sound processor (dstimA[5:0]) modulates the height of the pulse, which has an arbitrary waveform shape determined entirely by w00 to w15. In the example shown in Figure 5-17, w00 to w03 specify the shape of the cathodic phase, while w04 to w07 specify the shape of the anodic phase (w08 to w15 are unused). The step counter is reset after each ϕ_{HI} period, and the process repeats for the next electrode. In the example, since dstimB[5:0] = 43 is larger than dstimA[5:0] = 20, the peak of the current pulse through E_2 is higher. Finally, the frequency of ϕ_{HI} and ϕ_W can both be tuned to vary the phase width and number of steps/phase respectively.

Although the waveform values (w00 to w15) can take on any arbitrary value, it is important to ensure that the charge delivered during the cathodic phase is as close as possible to the charge delivered during the anodic recovery phase. In practice, this limits the flexibility of the waveform shape. However, the stimulator in this work uses DC blocking capacitors and a shorting phase to ensure that no DC current is delivered to the tissue.

5.6 Prototype Measurement Results

A prototype test chip of the fully-implantable CI SoC was implemented in a 0.18μm high-voltage CMOS technology, and the die photo is shown in Figure 5-18. The entire chip including pads measures 3.6mm × 3.6mm, while the active die area is 3.36mm^2. The chip was packaged in a 128-lead thin QFP package and the PCB used for chip characterization is shown in Figure 5-19. Models of the piezoelectric sensor

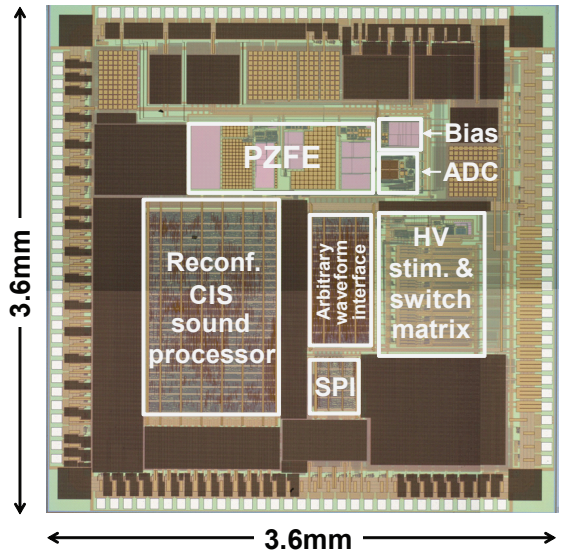

Figure 5-18: Die photo of the fully-implantable CI SoC prototype implemented in a 0.18μm high-voltage CMOS technology.

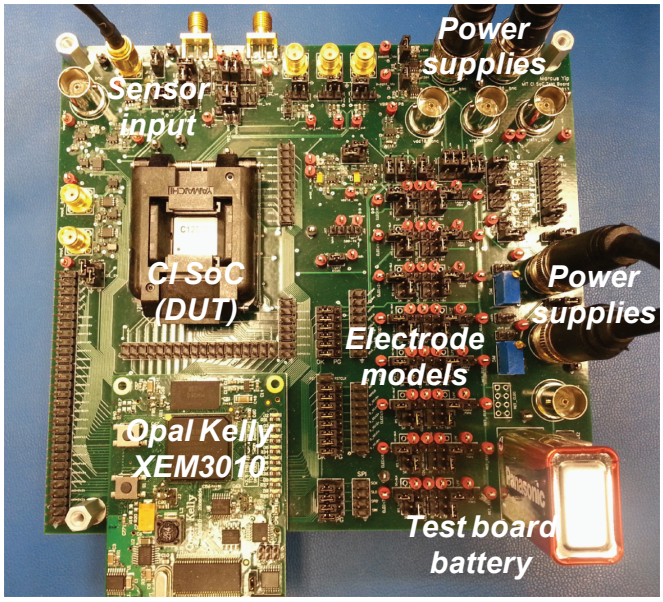

Figure 5-19: Printed circuit board for CI SoC testing and characterization.

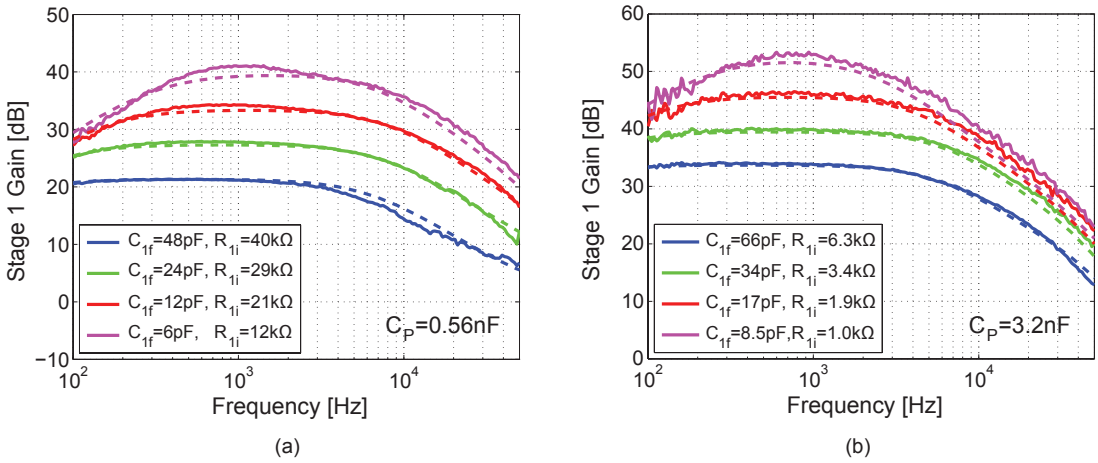

Figure 5-20: Measured gain response of stage 1 of the piezoelectric sensor front-end with (a) C_P=0.56nF, and (b) C_P=3.2nF for various gain settings. Simulation results are shown with dotted lines.

and stimulation electrodes were included on the PCB, and an Opal Kelly XEM3010 development board was used to provide communication between the prototype and a Matlab API on a laptop computer.

5.6.1 Piezoelectric Sensor Interface

The frequency response of the PZFE was measured with a 35670A Dynamic Signal Analyzer (Agilent Technologies). Outputs from the PZFE were buffered with high input-impedance op-amps (Analog Devices AD8605). A 16-bit 100kS/s ADC (Analog Devices AD7683) was also used to digitize the PZFE outputs for noise measurements. All measurements of the PZFE were made with an analog supply voltage of 1.5V.

Frequency Response

Figure 5-20 shows the measured frequency response of stage 1 of the PZFE for C_P = 0.56nF and 3.2nF which span the expected values for reasonable sizes of the piezoelectric sensor. For C_P = 0.56nF, the gain can be set from 21.2 to 41dB, while for C_P = 3.2nF, the gain can be set from 33.8 to 53dB by varying C_{1f}. At the lowest value for C_{1f} (i.e., highest gain setting), the high-pass corner frequency is 300Hz which is the lowest frequency processed by the CI sound processor. R_{1i} can be set to provide

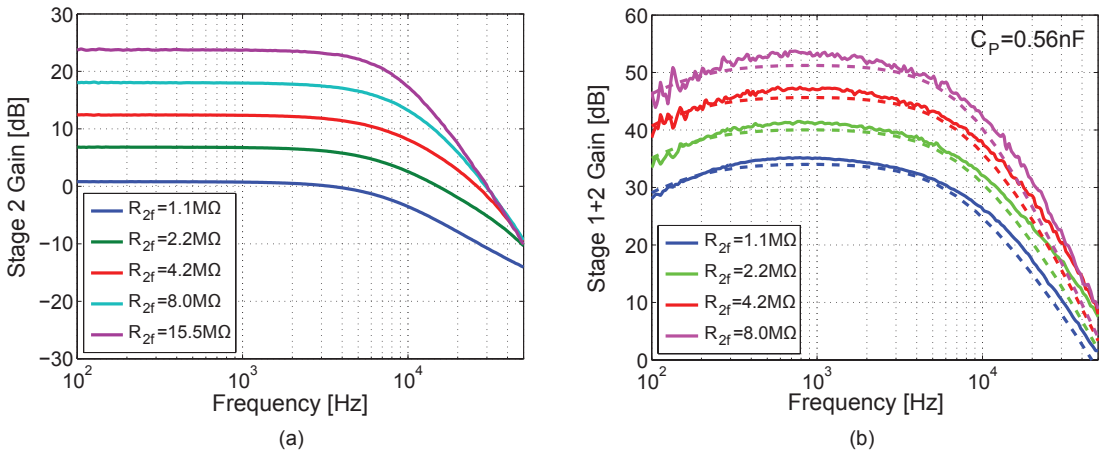

Figure 5-21: Measured gain response of (a) stage 2, and (b) stage 1 and 2 (C_P=0.56nF) of the piezoelectric sensor front-end. Simulation results are shown with dotted lines.

the desired low-pass cut-off. Simulation results are shown with dotted lines, showing good agreement with measured results.

Figure 5-21(a) shows the frequency response of stage 2 which provides additional programmable gain in 6dB steps by setting R_{2f}. Figure 5-21(b) shows the combined response of stage 1 and 2 for $C_P = 0.56$nF, $C_{1f} = 12$pF, $R_{1i} = 21$kΩ and various R_{2f} settings.

Input-Referred Noise

The overall noise performance of the PZFE is dominated by the noise of stage 1. Figure 5-22(a) shows the measured input-referred noise density for stage 1 with C_P = 3.2nF for various gain settings with a total current of $I_{STG1} = 4.5\mu$A. Small spurs are present in the spectrum at 300Hz, 420Hz, and 540Hz which are the 5$^{\text{th}}$, 7$^{\text{th}}$, and 9$^{\text{th}}$ harmonics of the 60Hz power-line noise. At a nominal gain setting of 40dB (C_{1f} = 34pF), the total integrated noise over the bandwidth of the CI sound processor from 300Hz to 5.5kHz is 1.93μV_{rms} which is less than the minimum expected signal at 40dB SPL as discussed in Section 4.1.3. It also shows good agreement with the noise analysis and simulation results presented in Section 5.3.1.

For the noise efficiency factor (NEF) calculation of the stage 1 charge amplifier [81],

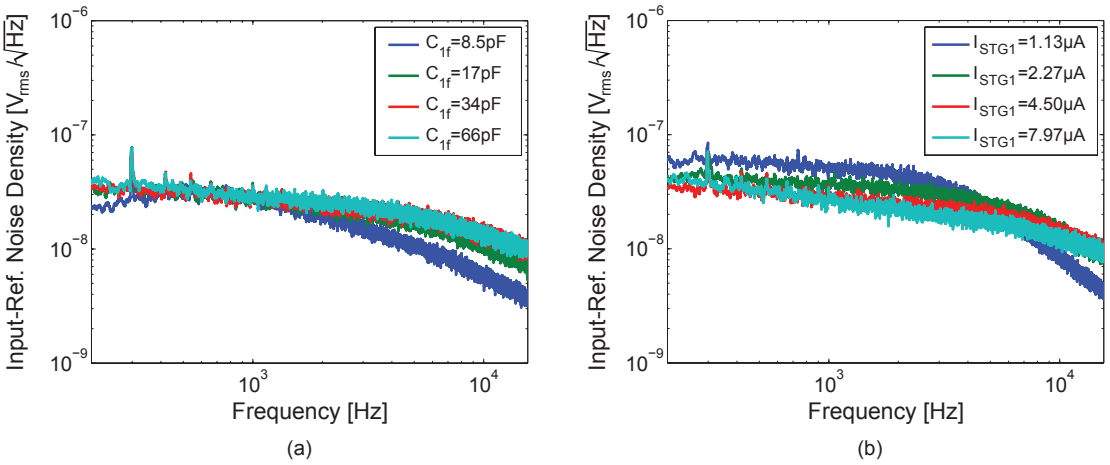

Figure 5-22: Measured input-referred noise density of stage 1 for C_P=3.2nF with (a) various gain settings, and (b) C_{1f}=34pF and various bias currents.

the total integrated noise from 300Hz to 25kHz is 2.65μV$_{\text{rms}}$, corresponding to a NEF of 2.68 (based on a 6.5kHz 3dB bandwidth). When factoring in the supply voltage of 1.5V, the PEF is 10.7 [26].

The noise within the signal bandwidth is dominated by thermal noise which can be lowered by increasing the total current in the stage 1 op-amp. Figure 5-22(b) shows the input-referred noise density when the total current is varied from 1.13μA to 7.97μA for C_{1f} = 34pF. When I_{STG1} is lowered to 1.13μA, the total integrated noise increases to 3.16μV$_{\text{rms}}$, whereas when I_{STG1} is set to its maximum value of 7.97μA, the total integrated noise is reduced to 1.76μV$_{\text{rms}}$.

The same set of measurements was repeated with C_P = 0.56nF and the results are summarized in Table 5.5 at the end of this chapter.

ADC

Since the ADC used in this work was ported from the ADC of the MSFE system for ECG monitoring described in Chapter 2, the measurement results are similar to those found in Section 2.4.5. The measured INL and DNL are -0.65LSB/+0.69LSB and -0.64LSB/+0.59LSB respectively at a sampling rate of 16kS/s, and the measured SNDR is 52dB for an ENOB of 8.35 bits.

THD and PSRR

Measured results show that the total harmonic distortion (THD) of the front-end is limited by the ADC. For a maximum THD of 1% (−40dB), the system can tolerate a maximum input signal of $V_P = 10\text{mV}_{\text{p-p}}$ for $C_P = 0.56\text{nF}$ at a nominal gain setting ($A_{STG1} = 33.4\text{dB}$, $A_{STG2} = 0.8\text{dB}$). Similarly, the maximum input signal is $V_P = 4.7\text{mV}_{\text{p-p}}$ for $C_P = 3.2\text{nF}$ at a nominal gain setting ($A_{STG1} = 39.5\text{dB}$, $A_{STG2} = 0.8\text{dB}$). This corresponds to an input dynamic range of 62dB and 58.7dB respectively for $C_P = 0.56\text{nF}$ and 3.2nF.

The PSRR of stage 1 was also measured by using the 35670A Dynamic Signal Analyzer to generate a $400\text{mV}_{\text{p-p}}$ ripple on the power supply. A PSRR of 46dB and 59dB is measured for $C_P = 0.56\text{nF}$ and 3.2nF respectively. This is important for a system such as a CI which may be in environments with large amounts of RF noise.

5.6.2 Reconfigurable Sound Processor

The reconfigurable CIS sound processor was tested by using an arbitrary function generator (AFG3102 from Tektronix) to generate arbitrary test signals at the input of the ADC. The outputs of the sound processor (dstimA[5:0] through dstimH[5:0]) were recorded on a laptop computer using the Opal Kelly/Matlab interface. Since the output of the sound processor is logarithmically compressed, spectrograms were reconstructed in Matlab by exponentiating the data. All measurements with the ADC and sound processor were made at a digital supply voltage of 0.6V.

To demonstrate the the reconfigurability in the number of channels of the processor, a logarithmic chirp signal was input to the ADC. Figure 5-23(a) shows the measured spectrogram at the output of the ADC, and panels (b), (c), and (d) show the measured spectrogram of the processor configured in 4-, 6-, and 8-channel modes. Since the processor uses a logarithmically spaced filter bank, its spectrogram looks linear as expected. A Matlab simulation of the 8-channel processor is shown in Fig. 5-23(e), showing good agreement with the measured result in (d).

The sound processor also has patient-fitting knobs like adjustable threshold (THR)

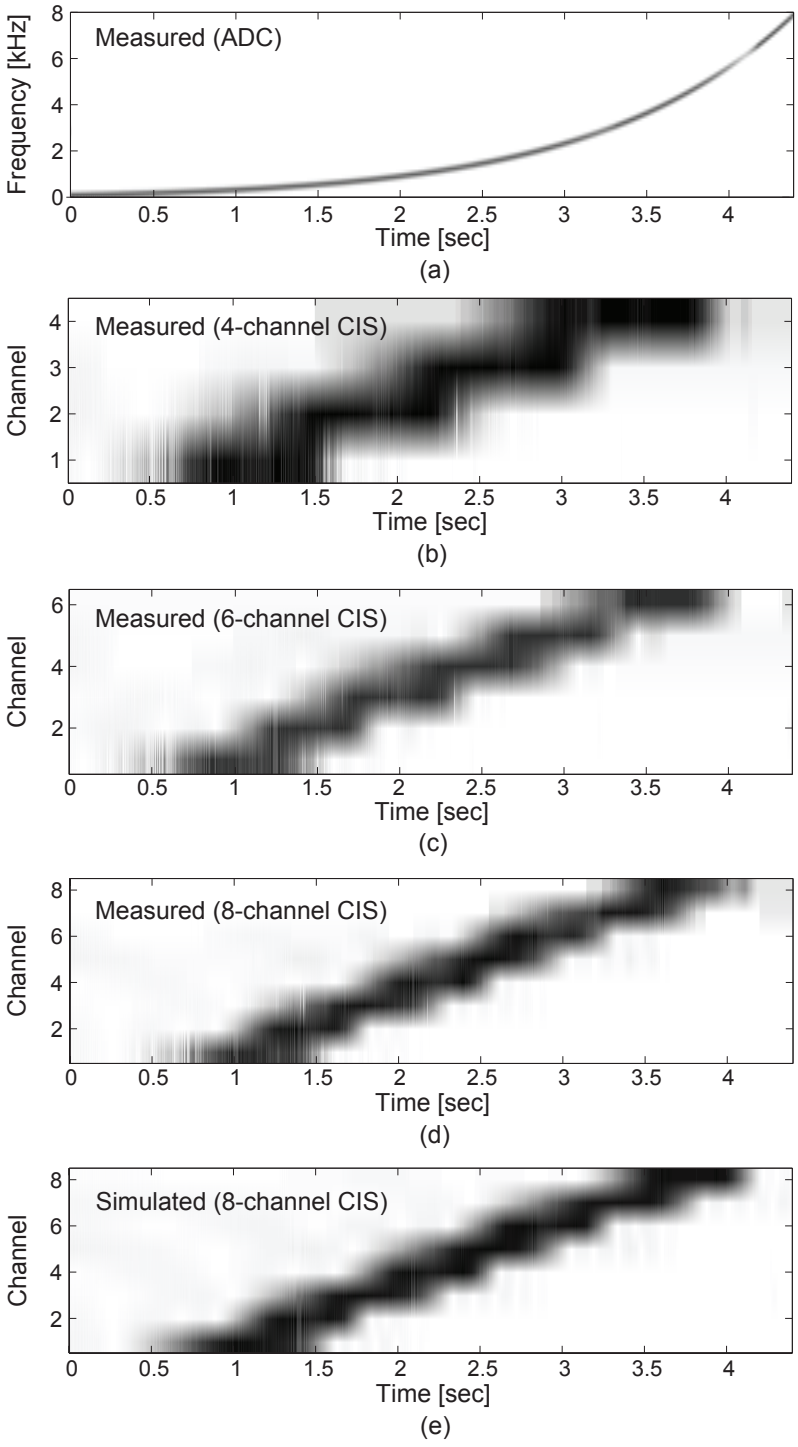

Figure 5-23: Measured spectrograms at the output of the (a) ADC, (b) 4-channel sound processor, (c) 6-channel sound processor, and (d) 8-channel sound processor when a logarithmic chirp signal is applied at the input. (e) Ideal Matlab simulation output to compare against the measured results shown in (d).

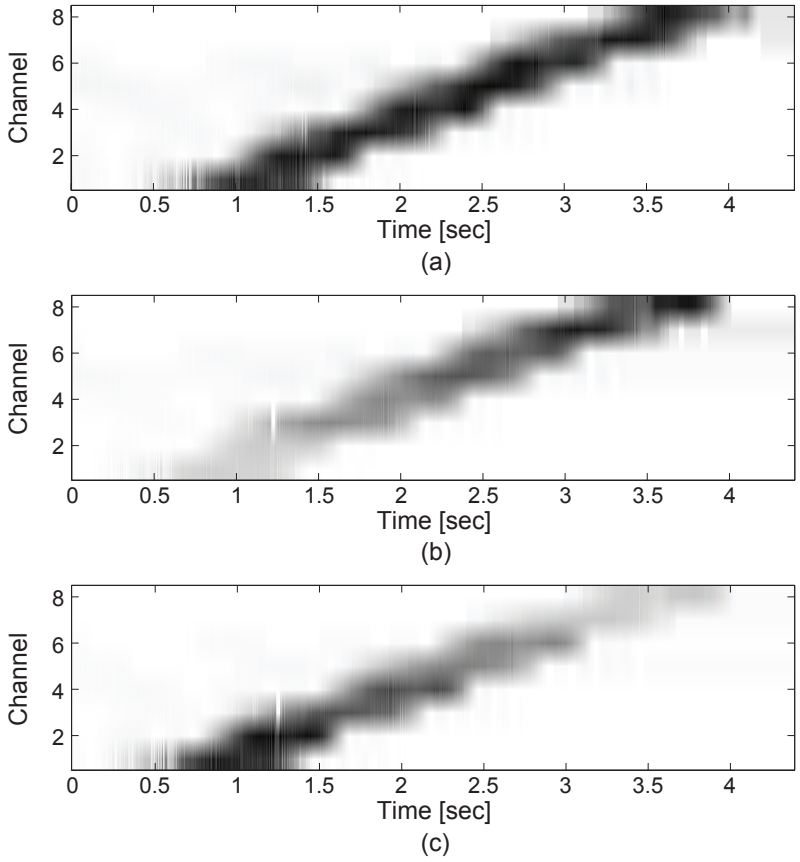

Figure 5-24: Measured spectrograms of a chirp at the output of the sound processor (8-channel mode) with (a) equal THR/MCL among all channels, (b) emphasis on high-frequency channels, and (c) emphasis on low-frequency channels.

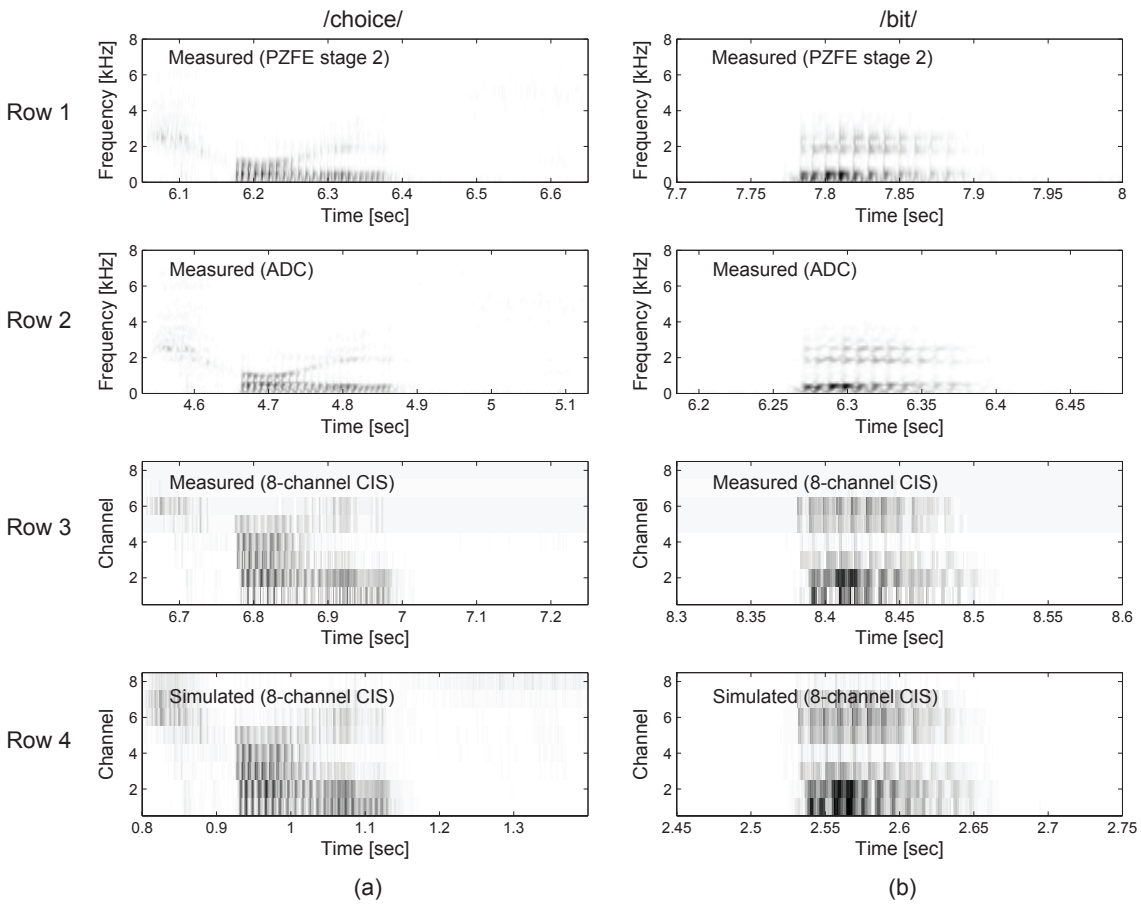

Figure 5-25: Measured spectrograms of the words (a) /choice/, and (b) /bit/ at the output of the piezoelectric sensor front-end (row 1), ADC (row 2), and 8-channel sound processor (row 3). The ideal Matlab simulation of the 8-channel processor is shown in row 4 for comparison.

and most-comfortable-level (MCL) which can be set individually for each channel. Figure 5-24 shows measured chirp spectrograms of the processor in 8-channel mode with (a) equal THR/MCL settings in all channels, (b) emphasis on high-frequency (basal) channels, and (c) emphasis on low-frequency (apical) channels.

Aside from the test chirp signal, the words "choice" and "bit" were generated at the input of the PZFE using an arbitrary function generator to demonstrate the operation of the full signal chain from the PZFE to the sound processor with speech signals. Figure 5-25 shows measured spectrograms of the words (a) "choice" and (b) "bit" at the output of the PZFE (row 1), ADC (row 2), and sound processor in 8-channel mode (row 3). A Matlab simulation of the processor is provided in row 4 for comparison. Note that the measured sound processor spectrogram also includes any non-idealities and bandwidth limitations of the PZFE and ADC.

5.6.3 Arbitrary Waveform Stimulator

This section presents measurement results of the arbitrary waveform stimulator and high-voltage electrode switch matrix.

Current Source

The core of the stimulator is the 6-bit current source which is required to have high output impedance and wide voltage compliance. Figure 5-26(a) shows the measured output current (I_{DAC}) versus V_{SINK} from 0 to 8V. The full-scale current at a digital input code of 63 is 1mA as designed. The output impedance of the current source is given by $\left(\frac{\mathrm{d}I_{DAC}}{\mathrm{d}V_{SINK}}\right)^{-1}$. The actual output impedance of the current source could not be measured beyond 20MΩ because of the limited accuracy of the Keithley 2400 SourceMeter used to measure the current. However, simulations show that the output impedance is 300MΩ which is well beyond the requirement.

Fig. 5-26(b) is a zoomed-in version of the plot in Fig. 5-26(a) which shows the headroom required in the current source. The voltage compliance is defined as the range of the supply voltage for which the current remains within 1% of its value at

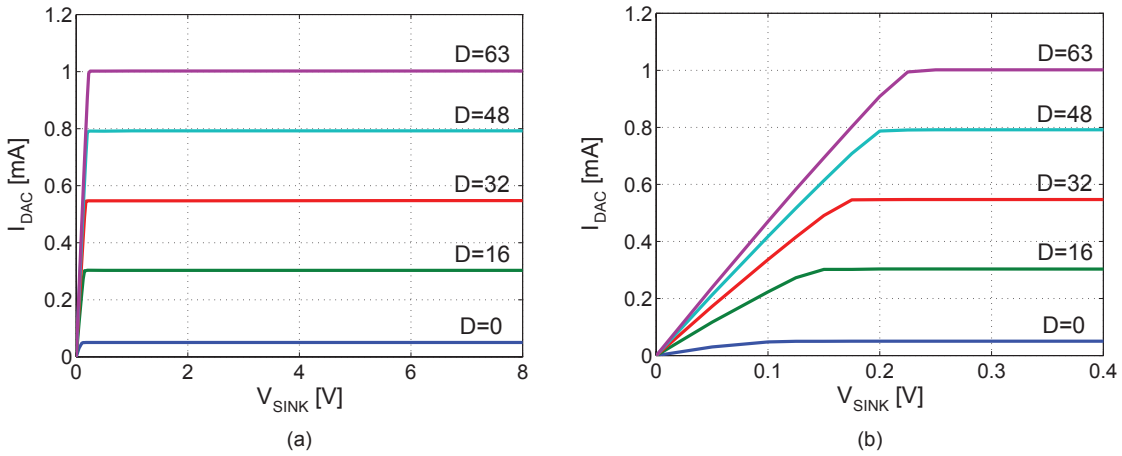

Figure 5-26: (a) Measured output current versus V_{SINK} of the stimulator current source for various input codes. (b) Zoomed-in version of the plot in (a) illustrating the large voltage compliance.

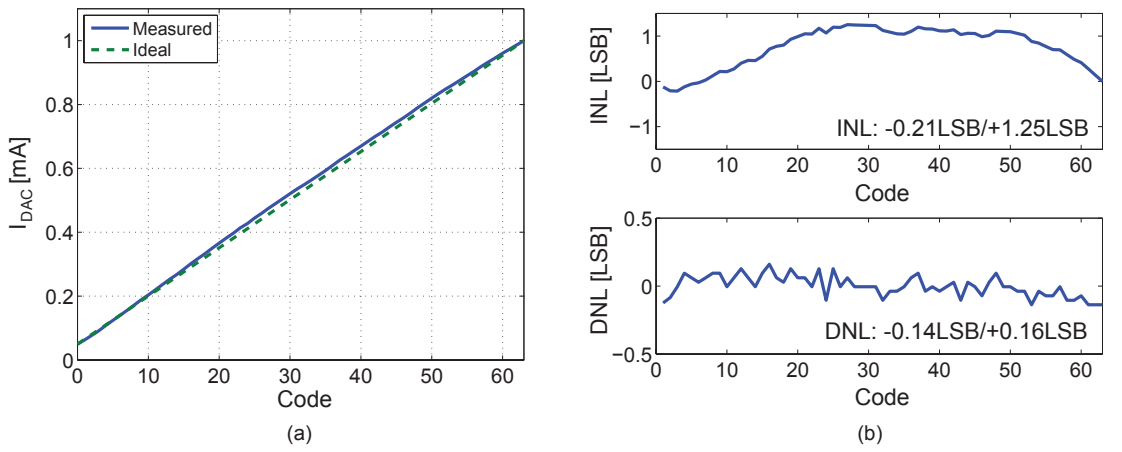

Figure 5-27: (a) Measured output current versus input DAC code for the stimulator current source. (b) Static linearity of the transfer curve in (a).

the nominal (high) end of the supply voltage [155]. The worst case compliance occurs at the full-scale of 1mA (D=63) where a headroom of 0.22V is required. For V_{MID} = 7V, this translates to a voltage compliance of 6.78V.

The linearity of the current source is plotted in Figure 5-27. The current I_{DAC} versus the input code is plotted in Fig. 5-27(a), while the INL and DNL are shown in Fig. 5-27(b). The 6-bit current source achieves an INL and DNL of -0.21LSB/+1.25LSB and -0.14LSB/+0.16LSB respectively which is adequate for neural stimulation applications.

Arbitrary Stimulation Waveform

When the low-voltage digital arbitrary waveform interface is used to control the current source, arbitrary biphasic stimulation waveforms can be easily generated. The shape of the waveform is programmed into the chip via a serial programming interface. Figure 5-28 shows the measured current and voltage from a model electrode with R_s = 3kΩ and C_d = 10nF for (a) a rectangular waveform, and (b) an energy-optimized waveform from a genetic algorithm (GA) discussed in Section 3.2.2. The waveforms are scaled in magnitude to achieve the same neural response based on nerve fiber simulations. Even though the GA waveform shape requires a higher peak current than the rectangular waveform, it consumes less energy and generates a smaller electrode voltage for this particular combination of electrode impedance and PWs. The latter observation implies that the stimulator supply voltage (V_{MID}) could potentially be reduced for the GA waveform which is similar to the findings in [156].

The phase width (PW) and number of steps per phase can also be adjusted by setting the frequency of the ϕ_{HI} and ϕ_W clocks, and PWs of 25μs, 31.25μs, and 50μs with 8 steps/phase are shown in Figure 5-28. Figure 5-29 shows measurements of additional arbitrary current waveforms to further demonstrate the flexibility of the stimulator. The waveforms were measured with a differential preamplifier (Tektronix ADA400A) and a small sense resistor (80.6Ω) was used to sense the current.

Measurements were also performed with a range of electrode impedances to highlight the high output impedance of the current source. Figure 5-30 shows mea-

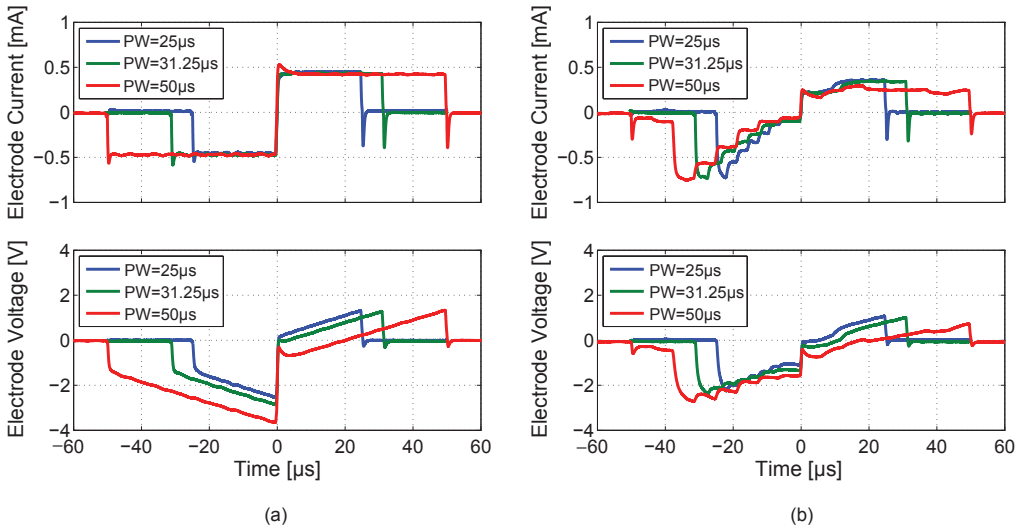

Figure 5-28: Measured current and voltage of a model electrode (R_s=3kΩ, C_d=10nF) for phase widths of 25μs, 31.25μs, and 50μs (8 time steps/phase) with (a) a rectangular waveform, and (b) optimized waveform from a genetic algorithm.

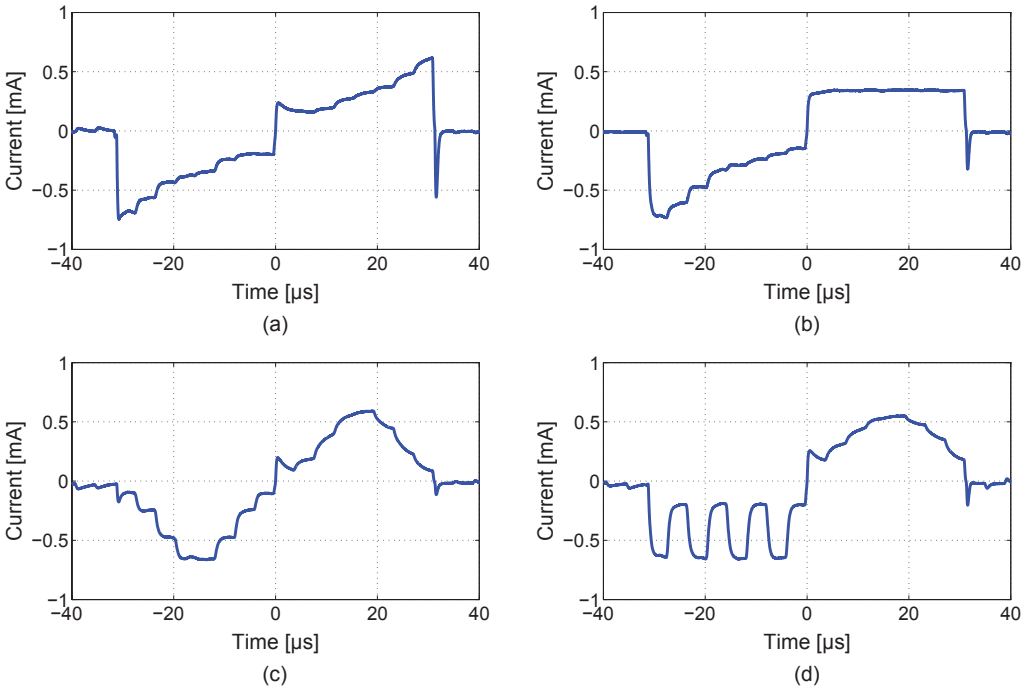

Figure 5-29: Measured electrode current through a model electrode (R_s=3kΩ, C_d=10nF) with a phase width of 31.25μs (8 time steps/phase) and arbitrary waveform shape: (a) decreasing exponential cathodic, increasing exponential anodic, (b) decreasing exponential cathodic, rectangular anodic, (c) triangular cathodic, triangular anodic, and (d) pulsed cathodic, sinusoidal anodic.

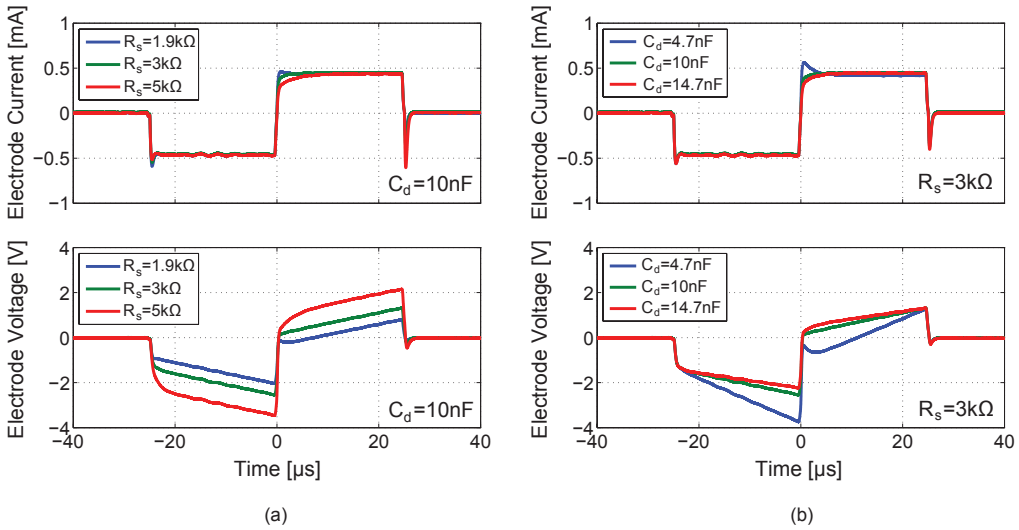

Figure 5-30: Measured electrode current and voltage for a rectangular waveform (25μs phase width) over different electrode impedances. (a) C_d=10nF, R_s=1.9, 3, and 5kΩ. (b) R_s=3kΩ, C_d=4.7, 10, and 14.7nF.

sured current and voltage waveforms for a rectangular pulse over a range of electrode impedances (C_d = 4.7nF to 14.7nF, R_s = 1.9kΩ to 5kΩ). The stimulator is able to maintain the same current across different electrode impedances due to its high output impedance.

Lastly, the average DC current error of the stimulator is well below the safety requirement of 100nA with the use of DC blocking capacitors (220nF) placed in series with the electrode, together with a shorting period.

Electrode Switch Matrix

The high-voltage electrode switch matrix is used to interleave the current source across all electrodes in monopolar configuration (i.e., all electrodes share a common return electrode that provides the return path for the current). Figure 5-31(a) shows a measurement of the interleaved current pulse trains at 1,000 pulses/sec per electrode through all 8 electrodes as required by the CIS sound processing strategy. The stimulation rate corresponds to the 1kHz update rate of the sound processor, and the height of the pulses is modulated by the processor output. In this measurement, the pulses are programmed to be rectangular with a PW of 31.25μs such that the current

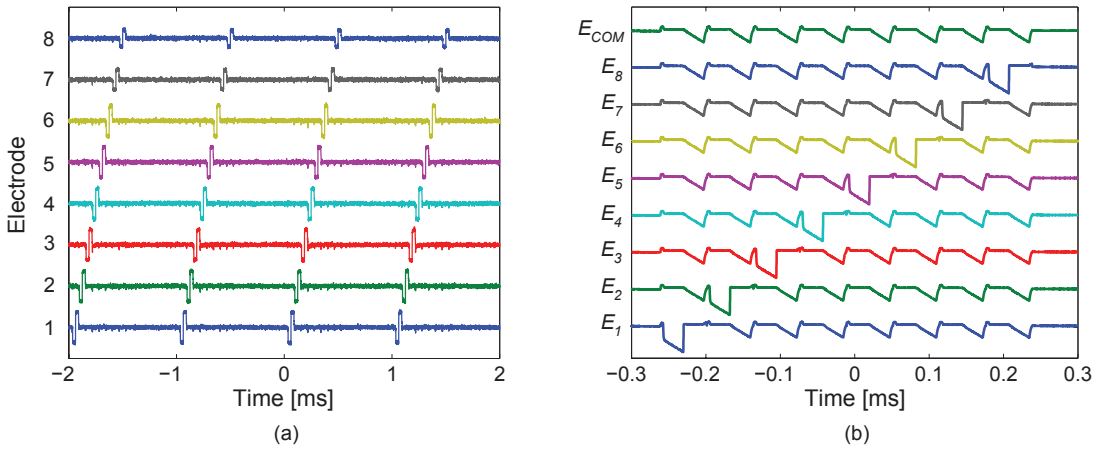

Figure 5-31: (a) Measured interleaved current pulse trains at 1,000 pulses/sec through the electrodes in 8-channel mode. (b) Measured voltages $E_1 - E_8$ and E_{com} of the high-voltage electrode switch matrix during one cycle of stimulation (electrodes are modeled by R_s=3kΩ, C_d=10nF).

source is active for 50% of the 1ms period ((2×31.25μs)×8 = 500μs).

Figure 5-31(b) shows the corresponding voltages measured on the electrode sites ($E_1 - E_8$ and E_{com}) of the switch matrix in Figure 5-12 for one cycle of stimulation.

Stimulator Power Consumption

In order to estimate the power consumption of the stimulator in a typical conversational setting, a two second long clip of speech ("her husband brought some flowers") downloaded from [157] was looped and played into the sound processor. Measurements were collected with the processor in 8-, 6-, and 4-channel modes, and stimulation phase widths of 25μs, 31.25μs, and 50μs for four different waveform shapes (rect/rect, dexp/iexp, dexp/rect, and GA). The magnitude of the waveforms were scaled to achieve the same neural response based on nerve fiber simulations. Figure 5-32 summarizes the measured power consumption with V_{MID} = 7V, which also includes the power consumption of the digital waveform interface at 0.6V, level shifters at 1.8V, and current source circuits at 3.3V.

For all waveforms, the power increases with the PW and number of channels because the current source must be active for a longer amount of time each period.

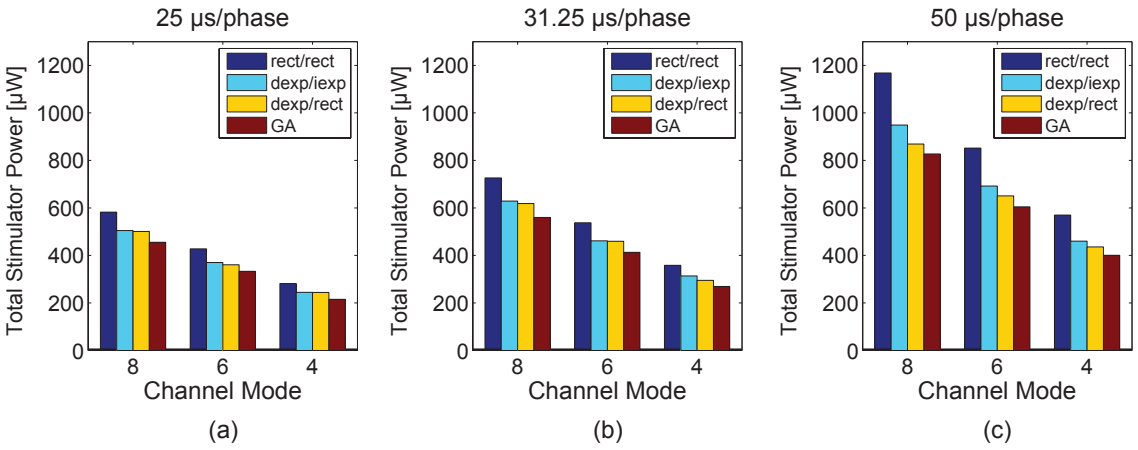

Figure 5-32: Measured total stimulator power at V_{MID}=7V across 8, 6, and 4-channel modes and various stimulation waveforms for phase widths of (a) 25μs, (b) 31.25μs, and (c) 50μs. (rect/rect: rectangular cathodic, rectangular anodic, dexp/iexp: decreasing exponential cathodic, increasing exponential anodic, dexp/rect: decreasing exponential cathodic, rectangular anodic, GA: genetic algorithm waveform)

The duty cycle of the current source is simply $(2\times\text{PW}\times N_{chan})/T_a$ where N_{chan} is the number of active channels and T_a = 1ms is the analysis period of the processor. Therefore, when the stimulator uses a PW of 25μs in 4-channel mode, the current source is active for only 20% of each period, while a PW of 50μs in 8-channel mode requires a duty cycle of 80%. However, the increase in power is not strictly linear with duty cycle because the waveforms at longer PWs typically have lower magnitudes because of the strength-duration characteristic of nerve fibers. That is, a pulse with longer duration requires a lower peak magnitude to initiate a spike in a nerve fiber.

When comparing across waveforms, both waveforms with a decreasing exponential cathodic phase (dexp/iexp, dexp/rect) require less power than the rectangular waveform. Furthermore, as expected, the genetic algorithm (GA) waveform is the most efficient, providing a savings of approximately 22% at PW=25μs and 29% at PW=50μs.

Aside from being able to program the THR and MCL patient fitting knobs for each channel, the sound processor can also reconfigure its global channel gain (applied to all channels) as well as the amount of logarithmic compression C (see Equation 5.13). Both parameters influence the output of the sound processor and therefore

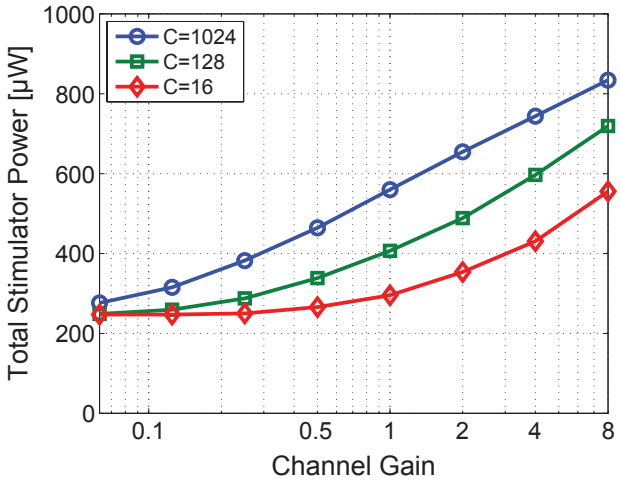

Figure 5-33: Measured total stimulator power (with optimized genetic algorithm waveform shape) versus channel gain setting in the CIS sound processor for compression factors of $C = 1024$, 128, and 16.

impacts the strength of the stimulus provided by the stimulator which is reflected by its power consumption. Figure 5-33 shows the power consumption trends (for the GA waveform) over channel gain settings for $C = 1024$ (most compressed), 128 (moderately compressed), and 16 (most linear).

5.6.4 Power Consumption

The power consumption of the entire SoC with a speech input is summarized in Table 5.4. The PZFE consumes 6.83μA from a 1.5V analog supply under nominal settings, resulting in 10.25μW of power consumption. The ADC and CIS sound processor consume 1.9/1.57/1.43μW of power from a 0.6V digital supply in 8/6/4-channel modes respectively. When the stimulator is programmed with the GA waveform with a PW of 31.25μs, the total stimulation power is 560/413/269μW in 8/6/4-channel modes for $V_{MID} = 7$V, corresponding to 98/97/96% of the total SoC power. Reconfigurability in the number of channels allows for a power-performance tradeoff.

At a phase width of 31.25μs, the GA waveform reduces the stimulation power by approximately 24% when compared to a rectangular waveform. Since the stimulation power dominates the power of the entire SoC, this savings is transferred directly to

Table 5.4: Measured breakdown of typical power consumption in the CI SoC. The stimulator power is measured with a speech input, using the genetic algorithm pulse shape with a phase width of 31.25μs, and nominal CIS processor settings.

Block	Supply Voltage [V]	Supply Current [μA]	Power [μW]
Piezoelectric sensor front-end	1.5		
Stage 1 (charge amplifier)		4.50	6.75
Stage 2 (PGA)		0.91	1.37
Stage 3 (ADC driver)		1.42	2.14
SAR ADC	0.6	0.51	0.31
Digital CIS sound processor	0.6		
(8/6/4-channel modes)		2.65/2.10/1.87	1.59/1.26/1.12
Stimulator and switch matrix		(8/6/4-channel)	(8/6/4-channel)
Digital waveform interface	0.6	0.58/0.54/0.49	0.35/0.32/0.30
Intermediate level shifters	1.8	0.21/0.16/0.11	0.38/0.29/0.20
Current source circuits	3.3	22.7/17.4/12.1	74.9/57.4/39.9
V_{MID} (5V to 10V)	7	68.4/50.2/32.3	479/351/226
V_{DDG} (7V to 12V)	9	0.60/0.45/0.30	5.40/4.05/2.70
Total (8-channel mode)			572
Total (6-channel mode)			425
Total (4-channel mode)			281

the overall system, resulting in an overall SoC power savings of 23%.

5.7 System Demonstration with a Mounted Sensor

The FICI SoC was tested with an actual piezoelectric sensor mounted on a human cadaveric temporal bone (bone098) with the help of Dr. Heidi Nakajima at the Massachusetts Eye and Ear Infirmary. The frequency response of the middle ear and attached sensor was characterized using the SoC following the same procedure outlined in Chapter 4 to check the condition of the bone. This section presents results of the system demonstration.

An arbitrary function generator (Tektronix AFG3102) was used to generate a clip of speech ("her husband brought some flowers") which was amplified by a Crest audio amplifier and played into the ear canal of the temporal bone through a speaker. The signal from the umbo-mounted piezoelectric sensor was picked up by the PZFE and processed by the CIS sound processor on the SoC. From the sound processor output, a spectrogram was generated, and sound was reconstructed (similar to sound synthesis in vocoders). For the sound reconstruction, the output of each channel is used to

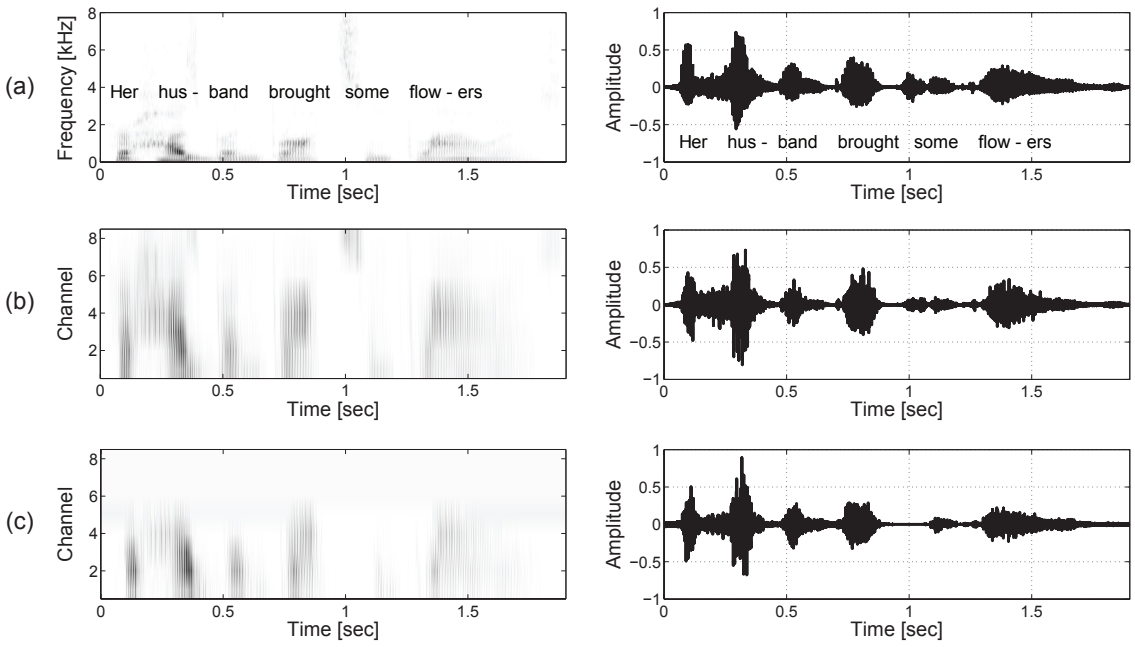

Figure 5-34: (a) Spectrogram and time-domain waveform of the input speech signal ("her husband brought some flowers") to the audio amplifier driving the speaker placed in the ear canal of the temporal bone. (b) Simulated spectrogram and reconstructed sound of the 8-channel processor in Matlab. (c) Measured spectrogram and reconstructed sound from the CI SoC with the piezoelectric sensor mounted on a temporal bone.

modulate a pure tone with a frequency equal to the center frequency of that channel, and all tones are summed up.

Figure 5-34(a) shows the spectrogram and time-domain waveform of the input speech signal in the ear canal, and Figures 5-34(b) and (c) show the simulated and measured spectrograms and reconstructed sound of the processor in 8-channel mode. The measured output includes some non-idealities like the bandwidth limitations of the PZFE, but it preserves the envelope information of the speech signal. Essentially, this demonstrates hearing with a cadaver ear using a mounted piezoelectric sensor and the FICI SoC.

5.8 Chapter Summary

This chapter presented the design, analysis, and characterization of a proof-of-concept SoC in a 0.18μm high-voltage CMOS technology for a fully-implantable cochlear implant. A custom piezoelectric sensor front-end (PZFE) for an umbo-mounted piezoelectric sensor is optimized for low-noise and ultra-low-power, and its key measured performance metrics are summarized in Table 5.5. Overall, the PZFE has the required dynamic range for speech, it can accommodate a range of sensor sizes with programmable gain settings, and it consumes only 10.3μW from a 1.5V supply. This is almost two orders of magnitude lower when compared to the discrete prototype described in Chapter 4 which uses five commercially-available AD8603 op-amps consuming 660μW from 3.3V.

Second, a reconfigurable multi-rate CIS sound processor is designed, consuming just 1.6μW from 0.6V in the 8-channel mode, and its performance summary is provided in Table 5.6. It leverages ultra-low-voltage digital processing to improve energy-efficiency, and it can be reconfigured in 8-, 6-, and 4-channel modes to enable a power-performance tradeoff. Furthermore, the processor parameters are highly reconfigurable to enable a patient-specific fit.

Third, the design of an arbitrary waveform stimulator allows the SoC to deliver energy-efficient neural stimulation waveforms to reduce the stimulation power. Because the PZFE and sound processor consume just over 10μW combined, the stimulator power represents more than 95% of the total SoC power. By being able to deliver alternate energy-efficient waveforms like the GA waveform from Chapter 3, the overall SoC power is reduced by 23% at a phase width of 31.25μs when compared to conventional rectangular stimulation. The performance summary of the stimulator subsystem is provided in Table 5.7.

Lastly, a demonstration of the SoC with a piezoelectric sensor mounted on the umbo of a human cadaveric temporal bone was completed by playing a clip of speech into the ear canal, and reconstructing the output of the SoC. This demonstrated hearing with a cadaver ear using an implantable sensor and the SoC. Overall, the

Table 5.5: Measured performance summary of the piezoelectric sensor front-end. [1]The input-referred noise is from stage 1 only (stage 1 dominates the overall noise budget). [2]The NEF and PEF calculations are based on the input-referred noise of stage 1 integrated from 300Hz to 25kHz, using a 6.5kHz 3dB bandwidth.

Sensor size	C_P**=0.56nF**	C_P**=3.2nF**
Supply voltage	1.5V	
Total front-end power	10.26μW	
Stage 1 gain (charge amplifier)	21.2 to 41dB (C_{1f}=48 to 6pF)	33.8 to 53dB (C_{1f}=66 to 8.5pF)
Stage 2 gain (PGA)	0.8, 6.8, 12.4, 18, 23.8, 29.4dB	
Stage 3 gain (ADC buffer)	12dB	
Tunable low-pass cut-off	6 to 7kHz	
[1]Input-referred noise (300-5500Hz)	2.81μV_{rms} (C_{1f}=12pF)	1.93μV_{rms} (C_{1f}=34pF)
Maximum input signal (1% THD)	10m$V_{p\text{-}p}$ (C_{1f}=12pF)	4.7m$V_{p\text{-}p}$ (C_{1f}=34pF)
Input dynamic range	62dB	58.7dB
[2]NEF/PEF (stage 1 only)	4.01/24.1	2.68/10.7
PSRR	46dB	59dB

Table 5.6: Measured performance summary of the ADC and reconfigurable CIS sound processor.

Supply voltage	0.6V
SAR ADC	
Resolution	9 bits
Sample rate	16kS/s
Linearity	INL: -0.65/0.69LSB DNL: -0.64/0.59LSB
SNDR/ENOB	52dB/8.35b
Figure-of-merit	66.7 fJ/conversion-step
Power consumption	0.31μW
Reconfigurable CIS Sound Processor	
Total bandwidth	300Hz to 5500Hz
Number of channels	8, 6, or 4
Analysis filters	Multi-rate decimation filter with FIR filter bank (-30dB stop band)
Input data rate and resolution	16kHz, 9 bits
Rectification	Full-wave or half-wave
Envelope detector low-pass cut-off	400Hz (2^{nd}-order Butterworth)
Envelope update rate	1kHz
Channel gain settings	1/16, 1/8, 1/4, 1/2, 1, 2, 4, 8
Compression function	$Y = \ln(1+CX)/\ln(1+C)$ C=1024, 128, 16
Threshold (THR) and most-comfortable-level (MCL) tuning resolution	3 bits
Output data rate and resolution	1kHz, 6 bits
Power consumption	1.59μW (8-channel) 1.26μW (6-channel) 1.12μW (4-channel)

Table 5.7: Measured performance summary of the arbitrary waveform stimulator and high-voltage switch matrix. [1]The measurement of the current source output impedance was limited by the accuracy of the measurement equipment.

Stimulation pulse rate/electrode	1000 pulses/sec
Number of electrodes	8, 6, or 4
Number of current sources	1, interleaved
Current source	
Resolution	6 bits
Full-scale	1mA
Linearity	INL: 1.25LSB, DNL: 0.16LSB
[1]Output impedance (measured)	>20MΩ
Output impedance (simulated)	300MΩ
Voltage compliance (V_{MID}=7V)	6.78V @ full-scale
Waveform	
Type	Biphasic, cathodic first
Phase width	25μs to 50μs
Arbitrary waveform	Yes, up to 8 time steps/phase
Target nominal electrode impedance	C_d=10nF, R_s=3kΩ
Electrode array configuration	Monopolar
DC blocking capacitors	Yes, required
Average DC current error/channel	$\ll$100nA (with DC blocking capacitors and shorting period)

SoC consumes less than 600μW when delivering the GA waveform during normal conversation. This would enable 21 hours of operation from a single 5g ultra-capacitor with an energy density of 5 W·hr/kg with 50% power conversion efficiency, therefore permitting a one-charge-per-day usage model.

Chapter 6

Conclusion and Future Directions

6.1 Summary of Contributions

From a user or clinician's point of view, there is always a desire to increase functionality, maximize intelligence, and extend lifetime, while minimizing the size and burden of medical devices. The drive for more functionality or *smarts* in a device, available over a longer period of time, usually requires more energy. On the other hand, miniaturizing the size of devices ultimately limits the volume available to energy-storage devices. These competing requirements of personalized medical devices present an opportunity for research and development in the area of biomedical applications that this thesis has tried to address.

The energy problem can be addressed in two ways: 1) increase the energy available to the system by increasing the density of energy-storage elements or by harvesting energy from the environment, or 2) improving the energy-efficiency of the system. This thesis focuses on the latter, by investigating the design and implementation of *ultra-low-power* and *energy-efficient* circuits and systems for two different medical applications. The goal for both systems is to minimize the overall system power to meet the stringent energy requirements. This is accomplished by applying four general low-power design principles for medical applications to a system design methodology that considers the system holistically. This can include optimizations made at the levels of physiology, sensors, architecture, and circuits.

The first principle is to *extend the scope* of the design beyond IC design in order to gain an understanding of the interface to the human body. Often times, this high-level knowledge can have a profound impact in shaping the system architecture. For example, by examining the electrode-body interface and understanding the sources of signal aggressors that can corrupt an ECG signal, the ECG front-end architecture was designed to specifically improve robustness against electrode offset and power-line interference. Extending the boundary of the work can also provide new opportunities to innovate at other layers of the design. For example, at the physiologic level, the optimization of the shape of electrical neural stimulation waveforms can reduce the overall power of a cochlear implant (CI) by 20-30%. This result was validated with *in-vivo* measurements of threshold and loudness with two human CI subjects. At the sensor level, the development of a low-power piezoelectric middle ear sensor with sensitivity from 40 to 90dB SPL is a key enabler of a fully-implantable CI.

A second guiding design principle is the use of *highly-digital, mixed-signal circuits and systems*. Although low-noise, robust analog design is required to interface to the human body or real world, the balance and blend of analog and digital computation is critical for improving the energy-efficiency and robustness of the system. For example, the ECG front-end system (conventionally an analog system) leveraged highly-digital techniques like oversampling and $\Delta\Sigma$-modulation to improve system robustness and dynamic range, and a mixed-signal feedback loop to improve robustness against power-line interference. Furthermore, in the CI SoC, a large amount of digital processing was used to accomplish *micro-Watt* reconfigurable sound processing, and to provide ultra-low-power control of the high-voltage arbitrary waveform stimulator that interfaces to neural tissue.

The use of highly-digital circuits and systems often leads directly to the third design principle, which is the use of *aggressive voltage scaling*. Voltage scaling is possible because medical applications are generally low-bandwidth. Scaling down to 0.6V improves the energy-efficiency of digital circuits by operating closer to their minimum energy point [147]. For example, a savings of greater than 5μW is achieved by scaling the supply from 1.5V to 0.6V for the mixed-signal ECG front-end in Chapter

2. Also, ultra-low-voltage digital computation is leveraged in the CI SoC to achieve *micro-Watt* sound processing and stimulation control. Furthermore, both prototyped systems featured a SAR ADC operating at 0.6V, which is close to its optimum energy efficiency point [1]. Finally, Chapter 2 highlighted four low-voltage analog techniques that enabled the analog circuits of the ECG front-end to operate from the same low voltage digital supply. The four techniques are summarized as follows: 1) the use of folded circuits and sub-threshold biasing to enable cascoding, 2) cascaded topologies, 3) fully-differential operation with local CMFB loops, and 4) avoiding mid-rail signal path switches and the use of bootstrap circuits.

The final important design principle is three-pronged: *integration, customization, and optimization.* First, the ECG front-end exhibited a high-level of integration by including all peripheral functions like voltage and current referencing and clock generation on-chip. This drastically simplifies the functional and power management requirements of the overarching band-aid system. Integration also provides opportunities to optimize the interface between sub-blocks. For example, a dual-DAC SAR ADC was merged with the *SINC* anti-aliasing filter in the ECG front-end system in order to eliminate the need for a fast-settling power-hungry ADC buffer. Second, customization of the system for a specific application can avoid wasting power on general purpose functions such as with FPGAs and microcontrollers, or generic op-amps and ADCs. In this thesis, this was illustrated with two custom IC prototypes for two specific medical applications. Third, the application of fine-grained circuit-level optimizations to all blocks can result in significant overall power reduction. For example, clock gating in the CI sound processor provided power scalability with the number of channels. Analog power gating was used to shut down the stimulator current source when inactive. Multi-rate signal processing was leveraged in the CI sound processor to operate only as fast as necessary. Also, digital optimizations on filter word length, coefficients, and multiplier-less structures were also employed. Finally, careful low-noise analog design was applied to the LNA of the ECG front-end and stage 1 of the piezoelectric sensor-front end to achieve good analog power-efficiency.

6.2 Future Work

Despite the contributions of this thesis, there are countless directions and opportunities for future work for both applications. Some of the more imminent directions are discussed next.

6.2.1 Robustness in Ambulatory ECG Monitoring

A wearable monitor can be subject to large amounts of motion artifact that arise from random motion such as walking or running. Any physical motion can cause changes at the electrode-skin interface, as well as deformation in the skin [158]. Due to the random nature of these artifacts, they are usually the most difficult to compensate for. A future generation of the front-end should try to address the issue of motion artifacts.

Techniques for motion artifact reduction typically leverage correlation with other sensors or electrodes. For example, principle component analysis using multiple leads can isolate motion artifact by projecting the signals onto a different set of basis vectors [159]. Impedance monitoring of the electrode site can also be done to log motion artifact events (assuming the motion artifacts are caused by changes in the impedance) [10]. Lastly, adaptive LMS filtering can be done with reference signals such as motion from accelerometers [160], electrode impedance [159], or skin stretch from an optical sensor [161].

6.2.2 Toward a Truly Invisible Hearing Prosthesis

The goal of the proof-of-concept FICI SoC in this thesis is to demonstrate feasibility. However, the SoC has many shortcomings, and in order to actually realize an invisible device that can be implanted inside a human ear, numerous challenges and future directions remain. Some of them are listed below:

- **Sensor mounting, packaging, and surgical techniques**: Perhaps the largest barrier to the commercialization of this technology is the development that

needs to be done on mounting the sensor, hermetically sealing and packaging the system, and developing the surgical techniques to implant the system. The piezoelectric sensor in this work was attached to a needle which was held in place by a micro-manipulator external to the temporal bone. The final implantable system will have to address the issue of mounting the sensor within the small space available in the middle ear cavity. Figure 6-1 shows a proposed method of anchoring the piezoelectric sensor to the mastoid using a tiny PCB clamp attached to a surgical Y-plate.

- **Further validation of energy-efficient neural stimulation waveform**: At this time, the experimental validation of the alternate stimulation waveforms using cat and human data is only at its infancy. The results from two cat subjects and two human CI subjects are only very preliminary. More subjects will have to be tested to gather statistics regarding the efficacy of the alternate stimulation waveforms.
- **Wireless power delivery**: This thesis addressed the design of the core circuit blocks for a FICI. A wireless power delivery system will have to be developed to rapidly and wirelessly recharge an implanted battery or ultra-capacitor once a day [43]. An efficient power management system is also needed to power the various analog, digital, and high voltage supply domains.
- **Combined electric-acoustic hearing**: Many studies have discussed the benefits of combined electric and acoustic hearing [138, 148, 162]. For example, any residual low-frequency acoustic hearing can be enhanced by a hearing aid, and combined with high-frequency electric hearing from a cochlear implant. Therefore, a future implantable device could provide electric hearing and enhance residual acoustic hearing to maximize performance.
- **Next-generation FICI**: A 2^{nd} generation of the FICI should achieve much higher performance. For example, automatic gain control can be added to the piezoelectric front-end, and the sound processor can be expanded to include up to 16 channels to provide increased spectral resolution. Furthermore, the energy-efficient stimulation waveforms are complementary to alternative sound

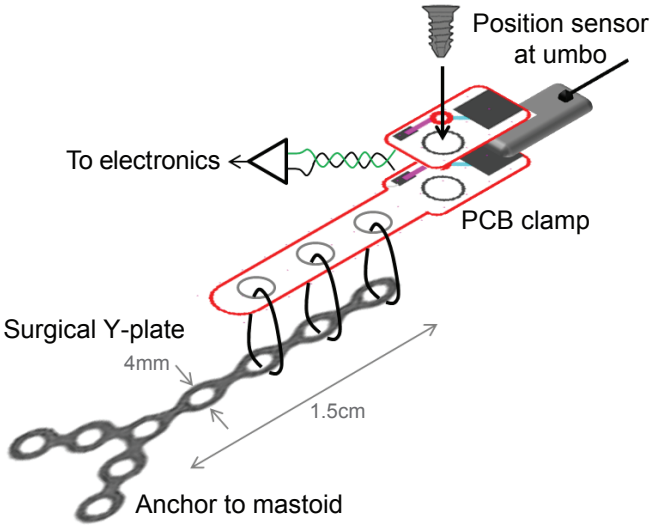

Figure 6-1: Proposed method of anchoring the piezoelectric sensor to the mastoid using a tiny PCB clamp attached to a surgical Y-plate.

processing strategies like Asynchronous Interleaved Sampling (AIS) which stimulates less frequently than the CIS strategy while potentially providing better performance [96]. A digital AIS processor with energy-efficient stimulation waveforms can be included in the next generation SoC to further reduce stimulation power. Lastly, in this work, a current source-based stimulator was implemented because they are simple, open-loop, have known methods of charge balancing, and can achieve fast settling for arbitrary waveform generation. However, a large amount of power is wasted in the headroom of the current source. Future work can explore the design of adiabatic or DC-DC converter-based stimulators that have enough bandwidth to provide arbitrary stimulation waveforms for CI applications [97]. Also, charge balancing techniques can be added to the stimulator to avoid the need for large DC blocking capacitors in order to reduce the size of the system [163, 164].

Appendix A

ECG Signal Aggressors and Specifications

A.1 Signal Aggressor Details

This section of the appendix supplements the discussion on ECG signal aggressors from Section 2.1.1. All off-chip and on-chip signal aggressors are illustrated in Figure 2-2.

A.1.1 Off-Chip Aggressors

Electrode offset: Figure A-1 shows a model of the electrode-electrolyte interface for ECG applications, where R_d and C_d model the impedance of the electrode-electrolyte interface, R_s is the series resistance of the electrolyte gel, and E_{hc} is the half-cell potential from charge accumulation arising from a chemical reaction at the electrode-electrolyte interface [60, 61]. The half-cell potential is dependent on the electrode material, skin anatomy, and sweat, and therefore any mismatch in E_{hc} between the two electrodes used in any differential measurement can result in a large differential DC electrode offset. As such, the front-end must have a high-pass filter (HPF) characteristic and be able to reject up to ±300mV of electrode offset [62].

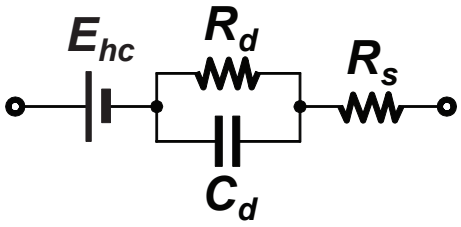

Figure A-1: Model of the electrode-electrolyte interface.

Signal attenuation: At low frequencies, the contact impedance of electrodes can range from 10's of kΩ for wet Ag/AgCl gel electrodes, up to the MΩ range for dry electrodes [20]. Therefore, the front-end must have high differential-mode input impedance (>100MΩ) to prevent significant signal attenuation due to the impedance divider effect.

50/60Hz power-line interference (PLI): Surrounding power lines can inject a large amount of 50Hz or 60Hz interference into the system via a number of different mechanisms: 1) magnetic induction, 2) capacitive coupling to the electrode leads, and 3) capacitive coupling to the human body. Of the three mechanisms, capacitive coupling to the body is the most difficult to mitigate. Since PLI cancellation is one of the main focuses of this work, a detailed discussion on PLI is provided in Appendix A.2.

Motion artifact: Motion artifacts are potentials that are generated when the skin-electrode interface is disturbed [158]. These artifacts are often due to random movement, making them the most difficult aggressor to compensate for.

A.1.2 On-Chip Aggressors

Device thermal noise: Transistors in any circuit (e.g., an op-amp) will contribute to its overall input-referred noise and limit the minimum detectable signal of the system. However, through appropriate device sizing, biasing, and amplifier topology, the input-referred noise of an amplifier can be minimized. A measure of the power effi-

ciency of an amplifier can be characterized by the noise efficiency factor (NEF) [81], which normalizes the noise of an amplifier with respect to its power consumption and noise bandwidth.

Low-frequency flicker noise: Flicker (or $1/f$) noise in MOSFETs arises from charge traps in the silicon-gate oxide interface and is independent of current density. Because of its $1/f$-shaped spectrum, flicker noise increases at lower frequencies making it particularly troublesome for biomedical signals which reside at low frequencies. However, flicker noise can be reduced by increasing transistor size and by using PMOS devices rather than NMOS devices. Alternatively, for applications that require very low noise, a modulation technique known as chopper-stabilization can be used [165].

Amplifier offset: Amplifier offset (not to be confused with electrode offset) arises mostly from threshold voltage mismatch in differential-pair transistors due to random dopant fluctuation. For biomedical applications where it may be necessary to achieve low offset, increasing the size of the input devices can reduce the threshold voltage mismatch. Alternatively, chopper-stabilization or auto-zeroing can also remove amplifier offset [165].

Substrate and power supply noise: Substrate and power supply noise are particularly important in mixed-signal systems, where the switching activity from noisy digital circuits can inject unwanted noise through the substrate and/or the power supply onto sensitive analog circuits such as the front-end. Therefore, it is important to achieve good CMRR and power supply rejection ratio (PSRR), often through a fully-differential architecture.

A.2 Power-Line Interference

Power-line interference in bio-potential recording has been studied extensively in literature over the years [2, 166–168], and a commonly used model of the impedances at the interface between the body and front-end circuits is shown in Figure A-2. The definitions of the model parameters are provided here for convenience:

- V_{PL}: The power-line voltage in the mains (e.g., 50Hz 230V_{rms} in Europe, 60Hz 120V_{rms} in North America)
- C_P, C_B: The stray coupling capacitance from the body to the power line (C_P), and to earth ground (C_B)
- i_P: The displacement current through the body as a result of the coupling capacitances C_P and C_B
- Z_{E1}, Z_{E2}: The electrode contact impedances for the two measurement electrodes ($E1$ and $E2$) involved with picking up the differential signal
- Z_{E3}: The electrode contact impedance for an optional third ($E3$) patient ground reference electrode (switch S_3 can be opened for two-electrode systems)
- Z_T: The internal body impedance between the two measurement electrodes
- Z_C, Z_D: The common-mode input impedance (Z_C) and differential-mode input impedance (Z_D) of the front-end circuitry
- GND_E, REF_{SYS}: The earth ground (GND_E) and the system reference (REF_{SYS}) for isolated systems (e.g., battery operated systems)
- C_{ISO}: The isolation capacitance between earth ground and the system reference (100's of pF for bench-top equipment, a few pF for isolated battery systems [168])

PLI can interfere with the system via three major mechanisms: 1) magnetic induction, 2) capacitive coupling to the electrode leads, and 3) capacitive coupling to the human body [166]. Magnetic induction can be a problem because a conductive loop is formed with the body, electrode leads, and front-end circuitry, and the time-varying magnetic field can induce an *ac*-potential in the loop that depends on the area and orientation of the loop, and amplitude of the magnetic field. This mechanism can be mitigated by minimizing the area of the loop by twisting the cables, or by placing the front-end circuits close to the body (e.g., wearable sensor nodes).

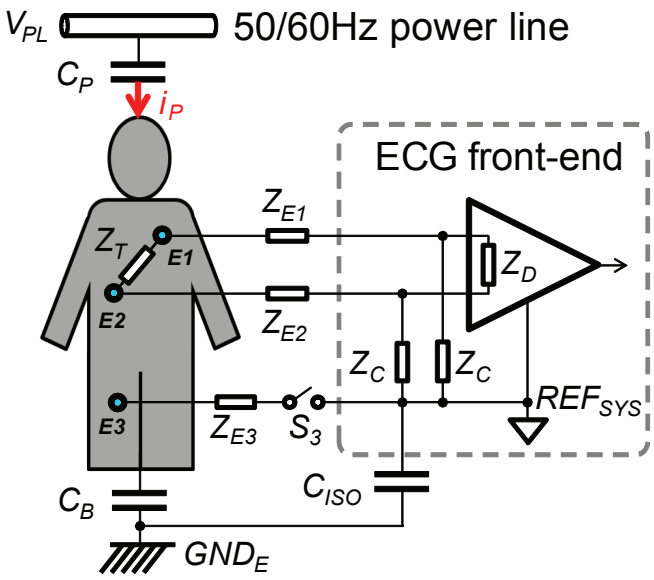

Figure A-2: Model of the impedances at the interface between the body and front-end for calculation of power-line interference [2].

The second mechanism is capacitive coupling to the electrode leads and this can be addressed by using shielded cables, or by using active buffering circuitry right at the electrode site to convert the high impedance node at the electrode-tissue interface, to a low impedance output that is unaffected by capacitive coupling from the power lines [52].

The third mechanism of capacitive coupling to the body is typically the most dominant because it is the most difficult to overcome. The displacement current (i_P) through the body can be calculated as

$$i_P = \frac{V_{PL}}{Z_{C_P} + Z_T + Z_{C_B}} \tag{A.1}$$

where Z_{C_P} and Z_{C_B} are the impedances of C_P and C_B respectively. For typical values of C_P (0.3pF to 10pF), C_B (120pF to 520pF) and Z_T (up to a few kΩ) [168], the impedance of C_P dominates at 50/60Hz, and so $i_P \approx V_{PL}\omega C_P$. For expected values of C_P, i_P can range from 10nA$_{\text{rms}}$ up to 0.5μA$_{\text{rms}}$. The displacement current i_P can result in differential-mode interference between the two measurement electrodes ($E1$ and $E2$) in two ways: 1) by passing through the internal body impedance Z_T

generating a differential voltage

$$V_{DI,1} = i_P Z_T \tag{A.2}$$

between $E1$ and $E2$, and 2) generating a large common-mode interference voltage (V_{CI}) that is converted to differential-mode interference ($V_{DI,2}$) through non-idealities in the electrodes (mismatch between Z_{E1} and Z_{E2}) and the front-end circuit (finite Z_C and CMRR). The total differential-mode interference can be expressed as $V_{DI,tot} = V_{DI,1} + V_{DI,2}$.

Figure A-2 can be used to determine expressions for V_{CI} and $V_{DI,2}$. Neglecting Z_T (which is typically much less than Z_{C_P} and Z_{C_B} at 50/60Hz), V_{CI} on the body for two-electrode systems (switch S_3 open) can be calculated as

$$V_{CI} = \frac{C_P}{C_P + C_B} \cdot V_{PL}. \tag{A.3}$$

While this common-mode interference (V_{CI}) would not affect a system with perfectly matched electrodes ($Z_{E1} = Z_{E2}$), infinite common-mode input impedance Z_C, and infinite CMRR, any practical system will suffer from an electrode impedance mismatch of $\Delta Z_E = Z_{E1} - Z_{E2}$, as well as finite Z_C and CMRR. Given these non-idealities and assuming that the percentage matching of Z_C on-chip is much better than the percentage matching of Z_{E1} and Z_{E2}, the resulting differential-mode interference $V_{DI,2}$ at the input of the front-end is given by

$$V_{DI,2} = \left(\frac{Z_C}{Z_C + Z_{E2}} - \frac{Z_C}{Z_C + Z_{E1}} \right) \cdot V_{CI} + \frac{V_{CI}}{CMRR}. \tag{A.4}$$

Furthermore, assuming that a properly designed front-end will have $Z_C \gg Z_{E1,2}$, then Equation A.4 reduces to

$$V_{DI,2} \approx \left(\frac{\Delta Z_E}{Z_C} + \frac{1}{CMRR} \right) \cdot V_{CI}. \tag{A.5}$$

Therefore, by combining Equations A.2, A.3 and A.5, the total differential-mode interference as a result of the displacement current is given by

$$V_{DI,tot} = V_{DI,1} + V_{DI,2} \approx \left[\omega C_P Z_T + \left(\frac{\Delta Z_E}{Z_C} + \frac{1}{CMRR}\right)\left(\frac{C_P}{C_P + C_B}\right)\right] \cdot V_{PL}. \quad \text{(A.6)}$$

Examining the parameters from Equation A.6, the only parameters that are controlled by the system designer are Z_C and CMRR. Even if these were ideal and infinite, the differential-mode interference would still be $V_{DI,tot} = \omega C_P Z_T V_{PL}$. As a numerical example, assuming worst case values of C_P = 15pF, Z_T = 2kΩ and a line voltage of V_{PL}=120V_{rms} at 60Hz, the differential-mode interference can be as large as $V_{DI,tot}$ = 1.36mV_{rms} (3.84$mV_{p\text{-}p}$) which can be as large as (or larger than) the ECG signal itself (0.1$mV_{p\text{-}p}$ to 5$mV_{p\text{-}p}$). If finite Z_C and CMRR are factored in, $V_{DI,tot}$ up to 10$mV_{p\text{-}p}$ can be possible.

Methods to Reduce Differential-Mode Interference

Considering the analysis in the previous section, it is clear that there is no way for the system designer to eliminate the differential-mode interference $V_{DI,1} = i_P Z_T$ which results from i_P passing through the body impedance Z_T. However, $V_{DI,2}$ (Equation A.5) can be minimized by ensuring that the front-end provides very large Z_C and CMRR. Specifications for values of Z_C and CMRR will be discussed in Appendix A.3.

Another method to reduce $V_{DI,2}$ is to limit the common-mode interference V_{CI}, which is usually necessary for practical purposes anyway. For example, using Equation A.3 with typical values of C_P=5pF, C_B=250pF, and V_{PL}=120V_{rms} at 60Hz, V_{CI}=2.35V_{rms} (6.65$V_{p\text{-}p}$), which is beyond the supply rails for any front-end implemented in a deep sub-micron CMOS technology. Therefore, a common way to limit V_{CI} is to use a third patient ground electrode ($E3$ in Figure A-2 with switch S_3 closed) to provide a low impedance path to the isolated system ground reference (REF_{SYS}) for i_P.

Another classic method to further reduce V_{CI} is to use the Driven-Right-Leg technique [169], which senses the common-mode voltage on the body and uses a feedback

amplifier to drive it to zero. This active feedback loop reduces the effective impedance between the body and REF_{SYS}, and limits V_{CI} even more when compared to using just a passive ground electrode.

However, as mentioned previously, even if both Z_C and CMRR were infinite, and the Driven-Right-Leg technique was applied, there would still be a residual amount of differential-mode interference due to the displacement current passing through Z_T. This residual amount of interference can only be eliminated by notch filtering [59] which is discussed in Section 2.2.2.

A.3 System Specifications Details

This section of the appendix supplements the discussion on system specifications for the ECG front-end from Section 2.2.1. Numerical examples are provided here to justify the system specifications listed in Table 2.2.

Input Impedance and CMRR

Both IEC and ANSI/AAMI specifications refer only to differential-mode input impedance (Z_D). The requirement for Z_D is tied mainly to signal attenuation. From Figure A-2, any differential signal across electrodes $E1$ and $E2$ will be scaled by $Z_D/(Z_D + 2Z_E)$ due to the voltage divider created by $Z_{E1,2}$ and Z_D. Therefore, if an attenuation of up to 10% (<1dB) is permitted, then the requirement is $Z_D \geq 18Z_E$. According to [20], the worst case electrode impedance occurs at the low end of the signal bandwidth and can reach 1MΩ at 1Hz. Therefore, Z_D should be larger than 18MΩ at 1Hz, which is slightly more strigent than required by the IEC and ANSI/AAMI standards.

As discussed in Appendix A.2, the requirements for the common-mode input impedance (Z_C) and CMRR of the front-end are tied to the reduction of 50/60Hz PLI. Following the analysis of [69] and using Equation A.5, if we assume that V_{CI} is limited to $1V_{p\text{-}p}$ by using a third patient ground electrode, then a CMRR value of 80dB would only add $0.1mV_{p\text{-}p}$ ($0.035mV_{rms}$) of PLI. Similarly, if we assume a worst case electrode

impedance mismatch of ΔZ_E=50kΩ [166], then $Z_C = 2000\Delta Z_E = 100\text{MΩ}$ would add $0.5\text{mV}_{\text{p-p}}$ ($0.18\text{mV}_{\text{rms}}$) of PLI.

If these two sources of interference are added to the worst case interference resulting from the displacement current passing through Z_T ($V_{DI,1}$=$1.36\text{mV}_{\text{rms}}$ calculated in Appendix A.2), then the total amount of PLI is $V_{DI,tot} = 1.36\text{mV}_{\text{rms}} + 0.035\text{mV}_{\text{rms}} + 0.18\text{mV}_{\text{rms}} = 1.58\text{mV}_{\text{rms}}$. Therefore, since $V_{DI,tot}$ is dominated by $V_{DI,1} = i_P Z_T$, even modest values of CMRR (≥60dB) and Z_C (≥50MΩ) are acceptable and do not increase the amount of PLI appreciably.

PLI and Electrode Offset Tolerance

Although the IEC and ANSI/AAMI standards do not have a hard requirement on PLI tolerance, it can be determined based on the analysis in Appendix A.2. Even if large values of CMRR and Z_C are achieved, the interference due to $V_{DI,1} = i_P Z_T$ still exists. Considering this in combination with finite values of CMRR (60-80dB) and Z_C (50-100MΩ), the front-end is required to tolerate up to $10\text{mV}_{\text{p-p}}$ of PLI without swamping out the ECG signal itself.

Both the IEC and ANSI/AAMI specifications do state that the front-end must also be able to reject up to ±300mV of electrode offset. This large value is required in order to accommodate the largest expected mismatch in the electrode half-cell potential (E_{hc}) shown in Figure A-1.

Noise and Dynamic Range

The more stringent ANSI/AAMI specification for input-referred noise is stated as $30\mu\text{V}_{\text{p-p}}$ over 10 seconds [63]. Assuming conservatively that the system is sampled at 500Hz, then N_{samp}=5,000 samples are recorded in a 10 second window. To translate this peak-to-peak noise specification to a RMS value, it is necessary to determine an acceptable confidence level to ensure that the peak-to-peak noise of the 5,000 samples will not exceed $30\mu\text{V}_{\text{p-p}}$ in the 10 second window [170]. The number of standard deviations (σ) within the peak-to-peak value can be calculated as $2\sqrt{2}\cdot\text{erfinv}(1 - 1/N_{samp})$. Applying this formula, we see that the peak-to-peak value must span

7.44σ. Since the RMS value is by definition 1-σ, the RMS noise specification is given by $30\mu V_{p\text{-}p}/7.44 = 4.03\mu V_{rms}$. This level of noise would allow the front-end to easily detect ECGs with a minimum amplitude of 0.1mV.

The maximum input signal specified by IEC and ANSI/AAMI are $6mV_{p\text{-}p}$ and $10mV_{p\text{-}p}$ respectively. However, it is also important to specify the amount of distortion acceptable at that level. According to the literature [70, 145], ensuring that the THD $< 1\%$ at the maximum input amplitude is sufficient for most biomedical applications. Combining the maximum input signal with the RMS noise requirement, the required dynamic range for the ECG system is approximately 50dB (IEC) to 59dB (ANSI/AAMI).

Supply Voltage and Power Consumption

One of the main focuses of this work is to develop an ECG front-end that is able to operate from the same supply as ultra-low-voltage DSPs. As such, the front-end in this work must be able to operate from a 0.6V supply.

The requirement for power consumption comes from the energy capacity of the chosen battery for the entire wearable ECG sensor node. In order to achieve a thin and small form factor, ultra-thin lithium polymer batteries are desirable. For example, a 12mm $\times$ 12.5mm battery that is just 2mm thick (PGEB201212) is available from [56]. However, it does suffer from low capacity, and the battery provides only 10mA·hr at 3.7V. In order to achieve 2 week operation, this requires the entire ECG sensor node to consume an average of 30μA. Assuming that 100% battery efficiency is possible (valid for low peak current draw), then the available input power to the power management subsystem is $P_{IN} = 3.7V \times 30\mu A = 110\mu W$. Assuming that the DC-DC converter generating 0.6V from 3.7V is 50% efficient, then just 55μW is available to the AFE, ADC, DSP and communication subsystems. Therefore, the target power consumption for the ECG front-end should be less than 10μW.

Gain and Bandwidth

ECGs have signal amplitudes on the order of 0.1mV to 5mV [61], with signal energy residing over a frequency range that starts around 0.05-0.67Hz, and ends around 100-200Hz [61–63]. Therefore, the bandwidth of the ECG front-end should be approximately 0.5Hz to 150Hz.

In order to specify the gain, assuming that the ADC full-scale is 0.5V (equal to the supply voltage minus 100mV of headroom), then the gain of the ECG front-end should be programmable such that signals ranging from 0.1mV to 5mV can all be gained up to utilize as much of the full-scale range as possible. This requires that the ECG front-end be able to provide gain between 40dB to 70dB.

Appendix B

Stimulator Efficiency Analysis

This appendix presents analysis and comparison of the energy efficiency of two neural stimulator circuit approaches: 1) a current-controlled (i.e., current source-based) stimulator [155,163,171], and 2) a DC-DC converter-based (i.e., inductor-based) stimulator [97].

B.1 Current Source Approach

Figure B-1 shows three possible approaches to implement a current-controlled biphasic current stimulator. In all three approaches, the tissue is biased at a mid-rail voltage of V_{MID}. Fig. B-1(a) shows a method that relies on two current sources I_P and I_N and a SPDT switch. During the cathodic phase (ϕ_C), I_P delivers negative current through the electrode from V_{DD} to V_{MID}, and during the anodic phase (ϕ_A), I_N delivers positive current from V_{MID} to ground. Figs. B-1(b) and (c) show current-controlled stimulators that require just a single current source and uses a H-bridge to achieve bi-directional current. An advantage of using the same current source for both phases is better current matching at the cost of more series switches. In Fig. B-1(b), current always flows from V_{DD} to V_{MID} through I_P, while in Fig. B-1(c), V_{DD} is not required and current always flows from V_{MID} to ground through I_N (this is the approach used in Section 5.5.1). As discussed in Section 3.1.6, current source-based stimulators can be very inefficient at low current levels because most of the power is

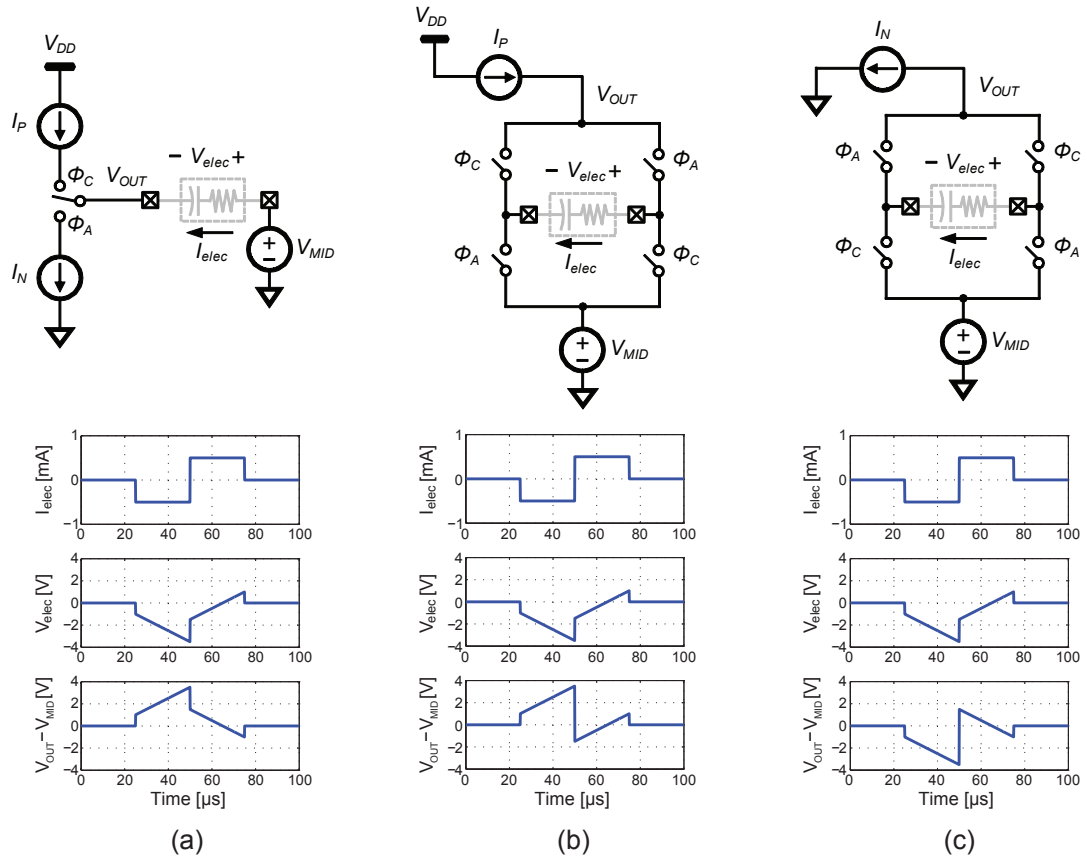

Figure B-1: Stimulator with a single power supply (V_{DD}) using (a) two current sources I_P and I_N, (b) a single current source I_P and a bridge, and (c) a single current source I_N and a bridge. The associated waveforms are shown below each schematic.

wasted in the headroom of the switches and current sources.

B.2 DC-DC Converter Approach

In order to address the poor efficiency of current source-based stimulators, a DC-DC converter-based stimulator using an inductor to recover energy from V_{MID} was designed in [97]. Figure B-2(a) shows a stimulator using a uni-directional power supply (with efficiency η) with a H-bridge to achieve bi-directional current. Power is delivered from an input voltage V_{IN} to V_{MID}. Since the power supply is uni-directional, an auxiliary DC-DC converter (shown in gray with efficiency η') can be used to return energy from V_{MID} to V_{IN}. If $V_{IN} > V_{OUT}$, then the power supply is a buck converter, and if $V_{IN} < V_{OUT}$, then it is a boost converter. Although not

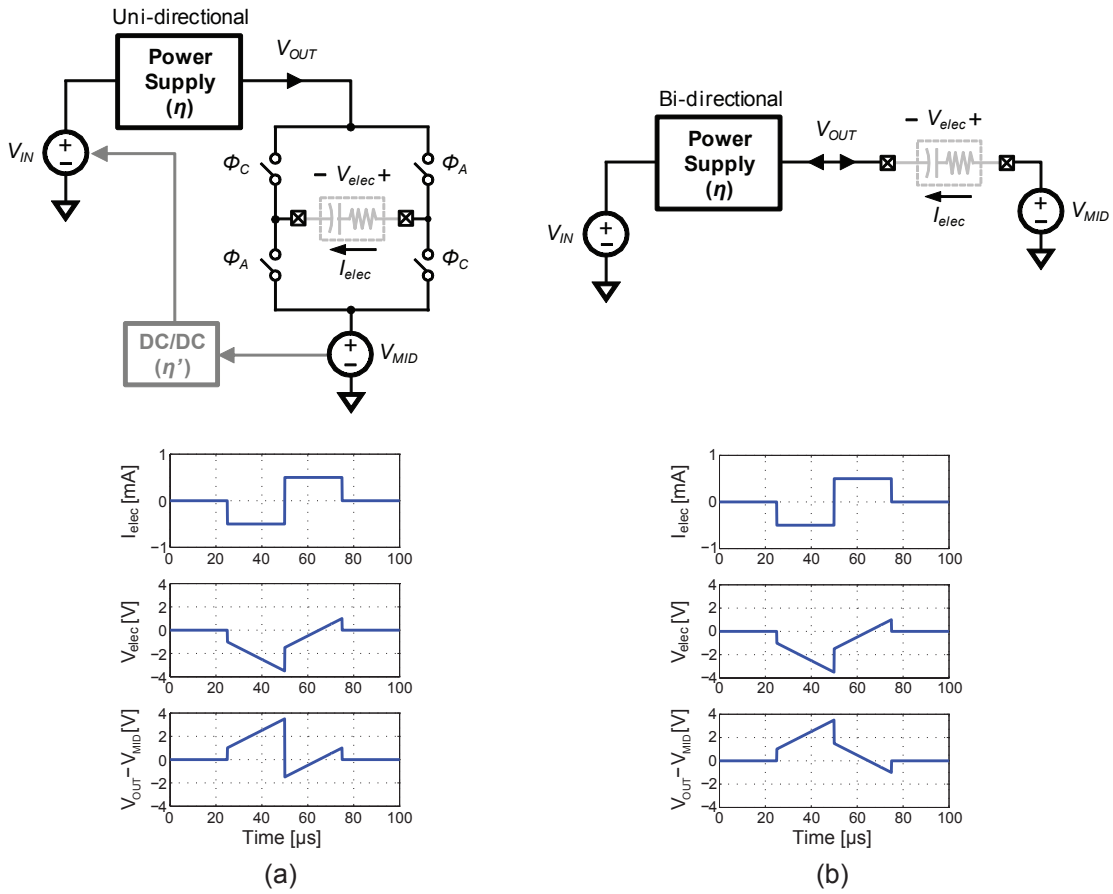

Figure B-2: (a) Stimulator using a uni-directional power supply and a bridge. (b) Stimulator using a bi-directional power supply. The associated waveforms are shown below each schematic.

shown, this approach requires feedback control in order to deliver the desired current waveform. During the stimulation phases, as the electrode voltage (V_{elec}) changes, the feedback control regulates V_{OUT} to a desired profile to maintain the desired current level.

Instead of a uni-directional power supply, the control logic can be modified to implement a bi-directional power supply as shown in Figure B-2(b). This is precisely the architecture presented in [97]. In this case, since the power supply can deliver power from V_{IN} to V_{MID}, and also return power from V_{MID} to V_{IN}, it does not require the H-bridge nor auxiliary DC-DC converter.

Figure B-3 shows simplified schematics of the bi-directional power supply for the

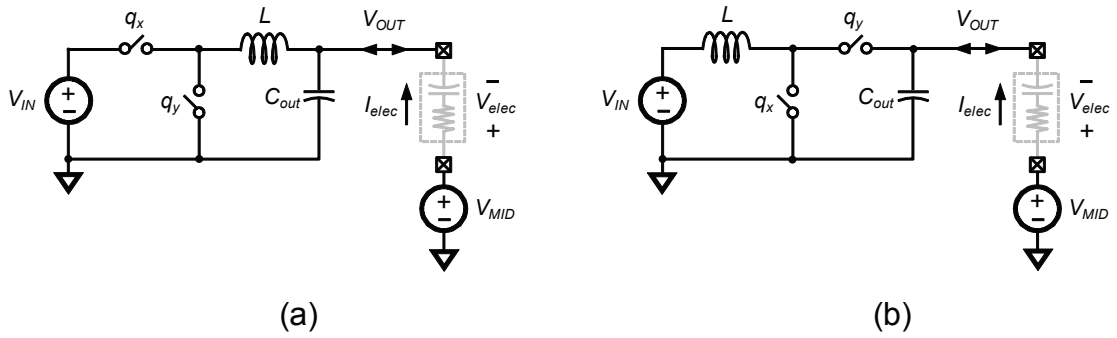

Figure B-3: Simplified schematic of a stimulator using a (a) forward buck (reverse boost) converter, and (b) forward boost (reverse buck) converter.

case of (a) $V_{IN} > V_{OUT}$ (as in [97]), and (b) $V_{IN} < V_{OUT}$. The control for switches q_x and q_y determines the direction of power flow for both cases. In Fig. B-3(a), during the cathodic phase (negative electrode current I_{elec}), the power supply is operating in forward mode as a buck converter to deliver power from V_{IN} to V_{MID} and the electrode capacitance. During the anodic phase (positive I_{elec}), the power supply is operating in reverse mode as a boost converter to return power from V_{MID} and the electrode capacitance to V_{IN}. Note that power is dissipated in the electrode resistance during both phases. Conversely, in Fig. B-3(b), during the cathodic phase, the power supply is in forward boost mode, while during the anodic phase, it is in reverse buck mode.

Intuitively, the DC-DC converter approach is more efficient than a current source approach because the inductor is used to return energy stored in V_{MID} and the electrode capacitance back to V_{IN}. With a current source approach, energy drawn from the supply is always dumped to ground. Therefore, for the DC-DC converter in periodic steady state, the energy drawn from V_{IN} is just the energy required to replenish energy lost from switching and conduction losses in the converter and electrode resistance.

Table B.1: Comparison of advantages and disadvantages for current source-based and DC-DC converter-based neural stimulators.

Stimulator approach	Advantages	Disadvantages
Current source	• Simple • Small and compact • Known charge balancing methods • Clean waveforms • Open loop, fast-settling	• Poor theoretical efficiency
DC-DC converter	• High theoretical efficiency	• Complex • Large size (require off-chip inductors and capacitors) • Poor charge control • Large current ripple • Feedback control (limited loop bandwidth)

B.3 Comparison of Stimulator Approaches

Table B.1 summarizes the advantages and disadvantages of both approaches. In general, the current source approach is simple, compact, has well established charge balancing methods, is able to provide clean waveforms, and can achieve fast settling because it is open loop. However, the main disadvantage is that it suffers from poor efficiency. Conversely, the DC-DC converter has high theoretical efficiency, but the design is more complex, it requires extra passives that may be too large to integrate, has poor charge control and large current ripple (depending on the switching frequency and control scheme), and relies on feedback control with limited loop bandwidth.

In theory, as discussed in Section 3.1.6, if the power supply efficiency $\eta = 100\%$, then the DC-DC converter approach requires no overhead power in the stimulation circuitry, and the only power drawn from V_{IN} is the power that is dissipated in the electrode resistance. However, practical values of η may be $<60\%$ when used as a stimulator for cochlear implant applications.

To understand why, let's consider the specifications provided in Table B.2. Assuming the circuit of Figure B-3(b) where $V_{IN} < V_{OUT}$, V_{IN} may be a low voltage supply of 2V to 3V available from a battery. V_{OUT} can typically range between 5V to 16V based on the stimulation current and expected values of electrode impedances

Table B.2: Example specifications for the DC-DC converter used in arbitrary-waveform stimulator applications.

Conversion ratio	V_{IN} = 2-3V, V_{OUT} = 5-16V
Load current	100μA to 1mA
Settling time	1μs

found in cochlear implants. Therefore, the DC-DC converter must be able to handle a large conversion ratio which may degrade its efficiency. Secondly, the desired load current ranges from 100μA to 1mA, which is extremely small. Therefore, the power budget for the control circuitry and other losses is extremely small. Lastly, in our work, we are interested in being able to provide arbitrary stimulation waveforms. For example, if each stimulation phase of 25μs is discretized into 10 time steps, then each time step is only 2.5μs long. This would require that the load current settle in approximately 1μs, which would require a switching frequency of at least 5MHz (e.g., 5 cycles per 1μs transition). Such a high switching frequency would result in large switching losses, further degrading the power supply efficiency.

In short, the specifications listed in Table B.2 are extremely difficult to achieve with high efficiency. Fundamentally, this is because the power supply is not being used like a DC-DC converter in the traditional sense. That is, instead of regulating the output voltage to a fixed value and delivering varying load current, the power supply when used as a stimulator in this work, has to dynamically and quickly change its output voltage to deliver the desired output current profile. To compound the problem, the load current ($<$ 1mA) range is extremely small. By looking at commercially available parts (e.g., Texas Instruments TPS61093 or Linear Technology LT1615) and academic literature, it is not unreasonable to assume that the maximum achievable efficiency is less than 60%, and more likely in the 30-50% range.

Considering practical values of the DC-DC converter efficiency for this stimulator application, we can now compare the overall efficiency of stimulation for both the current source and the DC-DC converter approaches. Following the efficiency analysis provided in [97] (which is done for biphasic rectangular waveforms), we can compare the efficiency of both approaches by defining an Energy Factor relative to a constant

Current Source stimulator (EF_{CS}). For the uni-directional power supply case shown in Figure B-2(a), the Energy Factor is given by

$$EF_{CS,uni} = \frac{E_{DCDC,uni}}{E_{CS}} = \frac{1 - \eta^2 + \alpha_R}{\eta}, \tag{B.1}$$

where $E_{DCDC,uni}$ is the energy consumed by the uni-directional power supply, and E_{CS} is the energy consumed by the current source. In this example, the efficiency of the auxiliary DC-DC converter that returns power from V_{MID} to V_{IN} is assumed to be equal to the efficiency of the main power supply (i.e., $\eta = \eta$').

For the bi-directional power supply case shown in Figure B-2(b), the Energy Factor (from reference [97]) is given by

$$EF_{CS,bi} = \frac{E_{DCDC,bi}}{E_{CS}} = \frac{1 - \eta^2}{2\eta}\left[1 + \frac{\alpha_C}{2} + \alpha_R \frac{(1 + \eta^2)}{(1 - \eta^2)}\right], \tag{B.2}$$

where $E_{DCDC,bi}$ is the energy consumed by the bi-directional power supply, $\alpha_C = V_C/V_{MID}$ and $\alpha_R = V_R/V_{MID}$, and V_C and V_R are the voltage drops across the electrode capacitance and resistance respectively. Based on this definition, a value of $EF_{CS} < 1$ indicates that the DC-DC converter approach is more energy efficient, while $EF_{CS} > 1$ suggests that the current source approach is more efficient.

From Equations B.1 and B.2, it can be seen that Energy Factor ratio is dependent on the power supply efficiency η, as well as α_R and α_C which depend on V_{MID}, the electrode impedance, the stimulation current amplitude, and the stimulation phase width. Assuming the typical parameter values listed in Table B.3[1], α_R can range from 0.1 to 0.5 at mid-scale current, and 0.4 to 0.8 at full-scale current. Similarly, α_C can range from 0.1 to 0.5 at mid-scale current, and 0.2 to 0.8 at full-scale current.

Figure B-4 plots the Energy Factor for a uni-directional power supply ($EF_{CS,uni}$) for typical values of α_R at (a) mid-scale and (b) full-scale stimulation current. At a mid-scale current of 500μA, the uni-directional power supply efficiency must be

[1]The electrode impedance and stimulation phase width in this work are very different from the values used in [97]. The work in [97] considered phase widths of 1ms, and an electrode model of approximately 0.93μF in series with 1kΩ.

Table B.3: Typical parameter values used to estimate α_R and α_C.

V_{MID}	6V to 8V
Electrode resistance	2kΩ to 5kΩ
Electrode capacitance	10nF
Phase width	10μs to 50μs
Mid-scale stimulation current	500μA
Full-scale stimulation current	1mA

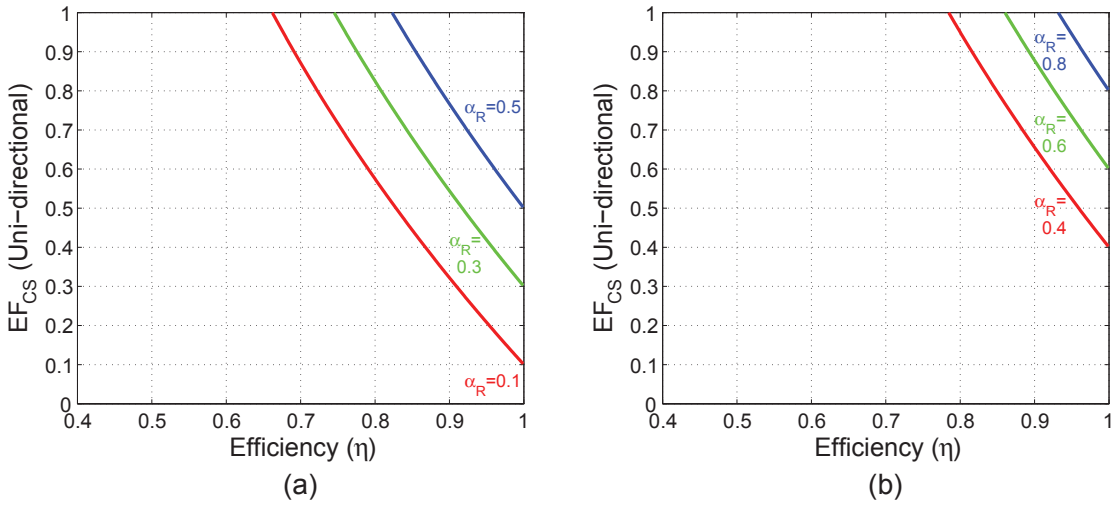

Figure B-4: Energy Factor as defined in [97] for typical values of α_R for a uni-directional power supply at (a) mid-scale and (b) full-scale stimulation current.

approximately 70% or greater in order for the DC-DC converter approach to be beneficial (i.e., $EF_{CS,uni} < 1$). At full-scale current, the requirement is increased to 80%. As discussed above, η is typically less than 60%. Therefore, for typical stimulation parameters found in cochlear implants, a current source-based stimulator is generally more efficient than using a uni-directional DC-DC converter.

If we consider the bi-directional power supply, the comparison with the current source approach is slightly more competitive. Figure B-5 plots $EF_{CS,bi}$ for typical values of α_R and α_C at (a) mid-scale and (b) full-scale stimulation current. At mid-scale, the bi-directional power supply efficiency has to be greater than approximately 50%, 60%, and 70% at $\alpha_R = 0.1$, 0.3, and 0.5 respectively. At full-scale, the requirement for η has to be even higher making it unlikely that the DC-DC converter approach will be beneficial. In general, this comparison suggests that the bi-directional DC-DC

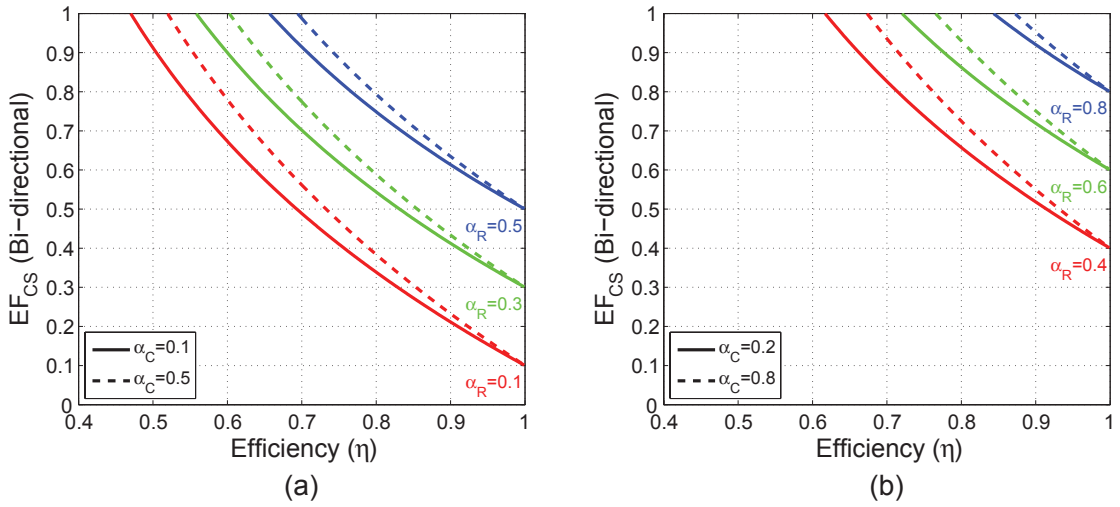

Figure B-5: Energy Factor as defined in [97] for typical values of α_R and α_C for a bi-directional power supply at (a) mid-scale and (b) full-scale stimulation current.

converter approach may be beneficial at lower current levels, but the benefits diminish or disappear altogether at higher current levels.

Therefore, although the DC-DC converter approach has high theoretical stimulation efficiency, its benefit over the current source approach is greatly diminished (especially at high current levels) by the typical stimulation parameters associated with cochlear implants, and the power supply efficiencies achievable in practice. Combined with the advantages and disadvantages listed in Table B.1, there is no clear answer as to which approach is the better solution, and the choice between using a current source or DC-DC converter approach is very much a matter of considering engineering tradeoffs.

Bibliography

[1] M. Yip and A. P. Chandrakasan, "A resolution-reconfigurable 5-to-10-bit 0.4-to-1 V power scalable SAR ADC for sensor applications," *IEEE J. Solid-State Circuits*, vol. 48, no. 6, pp. 1453–1464, Jun. 2013.

[2] E. M. Spinelli and M. A. Mayosky, "Two-electrode biopotential measurements: Power line interference analysis," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 8, pp. 1436–1442, Aug. 2005.

[3] Continua Health Alliance, [Online], Available: http://www.continuaalliance.org, Apr. 2011.

[4] G. Williams, K. Doughty, and D. A. Bradley, "A systems approach to achieving CarerNet - an integrated and intelligent telecare system," *IEEE Trans. Inf. Technol. Biomed.*, vol. 2, no. 1, pp. 1–9, Mar. 1998.

[5] H. Kim, R. F. Yazicioglu, T. Torfs, P. Merken, H. J. Yoo, and C. Van Hoof, "A low power ECG signal processor for ambulatory arrhythmia monitoring system," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, Jun. 2010, pp. 19–20.

[6] J. Kwong and A. P. Chandrakasan, "An energy-efficient biomedical signal processing platform," *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1742–1753, Jul. 2011.

[7] R. F. Yazicioglu, T. Torfs, J. Penders, I. Romero, H. Kim, P. Merken, B. Gyselinckx, H. J. Yoo, and C. Van Hoof, "Ultra-low-power wearable biopotential sensor nodes," in *Proc. IEEE EMBS*, Minneapolis, MN, Sep. 2009, pp. 3205–3208.

[8] N. Verma, A. Shoeb, J. Bohorquez, J. Dawson, J. Guttag, and A. P. Chandrakasan, "A micro-power EEG acquisition SoC with integrated feature extraction processor for a chronic seizure detection system," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 804–816, Apr. 2010.

[9] H. Kim, R. F. Yazicioglu, P. Merken, C. Van Hoof, and H. J. Yoo, "ECG signal compression and classification algorithm with quad level vector for ECG holter system," *IEEE Trans. Inf. Technol. Biomed.*, vol. 14, no. 1, pp. 93–100, Jan. 2010.

[10] R. F. Yazicioglu, S. Kim, T. Torfs, H. Kim, and C. Van Hoof, "A 30 μW analog signal processor ASIC for portable biopotential signal monitoring," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 209–223, Jan. 2011.

[11] F. Chen, A. P. Chandrakasan, and V. M. Stojanovic, "Design and analysis of a hardware-efficient compressed sensing architecture for data compression in wireless sensors," *IEEE J. Solid-State Circuits*, vol. 47, no. 3, pp. 744–756, Mar. 2012.

[12] D. B. Shire, S. K. Kelly, J. Chen, P. Doyle, M. D. Gingerich, S. F. Cogan, W. A. Drohan, O. Mendoza, L. Theogarajan, J. L. Wyatt, and J. F. Rizzo, "Development and implantation of a minimally invasive wireless subretinal neurostimulator," *IEEE Trans. Biomed. Eng.*, vol. 56, no. 10, pp. 2502–2511, Oct. 2009.

[13] F. G. Zeng, S. Rebscher, W. Harrison, X. Sun, and H. Feng, "Cochlear implants: System design, integration, and evaluation," *IEEE Reviews in Biomed. Eng.*, vol. 1, pp. 115–142, 2008.

[14] Proteus Digital Health, [Online], Available: http://www.proteusdigitalhealth.com, Jul. 2013.

[15] K. D. Kochanek, J. Xu, S. L. Murphy, A. M. Minino, and H. C. Kung, "Deaths: Final data for 2009," *National Vital Statistics Reports*, vol. 60, no. 3, Dec. 2011.

[16] American Heart Association Statistics Committee and Stroke Statistics Subcommittee, "Heart disease and stroke statistics - 2013 update," *Circulation: Journal of the American Heart Association*, pp. e6–e247, Dec. 2012.

[17] H. Zhou, K. M. Hou, L. Gineste, and C. De Vaulx, "A new system dedicated to real-time cardiac arrhythmias tele-assistance and monitoring," *Journal of Universal Computer Science*, vol. 12, no. 1, pp. 30–44, Jan. 2006.

[18] E. B. Bass, E. I. Curtiss, V. C. Arena, B. H. Hanusa, A. Cecchetti, M. Karpf, and W. N. Kapoor, "The duration of holter monitoring in patients with syncope: Is 24 hours enough?" *Archives of Internal Medicine*, vol. 150, no. 5, pp. 1073–1078, May 1990.

[19] G. Yang, J. Chen, F. Jonsson, H. Tenhunen, and L. R. Zheng, "A 1.0 V 78 μW reconfigurable ASIC embedded in an intelligent electrode for continuous remote ECG applications," in *Proc. IEEE EMBS*, Minneapolis, MN, Sep. 2009, pp. 2316–2319.

[20] J. Yoo, L. Yan, S. Lee, H. Kim, and H. J. Yoo, "A wearable ECG acquisition system with compact planar-fashionable circuit board-based shirt," *IEEE Trans. Inf. Technol. Biomed.*, vol. 13, no. 6, pp. 897–902, Nov. 2009.

[21] A. C. W. Wong, D. McDonagh, O. Omeni, C. Nunn, M. Hernadez Silveria, and A. J. Burdett, "Sensium: An ultra-low-power wireless body sensor network platform: Design & application challenges," in *Proc. IEEE EMBS*, Minneapolis, MN, Sep. 2009, pp. 6576–6579.

[22] L. Yan, J. Bae, S. Lee, T. Roh, K. Song, and H. J. Yoo, "A 3.9 mW 25-electrode reconfigured sensor for wearable cardiac monitoring system," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 353–364, Jan. 2011.

[23] INA333 Datasheet, [Online], Available: http://www.ti.com/lit/ds/symlink/ina333.pdf, 2008.

[24] ADS7866 Datasheet, [Online], Available: http://www.ti.com/lit/ds/symlink/ads7866.pdf, 2005.

[25] ADS1298 Datasheet, [Online], Available: http://www.ti.com/lit/ds/symlink/ads1298.pdf, 2012.

[26] R. Muller, S. Gambini, and J. M. Rabaey, "A 0.013 mm^2, 5 μW, DC-coupled neural signal acquisition IC with 0.5 V supply," *IEEE J. Solid-State Circuits*, vol. 47, no. 1, pp. 232–243, Jan. 2012.

[27] M. Yip, J. L. Bohorquez, and A. P. Chandrakasan, "A 0.6V 2.9μW mixed-signal front-end for ECG monitoring," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, Jun. 2012, pp. 66–67.

[28] N. Van Helleputte, S. Kim, H. Kim, J. P. Kim, C. Van Hoof, and R. F. Yazicioglu, "A 160μA biopotential acquisition ASIC with fully integrated IA and motion-artifact suppression," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2012, pp. 118–119.

[29] MSP430F1611 Datasheet, [Online], Available: http://www.ti.com/lit/ds/symlink/msp430f1611.pdf, 2011.

[30] TMS320C5504 Datasheet, [Online], Available: http://www.ti.com/lit/ds/symlink/tms320c5504.pdf, 2012.

[31] N. Ickes, Y. Sinangil, F. Pappalardo, E. Guidetti, and A. P. Chandrakasan, "A 10 pJ/cycle ultra-low-voltage 32-bit microprocessor system-on-chip," in *Proc. IEEE European Solid-State Circuits Conf.*, Helsinki, Finland, Sep. 2011, pp. 159–162.

[32] CC2550 Datasheet, [Online], Available: http://www.ti.com/lit/ds/symlink/cc2550.pdf, 2011.

[33] nRF24L01 Datasheet, [Online], Available: http://www.nordicsemi.com/eng/Products/2.4GHz-RF/nRF24L01, 2007.

[34] J. L. Bohorquez, A. P. Chandrakasan, and J. L. Dawson, "A 350 μW CMOS MSK transmitter and 400 μW OOK super-regenerative receiver for medical implant communications," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1248–1259, Apr. 2009.

[35] S. Drago, D. M. W. Leenaerts, F. Sebastiano, L. J. Breems, K. A. A. Makinwa, and B. Nauta, "A 2.4GHz 830pJ/bit duty-cycled wake-up receiver with -82dBm sensitivity for crystal-less wireless sensor nodes," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2010, pp. 224–225.

[36] A. Paidimarri, P. M. Nadeau, P. P. Mercier, and A. P. Chandrakasan, "A 440pJ/bit 1Mb/s 2.4GHz multi-channel FBAR-based TX and an integrated pulse-shaping PA," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, Jun. 2012, pp. 34–35.

[37] P. M. Nadeau, A. Paidimarri, P. P. Mercier, and A. P. Chandrakasan, "Multi-channel 180pJ/b 2.4GHz FBAR-based receiver," in *Proc. IEEE Radio Frequency Integrated Circuits Symp.*, Montreal, QC, Jun. 2012, pp. 381–384.

[38] P. P. Mercier and A. P. Chandrakasan, "A supply-rail-coupled eTextiles transceiver for body-area networks," *IEEE J. Solid-State Circuits*, vol. 46, no. 6, pp. 1284–1295, Jun. 2011.

[39] S. C. Jocke, J. F. Bolus, S. N. Wooters, A. D. Jurik, A. C. Weaver, T. N. Blalock, and B. H. Calhoun, "A 2.6-μW sub-threshold mixed-signal ECG SoC," in *Proc. IEEE Symp. VLSI Circuits*, Kyoto, Japan, Jun. 2009, pp. 60–61.

[40] D. J. Young, M. A. Zurcher, T. Trang, C. A. Megerian, and W. H. Ko, "Characterization of ossicular chain vibration at the umbo: implications for a middle ear microelectromechanical system design," *Ear, Nose and Throat Journal*, vol. 89, no. 1, pp. 21–26, Jan. 2010.

[41] J. Georgiou and C. Toumazou, "A 126-μW cochlear chip for a totally implantable system," *IEEE J. Solid-State Circuits*, vol. 40, no. 2, pp. 430–443, Feb. 2005.

[42] C. L. Searle, L. D. Braida, D. R. Cuddy, and M. F. Davis, "Binaural pinna disparity: another auditory localization cue," *J. Acoust. Soc. Am.*, vol. 57, no. 2, pp. 448–455, Feb. 1975.

[43] P. P. Mercier and A. P. Chandrakasan, "Rapid wireless capacitor charging using a multi-tapped inductively-coupled secondary coil," *IEEE Trans. Circuits Syst. I*, vol. 60, no. 9, pp. 2263–2272, Sep. 2013.

[44] W. H. Ko, R. Zhang, P. Huang, J. Guo, X. Ye, D. J. Young, and C. A. Megerian, "Studies of MEMS acoustic sensors as implantable microphones for totally implantable hearing-aid systems," *IEEE Trans. Biomed. Circuits Syst.*, vol. 3, no. 5, pp. 277–285, Oct. 2009.

[45] H. Zenner, M. Maassen, R. Lehner, J. Baumann, and H. Leysieffer, "An implantable hearing aid for inner ear hearing loss: Short-term implantation of microphone and transducer," *Otolaryng. Head Neck Surgery*, vol. 45, no. 10, pp. 872–880, Oct. 1997.

[46] Invisible Hearing® with Esteem®, [Online], Available: http://www.envoymedical.com, Apr. 2011.

[47] S. K. Kelly and J. L. Wyatt, Jr., "A power-efficient neural tissue stimulator with energy recovery," *IEEE Trans. Biomed. Circuits Syst.*, vol. 5, no. 1, pp. 20–29, Feb. 2011.

[48] W. Germanovix and C. Toumazou, "Design of a micropower current-mode log-domain analog cochlear implant," *IEEE Trans. Circuits Syst. II*, vol. 47, no. 10, pp. 1023–1046, Oct. 2000.

[49] R. Sarpeshkar, M. W. Baker, C. D. Salthouse, J. J. Sit, L. Turicchia, and S. M. Zhak, "An analog bionic ear processor with zero-crossing detection," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 78–79.

[50] R. Sarpeshkar, *Ultra Low Power Bioelectronics: Fundamentals, Biomedical Applications, and Bio-Inspired Systems*, 1st ed. Cambridge University Press, 2010.

[51] M. E. Sinangil, M. Yip, M. Qazi, R. Rithe, J. Kwong, and A. P. Chandrakasan, "Design of low-voltage digital building blocks and ADCs for energy-efficient systems," *IEEE Trans. Circuits Syst. II: Express Briefs*, vol. 59, no. 9, pp. 533–537, Sep. 2012.

[52] E. S. Winokur, M. K. Delano, and C. G. Sodini, "A wearable cardiac monitor for long-term data acquisition and analysis," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 1, pp. 189–192, Jan. 2013.

[53] V-Patch Medical Systems, [Online], Available: http://www.vpatch.com, Aug. 2013.

[54] AliveCor, [Online], Available: http://www.alivecor.com, Aug. 2013.

[55] P. R. Kinget, "Designing analog and RF circuits for ultra-low supply voltages," in *Proc. IEEE European Solid-State Circuits Conf.*, Munich, Germany, Sep. 2007, pp. 58–67.

[56] PowerStream Ultrathin Rechargeable Lithium Polymer Batteries, [Online], Available: http://www.powerstream.com/thin-lithium-ion.htm, 2013.

[57] B. H. Calhoun and A. P. Chandrakasan, "Characterizing and modeling minimum energy operation for subthreshold circuits," in *Proc. Int. Symp. Low Power Electronics and Design*, Newport Beach, CA, Aug. 2004, pp. 90–95.

[58] ADAS1000 Datasheet, [Online], Available: http://www.analog.com/static/imported-files/data_sheets/ADAS1000_1000-1_1000-2.pdf, 2013.

[59] J. L. Bohorquez, M. Yip, A. P. Chandrakasan, and J. L. Dawson, "A biomedical sensor interface with a sinc filter and interference cancellation," *IEEE J. Solid-State Circuits*, vol. 46, no. 4, pp. 746–756, Apr. 2011.

[60] R. F. Yazicioglu, C. Van Hoof, and R. Puers, *Biopotential Readout Circuits for Portable Acquisition Systems*, 1st ed. Springer Netherlands, 2009.

[61] J. G. Webster, *Medical Instrumentation: Application and Design*, 2nd ed. Boston, MA: Houghton Mifflin, 1992.

[62] *Medical electrical equipment - Part 2-47: Particular requirements for the safety, including essential performance of ambulatory electrocardiographic systems*, 1st ed. IEC International Standard 60601-2-47, 2001.

[63] *Cardiac monitors, heart rate meters, and alarms.* ANSI/AAMI EC13: American National Standards Institute/Association for the Advancement of Medical Instrumentation, 2002.

[64] ADS1222 Datasheet, [Online], Available: http://www.ti.com/lit/ds/symlink/ads1222.pdf, 2009.

[65] AD7799 Datasheet, [Online], Available: http://www.analog.com/static/imported-files/data_sheets/AD7798_7799.pdf, 2013.

[66] B. Murmann, "A/D converter trends: Power dissipation, scaling and digitally assisted architectures," in *Proc. IEEE Custom Integrated Circuits Conf.*, San Jose, CA, Sep. 2008, pp. 105–112.

[67] R. Sarpeshkar, "Universal principles for ultra low power and energy efficient design," *IEEE Trans. Circuits Syst. II: Express Briefs*, vol. 59, no. 4, pp. 193–198, Apr. 2012.

[68] A. M. Abo and P. R. Gray, "A 1.5-V, 10-bit, 14.3-MS/s CMOS pipeline analog-to-digital converter," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 599–606, May 1999.

[69] J. L. Bohorquez, "Digitally-assisted, ultra-low-power circuits and systems for medical applications," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, Feb. 2010.

[70] R. R. Harrison and C. Charles, "A low-power low-noise CMOS amplifier for neural recording applications," *IEEE J. Solid-State Circuits*, vol. 38, no. 6, pp. 958–965, Jun. 2003.

[71] M. van Elzakker, E. van Tuijl, P. Geraedts, D. Schinkel, E. Klumperink, and B. Nauta, "A 10-bit charge-redistribution ADC consuming 1.9 μW at 1 MS/s," *IEEE J. Solid-State Circuits*, vol. 45, no. 5, pp. 1007–1015, May 2010.

[72] A. Shikata, R. Sekimoto, T. Kuroda, and H. Ishikuro, "A 0.5V 1.1MS/sec 6.3fJ/conversion-step SAR-ADC with tri-level comparator in 40nm CMOS," in *Proc. IEEE Symp. VLSI Circuits*, Kyoto, Japan, Jun. 2011, pp. 262–263.

[73] H. Y. Tai, H. W. Chen, and H. S. Chen, "A 3.2fJ/c.-s. 0.35V 10b 100kS/s SAR ADC in 90nm CMOS," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, Jun. 2012, pp. 92–93.

[74] P. Harpe, E. Cantatore, and A. van Roermund, "A 2.2/2.7fJ/conversion-step 10/12b 40kS/s SAR ADC with data-driven noise reduction," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2013, pp. 270–271.

[75] C. Y. Liou and C. C. Hsieh, "A 2.4-to-5.2fJ/conversion-step 10b 0.5-to-4MS/s SAR ADC with charge-average switching DAC in 90nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2013, pp. 280–281.

[76] N. Kurosawa, H. Kobayashi, K. Maruyama, H. Sugawara, and K. Kobayashi, "Explicit analysis of channel mismatch effects in time-interleaved ADC systems," *IEEE Trans. Circuits Syst. I*, vol. 48, no. 3, pp. 261–271, Mar. 2001.

[77] B. P. Ginsburg and A. P. Chandrakasan, "An energy-efficient charge recycling approach for a SAR converter with capacitive DAC," in *Proc. IEEE Int. Symp. Circuits and Systems*, may 2005, pp. 184–187.

[78] L. Cordesses, "Direct digital synthesis: A tool for periodic wave generation (part 1)," *IEEE Sig. Proc. Mag.*, vol. 21, no. 4, pp. 50–54, Jul. 2004.

[79] E. B. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-29, no. 2, pp. 155–162, Apr. 1981.

[80] A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-Time Signal Processing*, 1st ed. Prentice Hall, 1999.

[81] M. S. J. Steyaert, W. M. C. Sansen, and C. Zhongyuan, "A micropower low-noise monolithic instrumentation amplifier for medical purposes," *IEEE J. Solid-State Circuits*, vol. SC-22, no. 6, pp. 1163–1168, Dec. 1987.

[82] R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 4, pp. 539–550, Apr. 1999.

[83] *NIDCD Fact Sheet: Cochlear Implants.* NIH Publication No. 11-4798, Mar. 2011.

[84] R. Sarpeshkar, C. Salthouse, J. J. Sit, M. W. Baker, S. M. Zhak, T. K. T. Lu, L. Turicchia, and S. Balster, "An ultra-low-power programmable analog bionic ear processor," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 4, pp. 711–727, Apr. 2005.

[85] I. Deligoz, S. R. Naqvi, T. Copani, S. Kiaei, B. Bakkaloglu, S. S. Je, and J. Chae, "A MEMS-based power-scalable hearing aid analog front end," *IEEE Trans. Biomed. Circuits Syst.*, vol. 5, no. 3, pp. 201–213, Jun. 2011.

[86] F. Serra-Graells, L. Gomez, and J. L. Huertas, "A true-1-V 300-μW CMOS-subthreshold log-domain hearing-aid-on-chip," *IEEE J. Solid-State Circuits*, vol. 39, no. 8, pp. 1271–1281, Aug. 2004.

[87] D. G. Gata, W. Sjursen, J. R. Hochschild, J. W. Fattaruso, L. Fang, G. R. Iannelli, Z. Jiang, C. M. Branch, J. A. Holmes, M. L. Skorcz, E. M. Petilli, S. Chen, G. Wakeman, D. A. Preves, and W. A. Severin, "A 1.1-V 270-μA mixed-signal hearing aid chip," *IEEE J. Solid-State Circuits*, vol. 37, no. 12, pp. 1670–1678, Dec. 2002.

[88] S. Kim, J. Y. Lee, S. J. Song, N. Cho, and H. J. Yoo, "An energy-efficient analog front-end circuit for a sub-1-V digital hearing aid chip," *IEEE J. Solid-State Circuits*, vol. 41, no. 4, pp. 876–882, Apr. 2006.

[89] A. A. Eshraghi, R. Nazarian, F. F. Telischi, S. M. Rajguru, E. Truy, and C. Gupta, "The cochlear implant: Historical aspects and future prospects," *The Anatomical Record*, vol. 295, no. 11, pp. 1967–1980, Nov. 2012.

[90] A. Wongsarnpigoon and W. M. Grill, "Energy-efficient waveform shapes for neural stimulation revealed with a genetic algorithm," *J. Neural Eng.*, vol. 7, no. 4, Aug. 2010.

[91] M. Sahin and Y. Tie, "Non-rectangular waveforms for neural stimulation with practical electrodes," *J. Neural Eng.*, vol. 4, no. 3, pp. 227–233, Sep. 2007.

[92] A. Wongsarnpigoon, J. P. Woock, and W. M. Grill, "Efficiency analysis of waveform shape for electrical excitation of nerve fibers," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 18, no. 3, pp. 319–328, Jun. 2010.

[93] R. D. Klafter and L. Hrebien, "An in vivo study of cardiac pacemaker optimization by pulse shape modification," *IEEE Trans. Biomed. Eng.*, vol. BME-23, no. 3, pp. 233–239, May 1976.

[94] S. Jezernik and M. Morari, "Energy-optimal electrical excitation of nerve fibers," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 4, pp. 740–743, Apr. 2005.

[95] M. G. Fishler, "Theoretical predictions of the optimal monophasic and biphasic defibrillation waveshapes," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 1, pp. 59–67, Jan. 2000.

[96] J. J. Sit, A. M. Simonson, A. J. Oxenham, M. A. Faltys, and R. Sarpeshkar, "A low-power asynchronous interleaved sampling algorithm for cochlear implants that encodes envelope and phase information," *IEEE Trans. Biomed. Eng.*, vol. 54, no. 1, pp. 138–149, Jan. 2007.

[97] S. K. Arfin and R. Sarpeshkar, "An energy-efficient, adiabatic electrode stimulator with inductive energy recycling and feedback current regulation," *IEEE Trans. Biomed. Circuits Syst.*, vol. 6, no. 1, pp. 1–14, Feb. 2012.

[98] S. K. Kelly, "A system for efficient neural stimulation with energy recovery," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, Oct. 2003.

[99] J. L. Wessale, L. A. Geddes, G. M. Ayers, and K. S. Foster, "Comparison of rectangular and exponential current pulses for evoking sensation," *Annals of Biomed. Eng.*, vol. 20, no. 2, pp. 237–244, 1992.

[100] T. J. Foutz and C. C. McIntyre, "Evaluation of novel stimulus waveforms for deep brain stimulation," *J. Neural Eng.*, vol. 7, no. 6, Dec. 2010.

[101] J. Simpson and M. Ghovanloo, "An experimental study of voltage, current, and charge controlled stimulation front-end circuitry," in *Proc. IEEE Int. Symp. Circuits and Systems*, 2007, pp. 325–328.

[102] J. C. Lilly, J. R. Hughes, E. C. Alvord, and T. W. Galkin, "Brief, noninjurious electric waveform for stimulation of the brain," *Science*, vol. 121, no. 3144, pp. 468–469, Apr. 1955.

[103] D. R. Merrill, M. Bikson, and J. G. R. Jefferys, "Electrical stimulation of excitable tissue: design of efficacious and safe protocols," *Journal of Neuroscience Methods*, vol. 141, pp. 171–198, Feb. 2005.

[104] D. R. McNeal, "Analysis of a model for excitation of myelinated nerve," *IEEE Trans. Biomed. Eng.*, vol. BME-23, no. 4, pp. 329–337, Jul. 1976.

[105] D. M. Whiten, "Electro-anatomical models of the cochlear implant," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, Feb. 2007.

[106] J. R. Schwarz and G. Eikhof, "Na currents and action potentials in rate myelinated nerve fibres at 20 and 37^{o}C," *European J. Phys.*, vol. 409, no. 6, pp. 569–577, Aug. 1987.

[107] J. H. M. Frijns, J. Mooij, and J. H. ten Kate, "A quantitative approach to modeling mammalian myelinated nerve fibers for electrical prosthesis design," *IEEE Trans. Biomed. Eng.*, vol. 41, no. 6, pp. 556–566, Jun. 1994.

[108] F. J. Vanpoucke, A. J. Zarowski, and S. A. Peeters, "Identification of the impedance model of an implanted cochlear prosthesis from intracochlear potential measurements," *IEEE Trans. Biomed. Eng.*, vol. 51, no. 12, pp. 2174–2183, Dec. 2004.

[109] C. Q. Huang, R. K. Shepherd, P. M. Carter, P. M. Seligman, and B. Tabor, "Electrical stimulation of the auditory nerve: Direct current measurement *in vivo*," *IEEE Trans. Biomed. Eng.*, vol. 46, no. 4, pp. 461–470, Apr. 1999.

[110] M. L. Hughes, K. R. Vander Werff, C. J. Brown, P. J. Abbas, D. M. R. Kelsay, H. F. B. Teagle, and M. W. Lowder, "A longitudinal study of electrode impedance, the electrically evoked compound action potential, and behavioral measures in Nucleus 24 cochlear implant users," *Ear & Hearing*, vol. 22, no. 6, pp. 471–486, Dec. 2001.

[111] P. A. Busby, K. L. Plant, and L. A. Whitford, "Electrode impedance in adults and children using the Nucleus 24 cochlear implant system," *Cochlear Implants Int.*, vol. 3, no. 2, pp. 87–103, Sep. 2002.

[112] B. Swanson, E. Van Baelen, M. Janssens, M. Goorevich, T. Nygard, and K. Van Herck, "Cochlear implant signal processing ICs," in *Proc. IEEE Custom Integrated Circuits Conf.*, San Jose, CA, Sep. 2007, pp. 437–442.

[113] R. K. Shepherd and E. Javel, "Electrical stimulation of the auditory nerve: II. effect of stimulus waveshape on single fibre response properties," *Hear. Res.*, vol. 130, no. 1-2, pp. 171–188, Apr. 1999.

[114] P. H. Gorman and J. T. Mortimer, "The effect of stimulus parameters on the recruitment characteristics of direct nerve stimulation," *IEEE Trans. Biomed. Eng.*, vol. BME-30, no. 7, pp. 407–414, Jul. 1983.

[115] J. T. Rubinstein, "An introduction to the biophysics of the electrically evoked compound action potential," *International Journal of Audiology*, vol. 43 Suppl. 1, pp. S3–S9, Dec. 2004.

[116] M. Barbara, V. Manni, and S. Monini, "Totally implantable middle ear device for rehabilitation of sensorineural hearing loss: preliminary experience with the esteem envoy," *Acta Oto-Laryngologica*, vol. 129, no. 4, pp. 429–432, Apr. 2009.

[117] K. Kroll, I. Grant, and E. Javel, "The Envoy totally implantable hearing system, St. Croix Medical," *Trends in Amplification*, vol. 6, no. 2, pp. 73–80, 2002.

[118] H. P. Zenner and J. Rodriguez Jorge, "Totally implantable active middle ear implants: ten years' experience at the University of Tubingen," *Adv. Otorhinolaryngol.*, vol. 69, pp. 72–84, 2010.

[119] D. J. Young, M. A. Zurcher, M. Semaan, C. A. Megerian, and W. H. Ko, "MEMS capacitive accelerometer-based middle ear microphone," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 12, pp. 3283–3292, Dec. 2012.

[120] W. T. Park, K. N. O'Connor, K. L. Chen, J. R. Mallon Jr., T. Maetani, P. Dalal, R. N. Candler, V. Ayanoor-Vitikkate, J. B. Roberson Jr., S. Puria, and T. W. Kenny, "Ultraminiature encapsulated accelerometers as a fully implantable sensor for implantable hearing aids," *Biomed. Microdevces*, vol. 9, no. 6, pp. 939–949, Dec. 2007.

[121] P. Huang, J. Guo, C. Megerian, D. J. Young, and W. H. Ko, "A laboratory study on a capacitive displacement sensor as an implant microphone in totally implant cochlear hearing aid systems," in *Proc. IEEE EMBS*, Aug. 2007, pp. 5691–5694.

[122] S. Nishihara, H. Aritomo, and R. L. Goode, "Effect of changes in mass on middle ear function," *Otolaryng. Head Neck Surgery*, vol. 109, no. 5, pp. 899–910, Nov. 1993.

[123] A. Maniglia, H. Abbass, T. Azar, M. Kane, P. Amantia, S. Garverick, W. H. Ko, W. Frenz, and T. Falk, "The middle ear bioelectronic microphone for a totally implantable cochlear hearing device for profound and total hearing loss," *The American Journal of Otology*, vol. 20, no. 5, pp. 602–611, 1999.

[124] A. Vujanic, R. Pavelka, N. Adamovic, C. Kment, S. Mitic, W. Brenner, and G. Popovic, "Development of a totally implantable hearing aid," in *Proc. 23rd Int. Conf. Microelectronics*, May 2002, pp. 235–238.

[125] A. Jenkins, J. Atkins, D. Horlbeck, M. Hoffer, B. Balough, J. V. Arigo, G. Alexiades, and W. Garvis, "U.S. phase I preliminary results of use of the Otologics MET fully-implantable ossicular stimulator," *Otolaryng. Head Neck Surgery*, vol. 137, no. 2, pp. 206–232, Aug. 2007.

[126] R. Perkins, J. P. Fay, P. Rucker, M. Rosen, L. Olson, and S. Puria, "The EarLens System: New sound transduction methods," *Hear. Res.*, vol. 263, no. 1-2, pp. 104–113, May 2010.

[127] R. L. Goode, G. Ball, and S. Nishihara, "Measurement of umbo vibration in human subjects - method and possible clinical applications," *The American Journal of Otology*, vol. 14, no. 3, pp. 247–251, May 1993.

[128] K. R. Whittemore, S. N. Merchant, B. B. Poon, and J. J. Rosowski, "A normative study of tympanic membrane motion in humans using a laser doppler vibrometer (LDV)," *Hear. Res.*, vol. 187, no. 1-2, pp. 85–104, Jan. 2004.

[129] J. J. Rosowski, H. H. Nakajima, and S. N. Merchant, "Clinical utility of laser-doppler vibrometer measurements in live normal and pathologic human ears," *Ear & Hearing*, vol. 29, no. 1, pp. 3–19, Jan. 2008.

[130] Piezo Systems Inc., "Piezo generator/sensor kit manual: KGS-006," 2011.

[131] M. E. Ravicz, J. J. Rosowski, and S. N. Merchant, "Mechanisms of hearing loss resulting from middle-ear fluid," *Hear. Res.*, vol. 195, no. 1-2, pp. 103–130, Sep. 2004.

[132] J. J. Rosowski, A. M. Huber, M. E. Ravicz, and R. L. Goode, "Are temporal bones useful models of human middle-ear mechanics?" in *Association for Research in Otolaryngology*, Feb. 2004.

[133] J. P. Beck, "US Patent Application No. 12/259,258: Implantable middle ear transducer having improved frequency response," Feb. 2010.

[134] K. Nie, A. Barco, and F. G. Zeng, "Spectral and temporal cues in cochlear implant speech perception," *Ear & Hearing*, vol. 27, no. 2, pp. 208–217, Apr. 2006.

[135] N. Hato, J. T. Welsh, R. L. Goode, and S. Stenfelt, "Acoustic role of the buttress and posterior incudal ligament in human temporal bones," *Otolaryng. Head Neck Surgery*, vol. 124, no. 3, pp. 274–278, Mar. 2001.

[136] F. G. Zeng, G. Grant, J. Niparko, J. Galvin, R. Shannon, J. Opie, and P. Segel, "Speech dynamic range and its effect on cochlear implant performance," *J. Acoust. Soc. Am.*, vol. 111, no. 1, pp. 377–386, Jan. 2002.

[137] P. Loizou, "Mimicking the human ear," *IEEE Signal Processing Magazine*, vol. 15, no. 5, pp. 101–130, Sep. 1998.

[138] B. S. Wilson and M. F. Dorman, "Cochlear implants: A remarkable past and a brilliant future," *Hear. Res.*, vol. 242, no. 1-2, pp. 3–21, Aug. 2008.

[139] J. K. Niparko, *Cochlear Implants Principles and Practices*, 2nd ed. Philadelphia, PA: Lippincott Williams and Wilkins, 2009.

[140] B. S. Wilson, C. C. Finley, D. T. Lawson, R. D. Wolford, D. K. Eddington, and W. M. Rabinowitz, "Better speech recognition with cochlear implants," *Letters to Nature*, vol. 352, pp. 236–238, Jul. 1991.

[141] R. V. Shannon, F. G. Zeng, V. Kamath, J. Wygonski, and M. Ekelid, "Speech recognition with primarily temporal cues," *Science*, vol. 270, no. 5234, pp. 303–304, Oct. 1995.

[142] P. Seligman and H. McDermott, "Architecture of the spectra 22 speech processor," *The Annals of Otology, Rhinology and Laryngology, Supplement*, vol. 166, pp. 139–141, Sep. 1995.

[143] K. E. Fishman, R. V. Shannon, and W. H. Slattery, "Speech recognition as a function of the number of electrodes used in the SPEAK cochlear implant speech processor," *J. Speech Lang. Hear. Res.*, vol. 40, no. 5, pp. 1201–1215, Oct. 1997.

[144] Q. J. Fu, R. V. Shannon, and X. Wang, "Effects of noise and spectral resolution on vowel and consonant recognition: Acoustic and electric hearing," *J. Acoust. Soc. Am.*, vol. 104, no. 6, pp. 3586–3596, Dec. 1998.

[145] W. Wattanapanitch, M. Fee, and R. Sarpeshkar, "An energy-efficient micropower neural recording amplifier," *IEEE Trans. Biomed. Circuits Syst.*, vol. 1, no. 2, pp. 136–147, Jun. 2007.

[146] R. Sarpeshkar, R. F. Lyon, and C. Mead, "A low-power wide-linear-range transconductance amplifier," *Analog Integrated Circuits Signal Process.*, vol. 13, no. 1-2, pp. 123–151, 1997.

[147] A. Wang and A. P. Chandrakasan, "A 180-mV subthreshold FFT processor using a minimum energy design methodology," *IEEE J. Solid-State Circuits*, vol. 40, no. 1, pp. 310–319, Jan. 2005.

[148] C. W. Turner, L. A. J. Reiss, and B. J. Gantz, "Combined acoustic and electric hearing: Preserving residual acoustic hearing," *Hear. Res.*, vol. 242, no. 1-2, pp. 164–171, Aug. 2008.

[149] L. M. Friesen, R. V. Shannon, D. Baskent, and X. Wang, "Speech recognition in noise as a function of the number of spectral channels: Comparison of acoustic hearing and cochlear implants," *J. Acoust. Soc. Am.*, vol. 110, no. 2, pp. 1150–1163, Aug. 2001.

[150] F. G. Zeng and R. V. Shannon, "Loudness balance between electric and acoustic stimulation," *Hear. Res.*, vol. 60, no. 2, pp. 231–235, Jul. 1992.

[151] B. Stobich, C. M. Zierhofer, and E. S. Hochmair, "Influence of automatic gain control parameter settings on speech understanding of cochlear implant users employing the continuous interleaved sampling strategy," *Ear & Hearing*, vol. 20, no. 2, pp. 104–116, Apr. 1999.

[152] Y. T. Kuo, T. J. Lin, Y. T. Li, and C. W. Liu, "Design and implementation of low-power ANSI S1.11 filter bank for digital hearing aids," *IEEE Trans. Circuits Syst. I*, vol. 57, no. 7, pp. 1684–1696, Jul. 2010.

[153] P. C. Loizou, M. Dorman, and Z. Tu, "On the number of channels needed to understand speech," *J. Acoust. Soc. Am.*, vol. 106, no. 4, pp. 2097–2103, Oct. 1999.

[154] K. S. Chong, B. H. Gwee, and J. S. Chang, "A 16-channel low-power nonuniform spaced filter bank core for digital hearing aids," *IEEE Trans. Circuits Syst. II*, vol. 53, no. 9, pp. 853–857, Sep. 2006.

[155] M. Ghovanloo and K. Najafi, "A compact large voltage-compliance high output-impedance programmable current source for implantable microstimulators," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 1, pp. 97–105, Jan. 2005.

[156] M. E. Halpern and J. Fallon, "Current waveforms for neural stimulation-charge delivery with reduced maximum electrode voltage," *IEEE Trans. Biomed. Eng.*, vol. 57, no. 9, pp. 2304–2312, Sep. 2010.

[157] Professor Philipos C. Loizou Group Website, [Online], Available: http://www.utdallas.edu/ loizou/cimplants/cdemos.htm, Nov. 2012.

[158] J. Ottenbacher, M. Kirst, L. Jatoba, M. Huflejt, U. Grossmann, and W. Stork, "Reliable motion artifact detection for ECG monitoring systems with dry electrodes," in *Proc. IEEE EMBS*, Vancouver, Canada, Aug. 2008, pp. 1695–1698.

[159] H. Kim, R. F. Yazicioglu, S. Kim, N. Van Helleputte, A. Artes, M. Konijnenburg, J. Huisken, J. Penders, and C. Van Hoof, "A configurable and low-power mixed signal SoC for portable ECG monitoring applications," in *Proc. IEEE Symp. VLSI Circuits*, Kyoto, Japan, Jun. 2011, pp. 142–143.

[160] D. A. Tong, K. A. Bartels, and K. S. Honeyager, "Adaptive reduction of motion artifact in the electrocardiogram," in *Proc. IEEE EMBS*, Houston, TX, Oct. 2002, pp. 1403–1404.

[161] Y. Liu and M. G. Pecht, "Reduction of skin stretch induced motion artifacts in electrocardiogram monitoring using adaptive filtering," in *Proc. IEEE EMBS*, New York, NY, Sep. 2006, pp. 6045–6048.

[162] K. N. Talbot and D. E. H. Hartley, "Combined electro-acoustic stimulation: a beneficial union?" *Clinical Otolaryngol.*, vol. 33, no. 6, pp. 536–545, Dec. 2008.

[163] J. J. Sit and R. Sarpeshkar, "A low-power blocking-capacitor-free charge-balanced electrode-stimulator chip with less than 6 nA DC error for 1-mA full-scale stimulation," *IEEE Trans. Biomed. Circuits Syst.*, vol. 1, no. 3, pp. 172–183, Sep. 2007.

[164] K. Sooksood, T. Stieglitz, and M. Ortmanns, "An active approach for charge balancing in functional electrical stimulation," *IEEE Trans. Biomed. Circuits Syst.*, vol. 4, no. 3, pp. 162–170, Jun. 2010.

[165] C. C. Enz and G. C. Temes, "Circuit techniques for reducing the effects of op-amp imperfections: Autozeroing, correlated double sampling, and chopper stabilization," *Proc. IEEE*, vol. 84, no. 11, pp. 1584–1614, Nov. 1996.

[166] J. C. Huhta and J. G. Webster, "60-Hz interference in electrocardiography," *IEEE Trans. Biomed. Eng.*, vol. BME-20, no. 2, pp. 91–101, Mar. 1973.

[167] B. B. Winter and J. G. Webster, "Reduction of interference due to common mode voltage in biopotential amplifiers," *IEEE Trans. Biomed. Eng.*, vol. BME-30, no. 1, pp. 58–62, Jan. 1983.

[168] R. E. Serrano, M. Gasulla, O. Casas, and R. Pallas-Areny, "Power line interference in ambulatory biopotential measurements," in *Proc. IEEE EMBS*, Cancun, Mexico, Sep. 2003, pp. 3024–3027.

[169] B. B. Winter and J. G. Webster, "Driven-right-leg circuit design," *IEEE Trans. Biomed. Eng.*, vol. BME-30, no. 1, pp. 62–66, Jan. 1983.

[170] TI Application Report: Thermal Noise Analysis in ECG Applications, [Online], Available: http://www.ti.com/lit/an/sbaa185/sbaa185.pdf, May 2011.

[171] D. Jiang, A. Demosthenous, T. A. Perkins, X. Liu, and N. Donaldson, "A stimulator ASIC featuring versatile management for vestibular prostheses," *IEEE Trans. Biomed. Circuits Syst.*, vol. 5, no. 2, pp. 147–159, Apr. 2011.