

**AN ELECTRONIC SYSTEM FOR EXTRACELLULAR NEURAL
STIMULATION AND RECORDING**

A Dissertation
Presented to
The Academic Faculty

by

Richard A. Blum

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Electrical and Computer Engineering

Georgia Institute of Technology
August 2007

AN ELECTRONIC SYSTEM FOR EXTRACELLULAR NEURAL
STIMULATION AND RECORDING

Approved by:

Stephen P. DeWeerth, Advisor
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Paul E. Hasler
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Robert H. Lee
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Mark G. Allen
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Bruce C. Wheeler
Department of Bioengineering
*University of Illinois at Urbana-
Champaign*

Date Approved: July 5, 2007

It is not up to you to complete the work, yet you are not free to desist from it.

—Pirkei Avot 2:21

To Cocoa, and to the memory of Serafina, Dijon, and Twinkle

ACKNOWLEDGEMENTS

During the course of my studies at Georgia Tech, I have benefited from interacting with many wonderfully talented people whose assistance and support has been essential to my work. I would like to take this opportunity to recognize those people who have helped me through my studies.

First, I would like to thank my parents and grandparents for instilling in me the importance of my education. I would like to extend this appreciation to my entire family, especially my sister Natalie Blum, for all their love and support.

I would like to thank everyone who worked on the 3-D Nets project, especially those involved in the development of electronics: Edgar Brown, James “Danger” Ross, Scott Buscemi, Samir Das, Kunal Gosrani, and Nakul Reddy. I would also like to thank Yoonkey Nam and Douglas Bakkum for conducting biological testing of the circuits developed as part of the project. Without their contributions, this work would not have been possible. I owe James Ross and Edgar Brown special gratitude for all their proofreading and editing of this document. Also, the National Institutes of Health supported much of this research through a Bioengineering Research Partnership grant (1 R01 EB00786-01).

I would like to thank Dr. Reid Harrison of the University of Utah for numerous discussions and for inspiration on the recording amplifier design.

I would like to thank the entire DeWeerth group, including past members, for their assistance and friendship: Edgar Brown, Stephan Clemens, Clinton Knight, J. Alex Bragg, Girish Patel, Mario Simoni, David Lin, Charles Wilson, Tina Hudson, Jason Meeks, Mara Carey, Joe Fernald, Michael Sorensen, Michael Reid, Kyla Ross, James Ross, Shane Migliore, Kartik Sundar, Carrie Williams, Kate Williams, Jevin Scrivens, JoAnna Todd, Scott Buscemi, Liang Guo, Samir Das, Bobby Brooke, Shawn O’Connor, Kunal Gosrani, and Nakul Reddy. I would especially like to thank Tina Hudson for first inviting me to the DeWeerth group meetings and Girish Patel and Charles Wilson for mentoring me. It has been a pleasure

working with these people, and I will always cherish the memories of lunches at Li'l Dino, racquetball, Monopoly, and so much more.

I would like to thank my committee: Dr. Stephen DeWeerth, my advisor; Dr. Paul Hasler; Dr. Mark Allen; Dr. Robert Lee; and Dr. Bruce Wheeler. Their wisdom and expertise has been essential to my work.

Additionally, I would like to thank all the faculty, and staff in the Laboratory for Neuro-engineering: Dr. Steve Potter, Dr. Lena Ting, Dr. Michelle LaPlaca, Dr. Robert Butera, and Dr. Ravi Bellamkonda, Amber Burriss, Bryan Williams, Jon Hall, and Jamie Lazin. I would also like to thank all the students in the Neurolab, and although there are far too many to be named here, I would like to mention Will Gerken, Murat Sekerli, Luke Purvis, Amanda Preyer, Nick Shapiro, Amanda Zimmerman, and Randy Weinstein as some of my many friends in the Neurolab.

I would also like to thank all my instructors over the years, including Dr. W. Marshall Leach, Jr.; Dr. John Peatman; Dr. Phillip Allen; Dr. George P. Burdell; and Dr. J. Alvin Connelly, to whom I owe thanks for the Texas Instruments Analog Fellows program that supported me through the master's program. I would also like to thank Marilou Mycko for the excellent administrative support that she provides to the School of Electrical and Computer Engineering.

I would like to thank all my brothers in the Alpha Epsilon Pi fraternity for the friendship and support that was invaluable to me over the years. I would especially like to thank my pledge brothers: David Gewertz, Alex Solodkin, Andrew Shutter, Howie Draisen, Adam Thompson, and Randy Weinstein. As both my fraternity brother and fellow Neurolab student, Randy Weinstein has provided me with so much advice, support, and doughnut runs over the years. I would also like to recognize some of my other close friends, including Tanah Barchichat, Jared Levy, Steve Raidbard, Brian Kime, Scot Humphreys, and Ben Dines.

Finally, and most importantly, I would like to thank my wonderful wife, Janna Blum, for all her support and encouragement that tremendously helped me through the difficult task of writing this dissertation.

TABLE OF CONTENTS

DEDICATION	iv
ACKNOWLEDGEMENTS	v
LIST OF TABLES	x
LIST OF FIGURES	xi
LIST OF SYMBOLS OR ABBREVIATIONS	xiv
SUMMARY	xv
1 INTRODUCTION	1
1.1 Development of Multi-Electrode Arrays	1
1.1.1 Planar Electrode Arrays	2
1.1.2 Non-planar Electrode Arrays	2
1.2 Development of Electronics for Neural Interfacing	3
1.2.1 Discrete Electronics	4
1.2.2 Integrated Circuits on Shared Substrates with MEAs	4
1.2.3 Integrated Circuits on Separate Substrates from MEAs	5
1.3 The Stimulation Artifact	5
1.4 A VLSI System for Multi-Electrode Stimulation and Recording	7
2 MODELING THE STIMULATION ARTIFACT	9
2.1 Overview of Electrode Theory	9
2.1.1 Electrode Classification	10
2.1.2 Electrical Models of Electrodes	11
2.1.3 Experimental Characterization of Electrode Impedance	17
2.2 Artifact Modeling	19
2.2.1 Nonlinear Model System	21
2.2.2 Linear Model System	24
2.2.3 Physical Test System	27
2.2.4 Comparison of Models of Stimulation Artifacts	27
2.3 Application of Models to Stimulation Artifact Elimination	29
2.4 Conclusions	30

3	INTEGRATED CIRCUITS FOR STIMULATION ARTIFACT ELIMINATION	33
3.1	Design of the First-Generation Integrated Circuit	33
3.1.1	Low-Noise Preamplifier	34
3.1.2	Stimulation Buffer	41
3.1.3	Artifact Elimination	43
3.1.4	Digital Control	49
3.2	Experimental Characterization of the First-Generation Integrated Circuit	51
3.2.1	Low-Noise Preamplifier	51
3.2.2	Stimulation Buffer	54
3.2.3	Artifact Elimination	55
3.3	Design of the Second-Generation Integrated Circuit	58
3.3.1	Preamplifier Noise Reduction	58
3.3.2	Increasing the Current Output of the Stimulation Buffer	59
3.3.3	Addition of Pole Shifting to the Artifact Elimination Circuitry . .	61
3.4	Experimental Characterization of the Second-Generation Integrated Circuit	63
3.4.1	Low-Noise Preamplifier	63
3.4.2	Stimulation Buffer	66
3.4.3	Application of Pole Shifting to Artifact Elimination	67
3.5	Conclusions	69
4	DIGITAL INTERFACES TO THE INTEGRATED CIRCUITS	70
4.1	Interfacing Analog Signals to the FPGA	71
4.1.1	Successive Approximation Analog-to-Digital Converter	72
4.1.2	Sigma-Delta Digital-to-Analog Converters	72
4.1.3	Data Converter Implementation and Testing	75
4.2	Signal Processing in FPGAs	78
4.2.1	Design of the Filtering System	81
4.2.2	Experimental Characterization	87
4.3	Conclusions	88
5	DISCUSSIONS	92
5.1	Novel Contributions of this Work	92

5.2	Summary of the Dissertation	93
5.2.1	Models of Stimulation Artifact Generation	93
5.2.2	Integrated Circuits for Elimination of the Stimulation Artifact	94
5.2.3	Filtering to Augment the Integrated Circuits	95
5.3	Future Work	96
APPENDIX A	SPICE MODELS FOR ELECTRODE SIMULATION	97
APPENDIX B	REVIEW OF ELECTRONIC NOISE	100
APPENDIX C	REVIEW OF THE WAVELET TRANSFORM	105
REFERENCES	111
VITA	117

LIST OF TABLES

2.1	Parameter Values for the Nonlinear Electrode Simulation	23
2.2	Model Parameters for the Linear Model SPICE Simulation	24
2.3	Model Parameters for the SPICE Diode Model Simulation	27
3.1	Transistor Sizes and Inversion Modes in the Recording System	39
3.2	Transistor Sizes in the Stimulation Buffer and Bias Network	43
3.3	Transistor Sizes in the Discharge Amplifier	44
3.4	Amplifier Activity During Operation Modes	47
3.5	Summary of Input Noise	53
3.6	Transistor Sizes in the Output Buffer	58
3.7	Transistor Sizes in the Revised Stimulation Buffer	61
3.8	Transistor Sizes in the Revised Feedback Amplifier	61
3.9	Summary of Input Noise in the Revised IC	65
4.1	Error Normalized to Signal Power	84
4.2	Wavelet De-noising Error, Normalized to Signal Power, with White Noise .	85

LIST OF FIGURES

2.1	Helmholtz layer	12
2.2	Circuit models of an electrode	16
2.3	Experimental setup for measuring electrode impedance	18
2.4	Photograph of an MEA after controlled platinum black deposition	19
2.5	Frequency–impedance plot of MEA electrodes	20
2.6	Simulink model system	22
2.7	Filtered artifacts from the Simulink model simulation	23
2.8	Linear electrode model system for simulation	24
2.9	Results of the SPICE simulation of linear electrode models	26
2.10	Charge stored on the electrode capacitance for the linear electrode model	26
2.11	Results of the SPICE simulation of electrode models with diodes for the charge-transfer resistance	28
2.12	Experimental system for generating stimulation artifacts	28
2.13	Artifacts generated by the physical system	29
2.14	Simulation of a post-stimulation discharge for artifact elimination	31
3.1	Block diagram of the IC	34
3.2	Topology of the recording amplifier	36
3.3	The feed-forward amplifier in the recording system.	37
3.4	Feedback amplifier	38
3.5	Schematic of the stimulation buffer	42
3.6	Bias network for the stimulation buffer	42
3.7	Discharge amplifier	45
3.8	Schematic of the recording, stimulation and artifact-elimination system	46
3.9	Linear analysis of the stability of the discharge loop	48
3.10	Simulated operation of the artifact-elimination circuitry	50
3.11	Die photograph	51
3.12	Frequency response of the recording system	52
3.13	Input-referred noise of the recording system	53
3.14	Frequency–impedance plot of the electrode	54

3.15	Current provided by the stimulation buffer	56
3.16	Effect of the artifact-elimination circuitry	57
3.17	Effect of the artifact-elimination circuitry with an external high-pass filter	57
3.18	Low output impedance buffer	59
3.19	Revised stimulation buffer	60
3.20	Revised feedback amplifier with pole shift	62
3.21	Frequency response of the revised IC	63
3.22	Input-referred spot noise of the new IC	64
3.23	Comparison of the NEF between the original and revised ICs	65
3.24	Stimulation current provided by the revised buffer	66
3.25	Effect of pole shifting	68
3.26	Artifacts for various stimulation current levels	68
4.1	Successive-Approximation ADC	73
4.2	First-order $\Sigma\Delta$ loop	75
4.3	Second-order $\Sigma\Delta$ loop	76
4.4	Comparison between first and second-order noise transfer functions	76
4.5	Sampling frequency constraint of the original ADC	77
4.6	Linear interpolation filter	79
4.7	Interpolation filter for the $\Sigma\Delta$ DAC	79
4.8	Power spectrum of the ADC–DAC cascade	80
4.9	Diagram of the wavelet filterbank for multi-resolution analysis	81
4.10	FFT of the symlet wavelet filters	82
4.11	Multi-resolution analysis applied to signals recorded from the first-generation IC	83
4.12	Simulated distortion of the wavelet filtering	85
4.13	Simulated distortion of the wavelet filtering for noisy signals	86
4.14	Block diagram of the FPGA filtering system	87
4.15	Test signal for the filtering algorithm	89
4.16	Comparison of the unfiltered, band-pass filtered, and wavelet-filtered signals	89
4.17	Effect of varying the wavelet coefficient threshold	90
4.18	Effect of the wavelet threshold on the noise levels and spike observations	91

B.1	Model circuit for deriving the thermal noise in a resistor	102
B.2	Models of noise in a resistor	102
B.3	Noise model of a MOSFET	104
C.1	Wavelet time-scale bins	106
C.2	Wavelet decomposition filterbank	110
C.3	Wavelet reconstruction filterbank	110

LIST OF SYMBOLS OR ABBREVIATIONS

ADC	Analog-to-Digital Converter.
CMOS	Complementary Metal Oxide Semiconductor.
DAC	Digital-to-Analog Converter.
DAQ	Data Acquisition System.
DSA	Dynamic Signal Analyzer.
FPGA	Field-Programmable Gate Array.
IC	Integrated Circuit.
LSB	Least Significant Bit.
MEA	Multi-Electrode Array.
MEMS	Micro-Electro-Mechanical Systems.
MOS	Metal Oxide Semiconductor.
MOSFET	Metal Oxide Semiconductor Field-effect Transistor.
MOSIS	Metal Oxide Semiconductor Implementation Service.
MSB	Most Significant Bit.
NEF	Noise Efficiency Factor.
NTF	Noise Transfer Function.
OTA	Operational Transconductance Amplifier.
PCB	Printed Circuit Board.
RHP	Right Half-plane.
RMS	Root-Mean-Square.
SNR	Signal-to-Noise Ratio.
SPI	Serial Peripheral Interface.
STF	Signal Transfer Function.
TSMC	Taiwan Semiconductor Manufacturing Company Ltd.
VLSI	Very Large Scale Integration or Integrated.
WRA	Wide-range Amplifier.

SUMMARY

We presented a system for extracellular neural interfacing that had the capability for stimulation and recording at multiple electrodes. As the core of this system, we designed a custom integrated circuit (IC) that contained low-noise amplifiers, stimulation buffers, and artifact-elimination circuitry. The artifact-elimination circuitry was necessary to prevent the activity of the stimulation buffers from interfering with the normal functioning of the low-noise amplifiers. As an aid in the design of the artifact-elimination circuitry, we developed models of the generation of the stimulation artifact, and we compared the models against physically generated artifacts.

We fabricated our integrated circuits in a $0.35\ \mu\text{m}$ CMOS process. We measured input-referred noise levels for the amplifiers as low as $3.50\ \mu\text{V}_{\text{rms}}$ in the in the bandwidth $30\ \text{Hz}$ – $3\ \text{kHz}$, corresponding to the frequency range of neural action potentials. The power consumption was $120\ \mu\text{W}$, corresponding to a noise–efficiency factor of 14.5. We we able to resume recording signals within $2\ \text{ms}$ of a stimulation, using the same electrode for both stimulation and recording.

After the activity of the artifact-elimination circuitry, a post-discharge artifact remained. We designed a filtering algorithm to remove the post-discharge artifact, and we implemented the filtering with a field-programmable gate array (FPGA). To connect the IC to the FPGA, we designed and built analog-to-digital converters (ADCs) using a mix of the computational resources of the FPGA and off-the-shelf analog components. The ADCs had an effective resolution of 10 bits. The filtering algorithm itself consisted of blanking for the duration of the stimulation and artifact-elimination, followed by a wavelet de-noising. The wavelet de-noising split the signal into frequency ranges, discarded those ranges that did not correspond to neural signals, applied a threshold to the retained signals, and recombined the different frequency ranges into a single signal. The combination of the filtering with the artifact-elimination IC resulted in the capability for artifact-free recordings.

CHAPTER 1

INTRODUCTION

The ability to observe the effects of stimulating neural tissue is essential to many scientific and engineering endeavors. Specific examples of applications for neural stimulation include studies of neural development and plasticity, clinical treatment of epilepsy and Parkinson's disease, retinal and cochlear implants, and the design of biosensors (Ruaro et al., 2005). Among the most important problems in neuroscience is that of understanding how networks of neurons develop over time and change in response to stimuli. Because of the interest in this field, many scientific studies of neural development and plasticity focus on the spatio-temporal dynamics of neural activity (DeMarse et al., 2001; Jimbo et al., 1999; Martinoia et al., 2005; van Pelt et al., 2004). Although neurons are complex electrochemical systems, they encode a large portion of the information that they process in quick voltage transients known as *action potentials* (Bialek et al., 1991; Hodgkin and Huxley, 1952). Because of the importance of neural electrical activity, combined with the limitless possibilities of electronics, observation of the electrical activity of neural tissue is one of the primary methods for determining the behavior and connectivity of the neural tissue.

1.1 Development of Multi-Electrode Arrays

Electrodes establish the link between bioelectrical signals and engineered, electronic systems, permitting the electronics to not only record the neural activity but also to alter it. Neuroscientists depend on electrodes and the associated electrical circuitry as essential technologies for a variety of experimental studies on neuronal plasticity and development. Although there are many studies of the neural dynamics that require only one or two electrodes, studying the system-level dynamics of neural tissue often requires many electrodes, separated by distances on the order of cellular dimensions, in order to observe the spatio-temporal dynamics of the electrical activity. This need for a large number of electrodes

has driven the fabrication of multi-electrode arrays (MEAs) through micromachining techniques similar to those that produce integrated circuits (ICs) or micro-electro-mechanical systems (MEMS). These techniques allow the fabrication of electrode arrays, consisting of hundreds of micron sized electrodes at spacings of tens to hundreds of microns. The fine spatial resolution of MEAs, combined with their long term biocompatibility, makes them optimally suited for studies of neural network development and plasticity.

1.1.1 Planar Electrode Arrays

Development of MEAs began in 1972, with the fabrication of an array of 30 metal electrodes, each $50\ \mu\text{m}^2$, etched onto a glass coverslip (Thomas et al., 1972). The electrode grid pattern consisted of two rows, $50\ \mu\text{m}$ apart, with a column spacing of $100\ \mu\text{m}$. This prototype MEA was capable of recording from cardiac cells that produced signal amplitudes in the range of $20\ \mu\text{V}$ – $2\ \text{mV}$. The electrodes were able to withstand stimulation currents up to $100\ \mu\text{A}$ without alteration of the impedance levels. Another early MEA, consisting of $12\ \mu\text{m}$ wide photoetched gold electrodes, was capable of recording 300 – $500\ \mu\text{V}$ signals from snail ganglia (Gross, 1979). Although these MEAs were important developments, they did not demonstrate the capability to record the lower-amplitude signals generated by vertebrate neural tissue.

An important development in MEAs was that of Pine, who, in 1980 recorded from dissociated neural cultures using an MEA consisting of 32 gold electrodes ($80\ \mu\text{m}^2$ size, $250\ \mu\text{m}$ spacing), obtaining recordings from cells within $40\ \mu\text{m}$ of the electrode centers (Pine, 1980). As that was the first use of MEAs to record from vertebrate neurons, it signaled a new era in neuroscience research. Since that development, planar MEAs have become a common tool for the study of the electrical activity of neural tissue, with many commercial varieties available.

1.1.2 Non-planar Electrode Arrays

An alternative to planar MEAs are those that consist of microfabricated towers. This approach began with an array of gold electrodes on a silicon substrate that was capable of recording action potentials from the cortex (Wise et al., 1970). Similar designs were

developed since that initial one (Hoogerwerf and Wise, 1994; Najafi et al., 1977; Nordhausen et al., 1996). These three-dimensional structures offered the possibility of electrode arrays with even more recording sites than possible with planar arrays. Additionally, because *in vivo* neural tissue is a three-dimensional network, the addition of height to the electrode arrays provided the ability to more accurately investigate the full spatial dependence of neural dynamics.

These MEMS structures also influenced the growth of neural tissue. More recently, the integration of fluidic channels into the electrodes added the ability to provide nutrients and growth factors and remove waste products (Choi et al., 2007; Cullen et al., 2007; Rowe et al., 2005). The addition of microfluidics created an environment closer to *in vivo* conditions, increasing the longevity of *in vitro* samples and making their dynamics more like those of *in vivo* tissue.

1.2 Development of Electronics for Neural Interfacing

Complexity and pervasive parallelization characterize both the nervous system and electronic systems—similarities that suggested neuroscientists could benefit from employing electronics as tools in the study of neuronal growth and development. Early MEAs relied on standard, readily available electronics for recording and stimulation; however, the increasing sophistication of MEA structures and experimental complexity demanded the use of specially designed electronics. Although these electronic systems have taken on many different forms; from ICs sharing a substrate with the electrodes to printed circuit boards (PCBs) with carefully chosen, standard ICs; they all have similar functions and design requirements.

Common to all these systems is the need for low-noise preamplifiers. The indirect connections inherent to extracellular electrodes result in attenuated signals in the range of 20–100 μV_{pp} (Claverol-Tinture and Pine, 2002; Pine, 2003). To record such small signals, the preamplifier must introduce less than 4 μV_{rms} of input-referred noise. The preamplifier must amplify the small input signals to a level above the noise floor of the data acquisition system (DAQ), as well as providing impedance transformation from the input impedance

of the electrodes to a low-impedance output.

Complementing recording functionality, many systems include stimulation amplifiers, adding experimental capability for two way communication with the neural culture. The capability for stimulation is essential for many applications, including studies of neural plasticity, clinical brain stimulation for the treatment of epilepsy and Parkinsons’s disease, sensory implants, biosensors, and closed-loop systems that neurons as controllers.

1.2.1 Discrete Electronics

One approach to electronic system design is to continue the historical precedent of system assembly from off-the-shelf components. This was the basis of the most popular commercial systems (Multi Channel Systems; Plexon). Academic researchers have also developed add-on interfaces to the commercial hardware to provide the capability for stimulation at arbitrary electrodes (Wagenaar and Potter, 2004).

1.2.2 Integrated Circuits on Shared Substrates with MEAs

Another approach to connecting electronics to MEAs is to combine circuits and electrodes onto a common substrate. An important advantage that this method offers is the ease of interconnecting the electrodes and circuits. Early efforts in this direction resulted in an array with ten recording sites and electronics with capability for amplification, multiplexing, and buffering (BeMent et al., 1986; Najafi and Wise, 1986). Later designs extended the size of the MEA to 32 electrodes, with the capability to select eight as active recording sites (Ji and Wise, 1992). More recent designs allowed for stimulation by bypassing the recording amplifiers, allowing recording as soon as 1 ms after stimulation and within 20 μm of the stimulation electrode (Olsson et al., 2005). These designs were effective, but they required the use of a custom MEMS/metal–oxide–semiconductor (MOS) fabrication technology that limited the size and complexity of the electronics.

Another approach to creating shared substrate MEA/amplifier systems is to modify a standard complementary metal–oxide–semiconductor (CMOS) technology to create electrodes out of the top layer metal, resulting in array elements containing the electrode and electronic circuits. An example of this was the fabrication of a 128×128 array in a 0.5 μm

CMOS process that was capable of measuring signals as small as $100 \mu\text{V}_{\text{pp}}$ from invertebrate neurons, although the noise level prohibited measurement of vertebrate neural activity (Eversmann et al., 2003).

1.2.3 Integrated Circuits on Separate Substrates from MEAs

Another approach is to connect separate MEAs and ICs together. This results in flexibility in the selection of ICs and MEAs: one IC can function in a system with a variety of different MEAs, from standard commercial MEAs to prototypes from academic research groups. An example of this type of design was an IC with 16 instrumentation amplifiers and stimulation circuitry in a $2.0 \mu\text{m}$ process (Pancrazio et al., 1998). This design had input-referred noise of $12\text{--}16 \mu\text{V}_{\text{rms}}$ in a 50 kHz bandwidth, which permitted recordings from neural tissue. Another design included 32 instrumentation amplifiers and output multiplexing in a $0.7 \mu\text{m}$ process, with input referred noise of $3 \mu\text{V}_{\text{rms}}$ (Dabrowski et al., 2004).

Harrison presented a neural amplifier design that optimized the transistor sizes of the signal path so that only the input pair introduced significant noise (Harrison and Charles, 2003). To prevent the dc offsets from causing amplifier saturation, the amplifier had an adaptive element in its feedback path, adding a high-pass pole (Delbrück and Mead, 1994). Among its advantages, this amplifier used very little power and die area while introducing noise at levels similar to larger amplifiers that consumed more power. Because of its efficient use of power and die area, it has been used as the basis for large arrays (Aziz et al., 2007; Harrison et al., 2007).

1.3 *The Stimulation Artifact*

Recording alone is insufficient to investigate neuronal behavior or the development of neural connectivity because many applications also require electrical input to the neural culture. Ideally, the experimenter should have the capability to switch the functionality of any electrode between stimulation and recording (Pancrazio et al., 1998; Wagenaar and Potter, 2004); however, an effect known as the *stimulation artifact* interferes with such flexibility by causing localized interference with recording for tens of milliseconds after stimulation (Mayer et al., 1992). The presence of the stimulation artifact has limited the study of neural

development and plasticity.

The stimulation artifact is a result of the properties of extracellular interfacing. Large signal losses are associated with extracellular recordings. The extracellular electrodes do not measure membrane potentials directly; rather, they record the electric field induced by ionic channel currents. This electric field decreases with distance from the cell, so the voltages present at the electrode are in the microvolt range, even though membrane potentials are in the millivolt range. Signal loss also takes place in the reverse path, so that extracellular stimulation requires voltages at the electrode that are many orders of magnitude larger than those due to cellular electrical activity (Pine, 1980). The stimulation voltages overwhelm the sensitive recording system, creating the stimulation artifact.

The undesired effects of the stimulation artifact has driven the development of methods for mitigating its interference, usually at the expense of functionality. In the simplest method, the system has designated electrodes that function as either stimulation or recording sites for the duration of the experiment, thus sidestepping the problem of recording at the site of the largest artifacts. Often, electronics designers placed sample and hold (S/H) circuitry at the input of the recording amplifier to prevent saturation of the electronic system during stimulation (Grumet et al., 2000; Novak and Wheeler, 1988). Another common technique blanked, or disabled, recording amplifiers near stimulation sites for up to 10 ms after stimulation (O’Keeffe et al., 2001). Many techniques focused on post-processing to filter out stimulation artifacts from neighboring electrodes (Gnadt et al., 2003; Wagenaar and Potter, 2002). These approaches all conceded the data closest to the stimulation, both temporally and spatially, as lost to the stimulation artifact; however, these data may represent the most significant response to the stimulation.

The difficulty of recording from saturated amplifiers raises the need to physically reduce the artifact itself. Dedicated physical circuitry is necessary to suppress the artifact. Jimbo et al. presented a design that brought the stimulation electrode back to its pre-stimulation voltage, which was kept in a S/H, immediately after stimulation (Jimbo et al., 2003). This method provided an effective stimulation while minimizing the artifact, both at neighboring electrodes and at the stimulation electrode. A limitation of this design was that the

area and power requirements rendered it unsuitable for very large scale integration (VLSI) technology—an important requirement as the natural scalability of VLSI systems keeps pace with growing sizes of MEAs.

1.4 A VLSI System for Multi-Electrode Stimulation and Recording

Although the wide variety of neural interfacing technology has enabled a wide variety of research, there are still opportunities for new technology to enhance experimental capability. Among the chief limitations of present systems are the interference from the stimulation artifact and the difficulties in assembling large-scale systems from off-the-shelf components. In light of the state of neural interfacing, a VLSI neural interfacing system capable of recording and stimulation in multiple electrodes, while preventing interference from the stimulation artifact, would be a significant contribution towards helping scientists investigate neural development and plasticity. In this work, we present the design, implementation, and characterization of such a system.

Our use of VLSI technology is important, because it promises that our system can scale with the increasing size of MEAs. As the development of three-dimensional electrodes continues, the advantage of interfacing systems built using VLSI technology over those that use off-the-shelf components becomes more pronounced.

The most distinguishing feature of our system is the method in which we deal with the stimulation artifact. We use novel circuitry to eliminate the stimulation artifact at its source, the electrode. Through the use of this circuitry, we can resume normal recording activity after stimulation after a much shorter duration than systems that rely on filtering techniques. Our system gives scientists the capability to deliver stimulation pulses at any electrode and still record neural action potentials after a brief duration. Through our novel device, we advance the possibilities for scientific inquiry, providing a significant contribution to the field of bioinstrumentation.

We begin the presentation of our system with an investigation into the origin of the stimulation artifact (Chapter 2). Developing a usable model of the generation of stimulation artifacts is crucial to the design of the artifact-elimination circuitry. After developing such a

model, we continue to the design and testing of the IC for neural stimulation and recording, including the capability for artifact-elimination (Chapter 3). Among the criteria that we use to evaluate the design are the noise levels, the stimulation currents, and the time necessary for artifact elimination. Following the presentation of the IC, we consider the augmentation of the IC with digital filtering to improve the performance of the artifact-elimination circuitry (Chapter 4). The use of digital circuitry also provides a means for future development of high-throughput data transfer from the IC to a computer system, which will aid in storage and analysis of the neural activity.

CHAPTER 2

MODELING THE STIMULATION ARTIFACT

Studies involving electrical stimulation and recording of *in vitro* neural cultures show promise of revealing how neuronal networks develop and respond to stimuli. These studies, however, must overcome numerous technical challenges. One common problem is the presence of the stimulation artifact, a long-lasting transient effect that obscures neural activity after stimulation. Because of stimulation artifacts, experimenters cannot observe the immediate, local response to stimulation. The design of an artifact elimination system that will alleviate the problems associated with stimulation artifacts is one of the major goals of this thesis.

In order to design artifact-elimination circuitry, we must first understand the nature of the stimulation artifact. In this chapter, we will discuss the stimulation artifact, its causes, and compensation methods. We will start our investigation with a summary of the physical properties of extracellular electrodes. From these properties, we will create computational models of electrodes, and show that these models are capable of producing stimulation artifacts. Armed with the knowledge of the generation of the artifacts, we will suggest methods to eliminate the artifact at the electrode.

2.1 Overview of Electrode Theory

Although neural activity is, to a large extent, an electrical phenomenon, establishing a connection between cells and electronics is considerably more complex than simply connecting a wire. A fundamental difference between biology and engineered systems is that the two employ different charge carriers. Bioelectrical currents are ionic, that is, they consist of chemical species moving in solution. Electrodes provide the necessary transduction of ionic currents into electrical currents, connecting bioelectrical activity and electronic systems to each other. Because of the critical role that electrodes play in a bio-instrumentation system, understanding their properties is a prerequisite for system design.

2.1.1 Electrode Classification

Electrodes come in a variety of physical forms, although they fall into two broad categories, *intracellular* and *extracellular* electrodes. The distinction between the two types of electrodes depends on the method of physical contact with electrically active cells. Intracellular electrodes puncture the cell membrane, making direct contact with the cell interior; extracellular electrodes, in contrast, contact only the cell exterior or extracellular medium. Not surprisingly, the differences between these two types result in differences in the electrical coupling that the electrodes make with cells. Intracellular electrodes directly measure the membrane voltage of the cells, while extracellular electrodes measure the membrane potential indirectly through its induced electric fields.

Just as the signals that the two types of electrodes can measure are different, the optimal electrical properties vary for intracellular and extracellular electrodes. There are two different idealizations of the electrical properties of an electrode. *Polarizable electrodes* are capable of sustaining an arbitrary dc voltage in the absence of current flow, and the other type, *non-polarizable electrodes*, permit dc current flow and their voltage is dependent on the current flow. The linear circuit model for an ideal polarizable electrode is a capacitor, for an ideal non-polarizable electrode, a resistor. Non-polarizable electrodes are naturally suited to intracellular measurements, which require high input impedance to prevent loading of the membrane capacitor. Similarly, polarizable electrodes are well suited for extracellular measurements, which require low noise (Gesteland et al., 1959).

Because this thesis concerns the design of instrumentation for the study of network plasticity and development, consideration must be made of the types of electrodes that neuroscientists employ for such studies. Plasticity and development are long term effects, so their study requires the ability to conduct chronic experiments, lasting weeks or months (Potter and DeMarse, 2001). Conducting such experiments requires care to ensure cell survival, precluding the use of electrode structures that causes acute injury to neurons; thus, intracellular electrodes, which make destructive transmembrane contact with neurons, are unsuitable for chronic studies, despite their high signal quality. Neuroscientists must, therefore, resort to extracellular electrodes, which do not cause cellular injury. Due to this requirement, we will

only consider extracellular electrodes in this work.

2.1.2 Electrical Models of Electrodes

In order to facilitate the design and simulation of electronics for neural interfacing, electrical models of electrode behavior are necessary. Development of models requires an analysis of the mechanisms by which transduction between electronic and ionic currents occurs. The processes that accomplish the transduction fall into two broad categories: charge transfer and capacitive coupling.

2.1.2.1 Capacitive Coupling in Electrodes

Capacitive coupling is the dominant effect in noble metal electrodes. The capacitive properties of the electrode are due to the formation of a layer of neutral water molecules that separates aqueous ions from the electrode. When the electrode makes contact with the solution, electrochemical reactions result in a voltage difference across the electrode–electrolyte interface, known as the *half-cell potential*. The electric field of the charged electrode acts on the polar water molecules, which form a hydration sheath (known as the *inner Helmholtz plane*) around the electrode. The hydration sheath limits the proximity of the hydrated ions to the electrode, and the closest approach of the ions to the electrode is known as the *outer Helmholtz plane*. Figure 2.1 gives a diagram of the structure of the hydration sheath. The structure of the hydration sheath, with an insulating region separating two conductors, is analogous to a capacitor.

In capacitive coupling, as in actual capacitors, equal but opposite amounts of charge accumulate on either side of the interface. A current on one side of the electrode (either biology or electronics) changes the charge on that side, inducing a corresponding change in the charge on the other side. Because no exchange of electrons across the electrode–electrolyte interface occurs in capacitive current flow, it does not affect the electrochemical equilibrium of the interface.

A first-order circuit model for capacitive coupling considers the hydration sheath as a parallel plate capacitor, using the distance to the outer Helmholtz plane, d_{OHP} , as the separation between the two conductors (the electrode and the electrolytic medium). This




Figure 2.1: Helmholtz layer, after Kovacs (1994). The hydration sheath limits the approach of the hydrated ions to the electrode, resulting in a structure similar to a parallel-plate capacitor

capacitance is known as the *Helmholtz capacitance*, and it follows the formula

$$C_H = \frac{\epsilon_r \epsilon_0 A}{d_{\text{OHP}}} \quad (2.1)$$

where ϵ_r is the relative permittivity of the medium, ϵ_0 is the permittivity of free space (8.854 pF/m), and A is the electrode area (Helmholtz).

Although an ideal capacitor is a linear device, the capacitance of the electrode–electrolyte interface is non-linear. As the voltage across the interface changes, the hydration sheath expands and contracts. The voltage-dependent thickness of the hydration sheath corresponds to a voltage-dependent capacitance. The circuit model for the capacitive effects must account for the voltage dependent nature of the space charge layer. The model consists of the

series combination of C_H with the voltage-dependent *Gouy–Chapman capacitance*, C_D :

$$\frac{1}{C_e} = \frac{1}{C_H} + \frac{1}{C_D} \quad (2.2)$$

where C_e is the total electrode capacitance, C_H is the voltage independent Helmholtz capacitance, and C_D is the voltage dependent Gouy–Chapman capacitance (Chapman, 1913; Gouy, 1910; Grahame, 1947). The Gouy–Chapman capacitance, which models the redistribution of ions due to electrode voltage, follows

$$C_D = \frac{\varepsilon_r \varepsilon_0 A}{L_D} \cosh\left(\frac{z\eta}{2U_T}\right) \quad (2.3)$$

where z is the charge on the ions in solution; η is the overpotential, the difference between the applied voltage and the equilibrium voltage, V_0 ; U_T is the thermal voltage ($U_T \equiv kT/q = 24.99$ mV at $T = 290$ K); and L_D is the *Debye length*. The Debye length, which is the space constant over which perturbations in potential decay, is

$$L_D = \sqrt{\frac{\varepsilon_r \varepsilon_0 U_T}{2n^0 z^2 q}} \quad (2.4)$$

where n^0 is the ionic concentration, and q is the charge on an electron (1.602×10^{-19} C) (Franks et al., 2005; Kovacs, 1994). At equilibrium voltage, $\frac{\partial C_D}{\partial V} = 0$, so the Gouy–Chapman capacitance does not contribute to a first-order, linear model of capacitive coupling.

An additional nonlinearity in the polarization capacitance is that it does not follow the frequency–reactance relationship of a true capacitor. Empirically, the interface capacitance is a constant-phase element, with impedance

$$Z_{CPA} = \frac{1}{(j\omega Q)^n} \quad (2.5)$$

where Q is a measure of the impedance magnitude, and $n \neq 1$ represents the deviation from ideal capacitive behavior (Franks et al., 2005; Fricke, 1932; McAdams et al., 1995).

The formulas (2.1) and (2.3) assume that the electrode is a smooth, uniform surface. In practice the surface is often irregular, resulting in an increase in the true surface area of the electrode. The electrical effect of surface roughness is to increase the interface capacitance (Daikhin et al., 1996; de Levie, 1965). In an approximation, the effects of surface

roughness may be accounted for by an approximation factor,

$$S = \phi A \tag{2.6}$$

where ϕ is the roughness factor that relates the true surface area, S , to the drawn area, A (Feltham and Spiro, 1971).

In many cases, it is desirable to intentionally apply surface treatments to increase the surface roughness of the electrode; doing so will lower the electrode impedance and improve the coupling between the electronics and neurons. Deposition of platinum black onto the electrode is especially common (Feltham and Spiro, 1971; Marrese, 1987). By controlling the duration of the electroplating procedure, it is possible to control the final electrode impedance precisely (Ross et al., 2004).

2.1.2.2 Charge Transfer in Electrodes

The second category of transduction between electronic and biological currents that occurs at the electrode is charge transfer. In contrast to capacitive coupling, in which electrons do not flow across the electrode–electrolyte interface, charge transfer involves the transfer of electrons, mediated by electrochemical reactions.

Under equilibrium conditions, a balance exists between oxidation and reduction reactions at the electrode, and although no net current flows through the electrode, there are equal, oppositely directed oxidation and reduction currents. As the electrode voltage, V , deviates from equilibrium, redox reactions take place that establish a current through the electrode of

$$I = J_o A \left(\exp \left[\frac{z(1-\beta)\eta}{U_T} \right] - \exp \left[\frac{-z\beta\eta}{U_T} \right] \right) \tag{2.7}$$

where J_o is the equilibrium current density and β is a symmetry factor that relates to the kinetics of the oxidation and reduction reactions (Kovacs, 1994). At equilibrium voltage $\eta = 0$, the oxidization and reduction currents each take on a value of $I_0 = J_o A$, resulting in no net current flow.

The form of (2.7) is similar to the ideal diode law,

$$I = I_s \left(e^{\frac{nV}{U_T}} - 1 \right) \tag{2.8}$$

suggesting that parallel, oppositely-oriented diodes suitably model the *charge transfer resistance* (Geddes et al., 1987). For small voltage deviations and symmetric reactions, a linear resistor of

$$R_t = \frac{U_T}{zI_0} \quad (2.9)$$

represents the effects of charge transfer.

2.1.2.3 Other Effects in Electrodes

A more complete model of the electrode must account for a variety of effects beyond capacitive coupling and charge transfer. Two of the most notable of these additional effects are the *Warburg impedance* and the *spreading resistance*. The Warburg impedance accounts for the limits of ionic diffusion in the medium by adding an element in series with the charge transfer resistance (Warburg, 1899). This impedance obeys

$$|Z_w| \propto \frac{1}{\sqrt{f}} \quad (2.10)$$

and

$$\angle Z_w = \frac{\pi}{4}, \quad (2.11)$$

resulting in a constant phase element.

A final effect requiring consideration is the spreading resistance, which is due to the resistance of the electrolyte. Because this is a measure of the resistance seen by a current leaving the electrode, the geometry of the electrode influences the value of the spreading resistance. Typical formulas for spreading resistance include

$$R_S = \frac{\rho}{4r} \quad (2.12)$$

for a circular electrode, where ρ is the electrolyte resistivity and r is the electrode radius, and

$$R_S = \frac{\rho \ln(4)}{\pi a} \quad (2.13)$$

for a square with sides of length a (Kovacs, 1994).

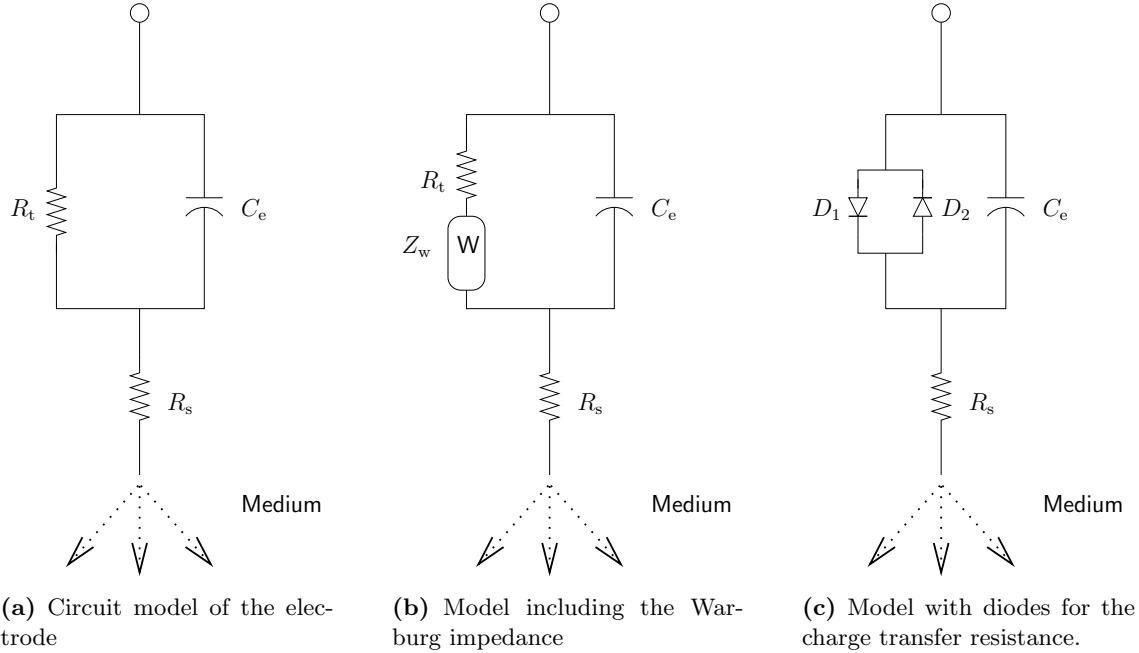


Figure 2.2: Circuit models of an electrode, consisting of a electrode–electrolyte capacitance, charge-transfer resistance, and spreading resistance. In general, the components in model (c) are non-linear. Model (b) adds an additional element to model the diffusion of ions in the medium. Model (c) replaces the non-linear charge-transfer resistance with diodes, whose non-linearity is easily modeled by circuit simulation programs.

2.1.2.4 Circuit Models of Electrodes

The circuit model of the electrode combines the separate effects of Section 2.1.2 into a single model. The Randels model represents the electrode as a parallel combination of the charge-transfer resistance and the interface capacitance, in series with the spreading resistance (Figure 2.2(a)) (Randels, 1947). The Warburg impedance may be included in series with the charge-transfer resistance (Figure 2.2(b)) (Kovacs, 1994).

In general, the components in these models are non-linear. For ease of incorporation into standard circuit analysis programs, linearized component values may be used. Modeling the charge-transfer resistance as two diodes (Figure 2.2(c)) preserves some of the non-linear electrode behavior without compromising the ease of using standard circuit analysis software.

2.1.3 Experimental Characterization of Electrode Impedance

Experimentally, electrodes are characterized by the magnitude of impedance as a function of frequency. One typical measurement technique places a known resistor in series with the electrode and measures the voltage across both elements (see Figure 2.3). By voltage division,

$$V_R = V_{in} \frac{R}{R + Z_{elec}} \quad (2.14)$$

and

$$V_{elec} = V_{in} \frac{Z_{elec}}{R + Z_{elec}}. \quad (2.15)$$

Combining (2.14) and (2.15) gives

$$Z_{elec} = R \frac{V_{elec}}{V_R}, \quad (2.16)$$

which may be separated into magnitude and phase,

$$|Z_{elec}| = R \frac{|V_{elec}|}{|V_R|} \quad (2.17)$$

$$\angle Z_{elec} = \angle V_{elec} - \angle V_R. \quad (2.18)$$

To avoid resistive loading, instrumentation amplifiers are used to measure the voltage across the resistor and the electrode because of their high-impedance, differential inputs. An SR785 dynamic signal analyzer (DSA) automates the frequency sweep, data collection, and math.

Using the impedance measurement circuitry, we characterized electrodes on two microfabricated electrode arrays, or multi-electrode arrays (MEAs) from Ayanda Biosystems. Both MEAs consisted of 60 planar electrodes of 30 μm diameter and 100 μm spacing. One MEA consisted of gold electrodes, the other, platinum. To provide a variety of electrode properties for testing purposes, some electrodes on the gold MEA had platinum black electroplated onto the electrode surface while others remained untreated (see Figure 2.4).

The impedance magnitude plot (Figure 2.5) shows that the untreated platinum and gold electrodes had similar impedances at the frequency range relevant to extracellular recordings

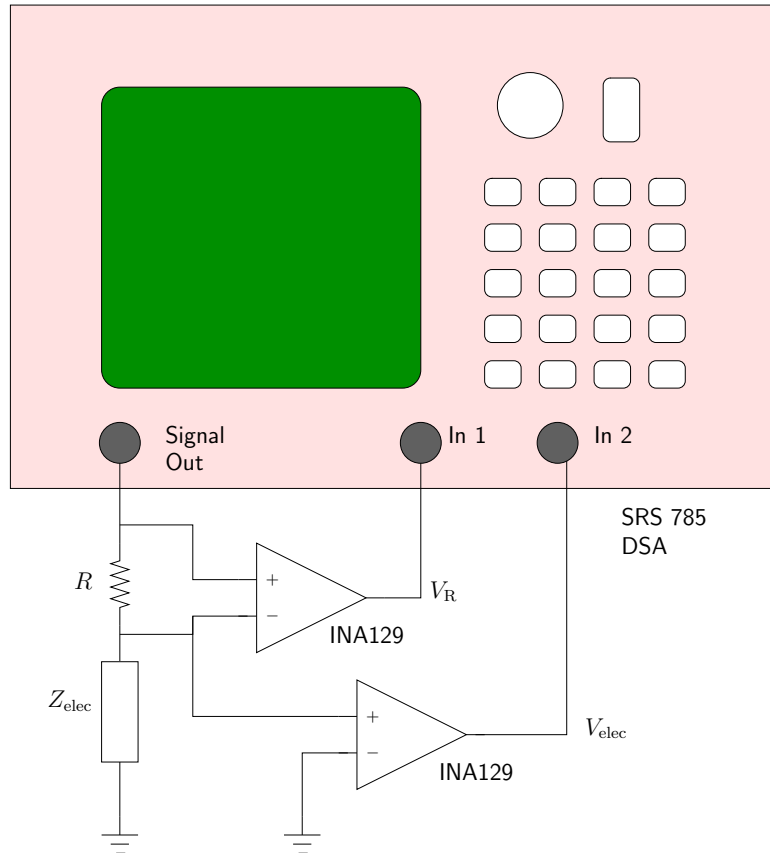


Figure 2.3: Experimental setup for measuring electrode impedance. This method compares the voltage drop across a known resistor to that across the electrode under test. The instrumentation amplifiers are necessary to prevent loading of the circuit. The dynamic signal analyzer automates sweeping the frequency of the sinusoidal input voltage.




Figure 2.4: Photograph of an MEA after controlled platinum black deposition. The treated electrodes were visibly darker than the untreated ones.

(300 Hz–3 kHz). This was consistent with the theoretical expectation that the impedance was dominated by capacitive coupling, which is dependent on electrode surface area. Measurements on the platinum black electrodes yielded a much lower impedance compared to the untreated electrodes, as expected by the increase in true surface area associated with their rough surface.

2.2 Artifact Modeling¹

The electrodes that are used in extracellular interfacing are primarily capacitive; thus, it is reasonable to hypothesize that we may understand the stimulation artifact as a phenomenon associated with the capacitive interface. As the state of a capacitor is determined by its stored charge, so might the artifact be generated by the charging of the electrode–electrolyte interface during stimulation. A good stimulation buffer is capable of providing large currents to quickly modify the charge stored on the interface, and any change in the charge from its equilibrium value would persist after the end of stimulation. With the stimulation buffer disconnected, the only discharge path remaining is through the charge transfer resistance, which is typically on the order of tens of megaohms. The discharge through the large

¹Portions of this section and the next section are from Blum et al. (2004)




Figure 2.5: Impedance magnitudes for the MEA. Most of the platinum black electrodes had impedances in a very narrow range (black traces), except for two outliers. The gold (red) and platinum (blue) electrodes also had similar impedances. The dashed, gray trace was the open-circuit limit of the measurement equipment.

impedance is capable of generating long-lasting transients like those observed.

In order to verify that trapped charge from stimulation causes the stimulation artifact, we studied model stimulation systems. We used two different models of the electrode. The first model used the detailed, non-linear equations for the behavior of the electrode. For the second model, we used the linear circuit approximation. In both of these systems, the model stimulation source was an ideal voltage source in series with a time dependent resistance that modeled the stimulation switch. We also considered a physical electrode with a stimulator constructed from off-the-shelf components. We compared the results of the two computational models to the physical system.

2.2.1 Nonlinear Model System

The most detailed model that we studied considered the detailed nonlinearities of the electrode (Borkholder, 1998). For implementation of the nonlinear model system, we used Simulink, a dynamic system simulator. The model (Figure 2.6) considered the electrode as an interface capacitance, $C_e(V)$, charge transfer current source, $I_T(V)$, and spreading resistance, R_s (Ross, 2003). The nonlinear components followed the formulas given in (2.2–2.7). We represented the stimulation circuitry with an ideal voltage source in series with a variable resistor that took on either a small resistance, which represented a closed stimulation switch, or a large resistance, which represented an open switch. As an addition to the model, there was a second-order, band-pass filter (100 Hz–30 kHz) that represented a typical recording preamplifier. Table 2.1 gives the parameter values used in the simulation, which were a mix of typical parameters available in literature (Borkholder, 1998) and parameters chosen to match the electrode impedance measurements.

The nonlinear simulations produced artifacts that were qualitatively similar to the linear model (Figure 2.7(a)). An additional step was to apply a 100 Hz–30 kHz bandpass filter to the electrode voltage. This simulated the artifacts that would be recorded by a typical bio-instrumentation system. The filtered artifacts (Figure 2.7(b)) were much smaller than the unfiltered artifacts, but would still have been capable of saturating a high-gain amplifier for over 10 ms.

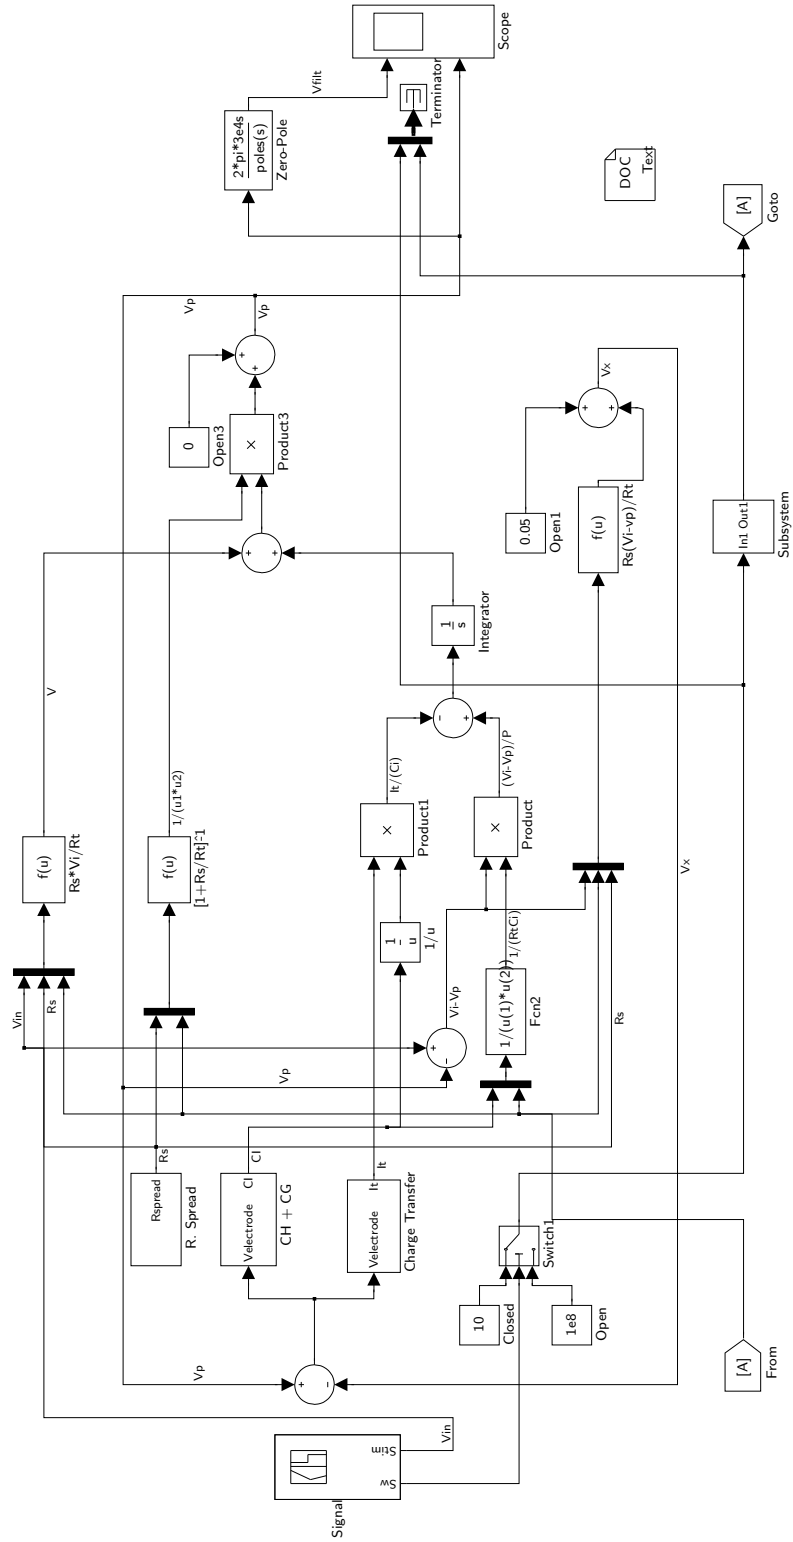


Figure 2.6: Simulink system for the nonlinear electrode model, from Ross (2003).

Table 2.1: Parameter Values for the Nonlinear Electrode Simulation, based on Borkholder (1998) and experimental data (Figure 2.5)

Parameter	Value			Units
	Platinum Black	Platinum	Gold	
r	15.0	15.0	15.0	μm
ϕ	27.0	1.0	1.0	—
ρ	73.0	73.0	73.0	$\Omega \cdot \text{cm}$
L_D	7.8	7.8	7.8	\AA
n^0	9.3×10^{22}	9.3×10^{22}	9.3×10^{22}	cm^{-3}
z	1.0	1.0	1.0	—
d_{OHP}	5.0	5.0	5.0	\AA
ϵ_r	78.5	78.5	78.5	—
V_0	50.0	50.0	50.0	mV
J_o	30.0	60.0	110.0	$\mu\text{A}/\text{cm}^2$
β	0.5	0.5	0.5	—

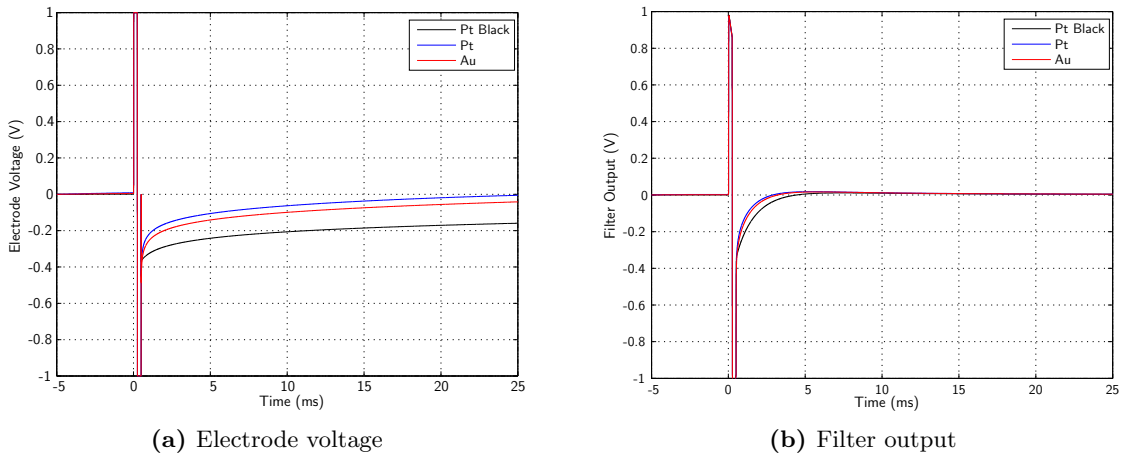


Figure 2.7: Artifacts produced by the nonlinear Simulink model. In the electrode voltages (a) we observed long-lasting stimulation artifacts. When band-pass filtered at 100 Hz–30 kHz (b), the artifacts decayed much more quickly, although they were still large enough to have saturated a high-gain amplifier for over 15 ms.

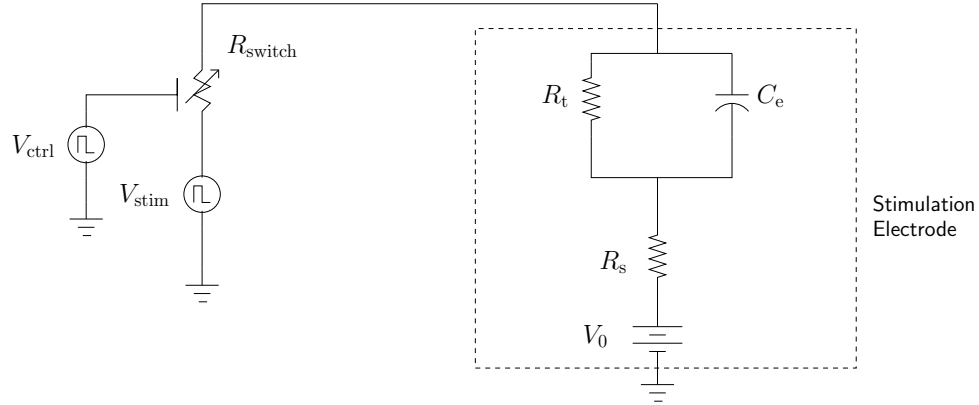


Figure 2.8: Linear circuit model of the electrode. The electrodes were represented as parallel RC circuits. The stimulation switch model was a resistor that took on a low value for a closed switch and a high value for an open switch.

Table 2.2: Model Parameters for the Linear Model SPICE Simulation, based on experimental data (Figure 2.5)

Parameter	Value	Units		
$R_{\text{switch,closed}}$	100	Ω		
$R_{\text{switch,open}}$	1	$G\Omega$		
R_s	10	$k\Omega$		
	Platinum Black	Platinum	Gold	
V_0	50	50	50	mV
C_e	8000	300	300	pF
R_t	30	70	150	$M\Omega$
τ_{stim}	88	3.3	3.3	μs
τ_{disch}	240	21	45	ms

2.2.2 Linear Model System

After our consideration of the nonlinear model, we considered a simpler, linear model that used an RC circuit for the electrode. In addition to the electrodes, the model system included a time-dependent switch resistance, R_{switch} ; a stimulation electrode; and spreading resistance (see Figure 2.8). The V_{stim} source was a biphasic pulse, with a duration of $250 \mu\text{s}$. The linear simulation included models for the gold, platinum, and platinum black electrodes. For the parameters used in the stimulation, see Table 2.2. These parameters were chosen to mimic the platinum black electrode in Figure 2.5.

We used SPICE, a common circuit simulator, to perform the linear simulation (see Appendix A). The development of an electrode model for a circuit simulation program was important because such a model can provide a powerful tool for the design of circuits for bio-electrical interfacing.

Transient simulations on the linear model demonstrated the formation of stimulation artifacts. As Figure 2.9 shows, the artifacts generated by stimulation prevented the electrode from returning to its equilibrium voltage for over 25 ms. Generation of a model artifact was as follows: First, the switch assumed a low value, $R_{\text{switch,closed}}$, which simulated connecting the stimulation voltage. The time constant associated with charging the electrode during stimulation was

$$\tau_{\text{stim}} = (R_s + R_{\text{switch,closed}}) C_e. \quad (2.19)$$

Next, the stimulation voltage source provided a biphasic pulse. Finally, the variable resistor assumed a high value. If the electrode capacitance had any remaining charge after the stimulation switch opens, that charge had to discharge through R_t , with the time constant

$$\tau_{\text{disch}} = R_t C_e. \quad (2.20)$$

Because R_t typically had a large value, discharging the capacitor was a slow process.

The linear model provided the simplest tool for exploring the cause of the artifact. Computing the charge stored in the electrode–electrolyte interface (the capacitor of the stimulation electrode) verified that stimulation did indeed charge the interface (Figure 2.10). The charge remained on the electrode for tens of milliseconds after stimulation, which resulted in the stimulation artifact. Although the stimulation pulse itself was symmetric, the voltage across the capacitor changed during the positive phase of stimulation, which led to asymmetric charging and discharging of the capacitor during stimulation.

As slight modification of the linear model, diodes were used to represent the charge-transfer resistance. This modeled the voltage-dependent nonlinearities of the charge transfer resistance, while still having permitted the use of standard circuit components. The diodes were ideal diodes with the scale current, I_s , chosen to model the charge transfer currents (see Table 2.3). The artifacts generated by this model (Figure 2.11) had significantly lower





Figure 2.9: Results of the SPICE simulation of linear electrode models. All the model electrodes tested exhibited stimulation artifacts that lasted for over 25 ms. As expected from the linear time constants, the artifact decayed most quickly for the platinum electrode and most slowly for the platinum-black electrode.



(a) Gold (red) and platinum (blue)

(b) Platinum black

Figure 2.10: Charge stored on the electrode capacitance in the linear SPICE model. Because of the small time constant involved in charging the gold and platinum electrodes, the charge on those electrodes changed so rapidly as to have appeared to change instantaneously at the timescale of these plots.

Table 2.3: Model Parameters for the SPICE Diode Model Simulation, based on experimental data (Figure 2.5)

Parameter	Value			Units
$R_{\text{switch,closed}}$	100			Ω
$R_{\text{switch,open}}$	1			$\text{G}\Omega$
R_s	10			$\text{k}\Omega$
	Platinum Black	Platinum	Gold	
V_0	50	50	50	mV
C_e	8000	300	300	pF
I_s	800	400	200	pA

amplitudes than those of the linear model, because the diodes shunted off much of the current that charged the capacitor in the linear model; however, these artifacts were still large enough to saturate a high-gain recording system for over 25 ms.

2.2.3 Physical Test System

To judge the accuracy of the models of stimulation-artifact generation, we required reference artifacts from a real, physical system. To take these artifacts, we designed a simple, single-channel stimulation system, shown in Figure 2.12. The stimulation voltage was buffered by a LF347 operational amplifier, and a DG212 switch connected the stimulation signal to the amplifier. A second operational amplifier on the same LF347 die amplified and high-pass filtered the signal. We used a DAC488HR digital-to-analog converter (DAC) to provide computer-controlled stimulation signals and a TDS3054B oscilloscope captured the amplified artifacts. Figure 2.13 shows the artifacts that the physical system produces.

2.2.4 Comparison of Models of Stimulation Artifacts

We considered a physical test system and three computational models for stimulation artifact generation. Each computational model generated qualitatively similar artifacts. Interestingly, in the physical model system, the platinum black electrode generated the smallest artifacts, while the computational models predicted that the platinum black electrodes




Figure 2.11: Results of the SPICE simulation of electrode models with diodes for the charge-transfer resistance. Because the diodes became very conductive for large applied voltages, the artifacts were smaller than those for the fully linear model (Figure 2.9).




Figure 2.12: Experimental system for generating stimulation artifacts. A computer-controlled DAC provided the stimulation pulse and an oscilloscope recorded the artifact.




Figure 2.13: Artifacts generated by the physical system. The apparent asymmetry in the stimulation voltage was due to the the removal of low-frequency offsets by a high-pass filter.

should have the largest artifacts. This suggested that exact prediction of the artifact behavior required either the consideration of additional effects in the model or a more detailed methodology for electrode parameter estimation.

2.3 Application of Models to Stimulation Artifact Elimination

Rather than develop a more detailed model for the electrode, it was useful to consider what we were able to learn from the simple models. Although exact prediction of stimulation artifact size and duration was not possible, the models still indicated that the artifact was closely linked to the capacitive nature of the electrode. As demonstrated in the linear model, the stimulation source charged the electrode capacitance. After stimulation, the electrode capacitance remained partially charged. The post-stimulation charge had no other discharge path but through R_t which set up a long time constant of $C_e R_t$. Based on this result, an additional conductive pathway in parallel with the electrode should have allowed for rapid charge removal (Jimbo et al., 2003).


The models provided a means of testing the efficacy of a discharge path in eliminating the artifact. By leaving the stimulation switch closed after the end of stimulation and setting the stimulation voltage to ground, the simulation provided a low-resistance discharge path equal to $R_s + R_{\text{switch,closed}}$ for the electrode capacitance. To explore the effect of a discharge path on the stimulation artifact, we conducted Simulink model simulations with varying discharge durations (Figure 2.14).

Examining the electrode voltage in response to stimulation (Figure 2.14(a)) showed that the size of the stimulation artifact remaining after a 1.0 ms discharge period was significantly less than the original artifact. Interestingly, a high-pass filter attenuated the artifacts (Figure 2.14(b)), as was expected from the slow time constant associated with the discharge process. Although the filtered artifacts decayed faster than the unfiltered electrode voltage, it was not clear from the simulation whether the filtering improved the ability to record neural activity.


A practical complication to the discharge procedure was that, due to the electrochemical offset voltage of the electrode, discharging the electrode to ground actually left a charge of $C_e V_0$ on the electrode. Simulations that included the electrode-offset voltage (Figure 2.14(c) and Figure 2.14(e)) exhibited a noticeable increase in the artifact that remained after the 1 ms discharge when compared with the simulations of the electrode without an offset voltage. These results indicated that an artifact-elimination system must sample the offset voltage of the electrode, and it must then discharge the electrode to the sampled voltage.

2.4 Conclusions


In this chapter, we presented circuit models of the generation of the stimulation artifact in electrodes. According to the models, the stimulation artifact was a result of a net charge that remained on the electrode–electrolyte interface after stimulation. Discharging the charge that accumulated during stimulation was limited by the large resistance in parallel with the electrode capacitance. We compared the artifacts that the models generated with models that a physical system generated, and we found that, although the model was not able to accurately predict the rates that the stimulation artifacts decayed, the predicted




(a) Electrode voltage, 0 mV offset




(b) Filtered electrode voltage, 0 mV offset




(c) Electrode voltage, 50 mV offset



(d) Filtered electrode voltage, 50 mV offset



(e) Electrode voltage, -50 mV offset



(f) Filtered electrode voltage, -50 mV offset

Figure 2.14: Simulation of a post-stimulation discharge for artifact elimination. The addition of a low-impedance path after stimulation significantly reduced the size of the stimulation artifact.

stimulation artifacts were qualitatively similar to the physical ones.

Although the accuracy of the models was limited, they were still a useful tool for the investigation of methods to reduce the stimulation artifact. Using the models, we demonstrated that the addition of a conductive pathway to the electrode was able to speed up the discharge process, reducing the stimulation artifact. Perhaps most importantly for our work, the we developed models that were suitable for inclusion in circuit simulators, so that they provide valuable tools for the design of electronics for eliminating the stimulation artifact.

CHAPTER 3

INTEGRATED CIRCUITS FOR STIMULATION ARTIFACT ELIMINATION ¹

Many neuroscience investigations into plasticity and development of neural tissue involve measuring and modifying the electrical activity of neural tissue. Conducting these experiments requires an electronic system capable of recording the minute extracellular action potentials, requiring low-noise amplifiers. Additionally, the electronic system must deliver stimulation signals to the culture that are sufficient to evoke neuronal activity. In most existing interfacing systems, the side-effects of the stimulation signals on the electrode and recording preamplifiers interfere with the ability to record signals. The resulting interference is known as the *stimulation artifact*. Because the immediate, local effect of a stimulus may contain information relevant to the experiment, it is highly desirable to reduce or eliminate the stimulation artifact.

In this chapter, we present the design and characterization of integrated circuits (ICs) for electronic interfacing to neurons. The IC consisted of multiple interfacing channels, each connecting to a single electrode. Each channel contained a recording system, stimulation buffer, and artifact-elimination circuitry. We present the design and testing of a first-generation IC, and then we present a second-generation IC that we designed to address some of the shortcomings of its predecessor.

3.1 Design of the First-Generation Integrated Circuit

The IC comprised the link between the neural tissue and the control and storage capabilities of computer systems. It contained an array of identical electrode interface channels, each containing stimulation buffers, recording preamplifiers, and artifact-elimination electronics (Figure 3.1). Grouping all the circuitry for one electrode together allowed easy expansion of

¹Portions of this chapter are from Blum et al. (2007)

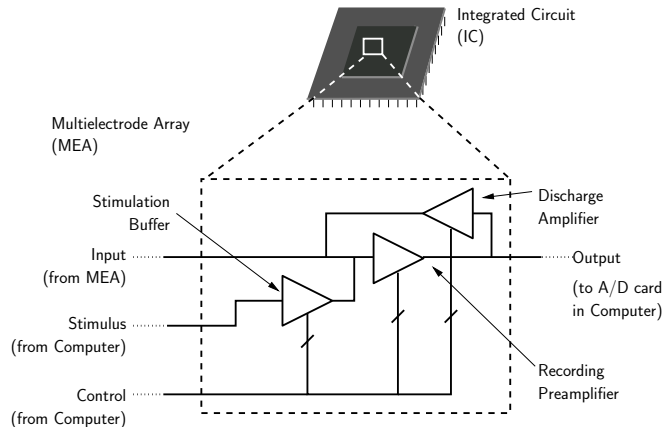


Figure 3.1: Block diagram of the IC. The IC contained an array of 16 stimulation, recording, and artifact-elimination elements, each matching with an electrode.

the IC for multi-electrode arrays (MEAs) with larger electrode counts. Digital circuitry, including shift registers and multiplexers, enabled the computer to apply independent control signals to each channel, allowing for independent stimulation and recording on all channels.

3.1.1 Low-Noise Pre-amplifier

In the design of the multichannel stimulation and recording system, our first goal was accurate recording of neural activity. There were a number of constraints that we had to consider, including those of recording from a single neuron and those of integrating the pre-amplifier into a larger system.

Recording consists of faithfully amplifying the small neural signals present at the electrodes while rejecting unwanted interference. The noise floor of the recording pre-amplifiers limits the ability to detect action potentials. Extracellularly recorded action potentials typically have maximum amplitudes $\leq 100 \mu\text{V}$ and frequency components in the range of 30 Hz–3 kHz, requiring input noise $\ll 20 \mu\text{V}_{\text{rms}}$ in that bandwidth.

Additionally, the pre-amplifier must reject the open circuit potential of the electrode. This electrochemical offset, which can be as large as 100 mV, can saturate the recording pre-amplifiers. Traditional ac coupling effectively blocks dc offsets to prevent saturation, but requires large capacitors that are not suitable for monolithic integration. We used a mixture of active circuitry and small value capacitors to block low frequency components.

Our first-generation design for the preamplifier was based around a low-noise transconductance amplifier with capacitive feedback, as Figure 3.2 shows. The feedback configuration limited the closed-loop gain to $-C_I/C_F$; we chose a gain of -50 (34.0 dB) using $C_F = 40$ fF and $C_I = 2.0$ pF. The choice of the closed-loop voltage gain was a compromise between maximizing the output signal level, conserving die area, and stability (see Section 3.1.3). Purely capacitive feedback would have left the inverting input of the feed-forward amplifier floating, so we had to introduce a dc path to that node. A weakly biased transconductance amplifier in the feedback path established the dc current path necessary to bias the input. The feedback amplifier also created a high pass filter that rejected the electrode open-circuit potential. Varying the bias current of the feedback amplifier tuned the high pass cutoff frequency of the recording system. Although the poles of the transfer function were complex, we made a close approximation with real poles,

$$p_{\text{HP}} \approx -\frac{g_{\text{mf}}}{C_F} \quad (3.1)$$

and

$$p_{\text{LP}} \approx -\frac{g_m}{C_L} \frac{C_F}{C_I + C_F + C_{\text{in}}} \quad (3.2)$$

where g_m is the transconductance of the feed-forward amplifier, g_{mf} is the transconductance of the feedback amplifier, C_I is the input capacitor, C_F is the feedback capacitor, and C_{in} is the parasitic input capacitance to the feed-forward amplifier. In addition to the two left half plane poles, the bandpass response also had a zero at the origin. There was an additional right half-plane (RHP) zero due to the parallel combination of the capacitive feed-through path and the feed-forward amplifier,

$$z_{\text{RHP}} = \frac{g_m}{C_F}. \quad (3.3)$$

This zero was at a very high frequency, so it had only a minor effect on the dynamics of the system.

The feed-forward recording amplifier topology (see Figure 3.3) is that of a wide-range amplifier (WRA) (Mead, 1989). Sizing the transistors to ensure that the input differential

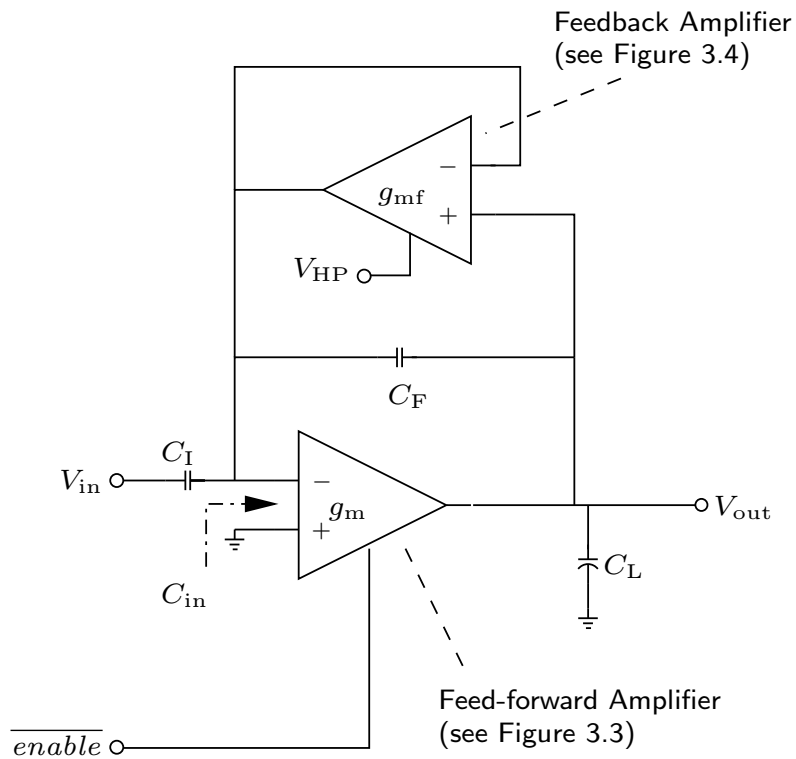


Figure 3.2: Topology of the recording preamplifier. Capacitive feedback around the feed-forward amplifier set the recording gain. The enable line of the feed-forward amplifier allowed disabling during stimulation as part of the artifact-elimination protocol. The feedback amplifier provided a dc pathway to the internal node of the capacitive divider, as well as an adjustable high pass filter to reject electrode offsets.

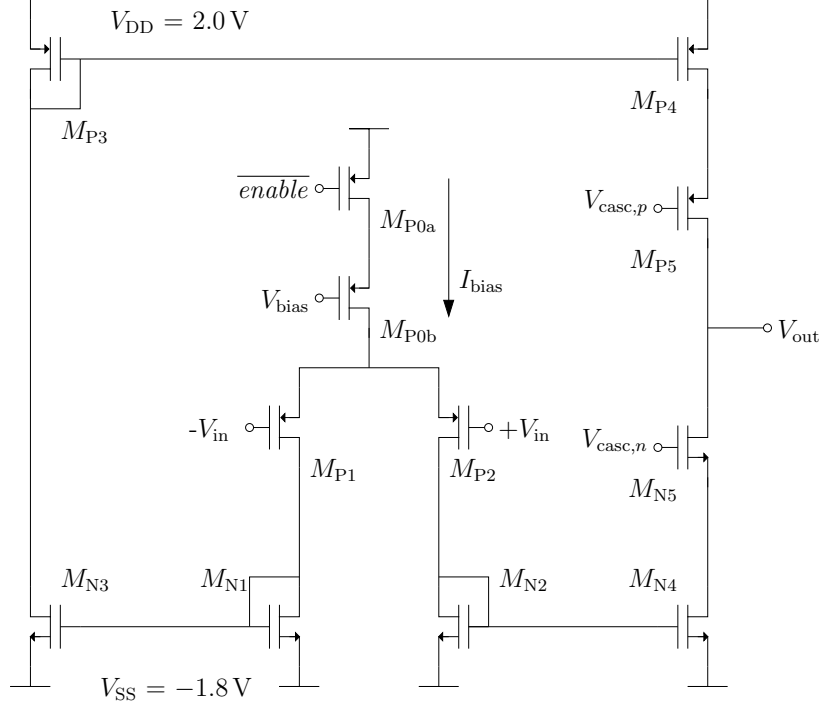


Figure 3.3: The feed-forward amplifier in the recording system. The amplifier used a WRA topology, which allowed for increased output voltage range compared to a simpler amplifier. The transistors were sized so that the input differential pair, M_{P1} and M_{P2} , operated in moderate inversion, while the other transistors operated in strong inversion. This sizing resulted in the noise from the input differential pair dominating over the noise produced by the other transistors.

pair operated in moderate inversion and the load current mirrors operated in strong inversion (see Table 3.1), such that the input stage transconductance was much greater than the transconductance of any of the current mirror transistors, minimized the noise contributions of the current mirror transistors, resulting in noise levels near the theoretical limit for a given power dissipation (Harrison and Charles, 2003). Neglecting the noise contributions from the current mirror transistors, the mid-band thermal noise contribution to the output of the recording preamplifier was

$$\overline{v_{no}^2} = \left(1 + \frac{C_I}{C_F}\right)^2 \frac{16kT}{3g_m} \Delta f. \quad (3.4)$$

At a bias level of $I_{bias} = 13 \mu A$, $g_m \approx 130 \mu S$, which resulted in an input-referred noise level of $13 \text{ nV}_{rms}/\sqrt{\text{Hz}}$, or 720 nV_{rms} in a 3 kHz bandwidth.

The feedback amplifier (see Figure 3.4) used an operational transconductance amplifier (OTA) topology. Previous implementations of capacitively coupled amplifiers had used a

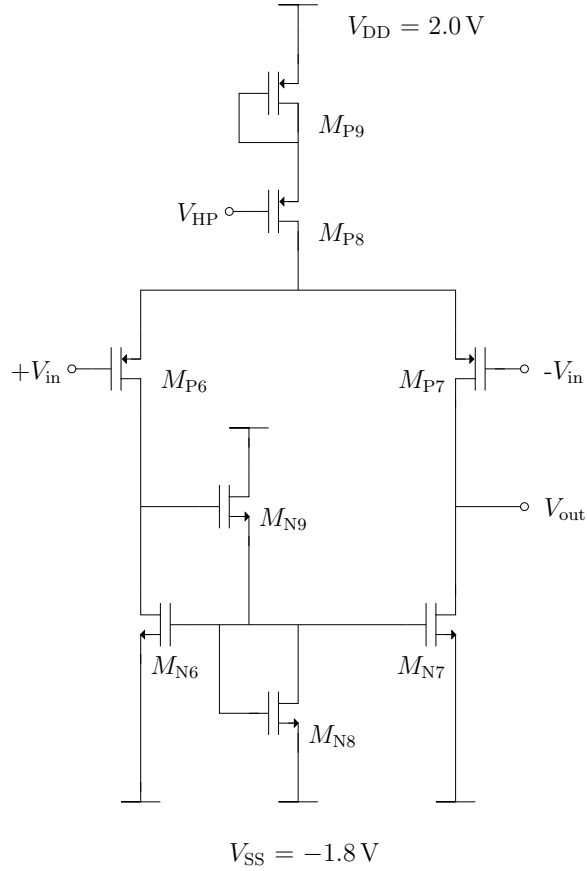


Figure 3.4: Feedback amplifier that provided the dc feedback around the feed-forward amplifier. M_{N8} and M_{N9} ensured that V_{ds} of M_{N6} was large enough for operation in sub-threshold saturation.

metal-oxide-semiconductor field-effect transistor (MOSFET) adaptive element to provide the dc feedback pathway (Aziz et al., 2007; Delbrück and Mead, 1994; Harrison and Charles, 2003). Such a topology was effective for normal recording operation; however, the conductance of such an element became very large for even a small applied voltage. As we show in Section 3.1.3, our artifact-elimination method relied on the charge at the inverting input remaining constant during stimulation; thus, we were unable to use a feedback element that would have acted as a low impedance path, as it would have dissipated the stored charge.

Due to the very small current levels necessary to provide large equivalent resistance values, modifications to the standard OTA topology were necessary. Had we used a standard OTA, the input transistor for the active load, M_{N6} , would have operated in the subthreshold triode mode, because it would have insufficient drain-source voltage $V_{ds} < 4U_T$, where U_T

Table 3.1: Transistor Sizes and Inversion Modes in the Recording System

Transistor	Width (μm)	Length (μm)	Inversion
Feed-forward Amplifier			
M_{P0a}	21.8	3.0	Strong
M_{P0b}	5.0	5.0	Strong
M_{P1}, M_{P2}	200.0	1.0	Moderate
M_{P3}, M_{P3}	4.0	3.2	Strong
M_{P5}	18.0	0.8	Strong
$M_{N1}, M_{N2}, M_{N3}, M_{N4}$	3.0	11.2	Strong
M_{N5}	6.0	0.8	Strong
Feedback Amplifier			
M_{P6}, M_{P7}	2.0	5.0	Weak
M_{P8}, M_{P9}	2.0	1.0	Weak
M_{N6}, M_{N7}	1.0	20.0	Weak
M_{N8}	10.0	1.0	Weak
M_{N9}	1.0	4.0	Weak

is the thermal voltage ($U_T \equiv \frac{kT}{q}$, where q is the charge of an electron). Two additional transistors, M_{N8} and M_{N9} , increased the drain-to-source voltage for the input transistor to the active load, M_{N6} . The diode-connected transistor, M_{N8} , sank current that M_{N9} had to source. To provide that current, the gate voltage of M_{N9} had rise, which increased V_{ds} of M_{N6} . The extra transistor in the tail current supply, M_{P9} acted to scale down the bias current relative to the reference current, which provided finer control of the high pass cutoff frequency than would have been possible with a standard current mirror.

The feedback amplifier was a significant impact on the noise level of the overall amplifier structure. Assuming all transistors in the signal path operated in weak inversion (a safe assumption given the extremely small bias currents that were necessary), the feedback amplifier produced noise according to

$$\overline{v_{n,fb}^2} = \frac{32kT}{3g_{mf}} \Delta f \quad (3.5)$$

which was larger than that of an ideal resistive element by a factor of 8/3. The feedback capacitors filtered this noise, so that the output noise contribution of the feedback amplifier

was

$$\overline{v_{\text{no,fb}}^2} = \left(\frac{1}{1 + \frac{2\pi f C_F}{g_{\text{mf}}}} \right)^2 \frac{32kT}{3g_{\text{mf}}} \Delta f. \quad (3.6)$$

This noise, which summed with that of the thermal noise due to the feed-forward amplifier (3.4), had a spectrum that was flat for frequencies below the high-pass pole of the recording preamplifier and fell off with a first order slope for higher frequencies. Assuming the real pole approximation of (3.1), the resulting input-referred thermal noise (in the neural signal bandwidth of 30 Hz–3 kHz) was

$$\begin{aligned} \overline{v_{\text{ni}}^2} = & \frac{(C_F + C_1)^2}{C_1^2} \frac{16kT}{3g_m} (3 \text{ kHz} - 30 \text{ Hz}) \\ & + \frac{C_F}{C_1^2} \frac{16kT}{3\pi} \left[\frac{f_{\text{HP}}}{f_{\text{HP}} + 30 \text{ Hz}} - \frac{f_{\text{HP}}}{f_{\text{HP}} + 3 \text{ kHz}} \right]. \end{aligned} \quad (3.7)$$

Setting a very low cutoff frequency resulted in a high level of spot noise at frequencies below the high-pass cutoff frequency. In the pass-band, the noise from the feedback amplifier appeared qualitatively similar to $1/f$ noise on top of the thermal noise of the main amplifier, as observation of the flat noise level at low frequencies required slow measurements. True $1/f$ noise in the system and measurement setup may have prevented observation of the flat noise spectrum at low frequencies.

It was intuitively attractive to assume that tuning of the high pass cutoff frequency would have shaped the noise of the feedback amplifier in a manner similar to the gain-bandwidth relationship of a standard feedback amplifier; however, this was not the case. Due to the relationships

$$\sqrt{\overline{v_{\text{no,fb}}^2}} \Big|_{f=0} \propto \sqrt{\frac{1}{g_{\text{mf}}}} \quad (3.8)$$

and

$$f_{\text{HP}} \propto \frac{1}{g_{\text{mf}}} \quad (3.9)$$

a one-decade change in cutoff frequency resulted in a half-decade change in the root-mean-square (RMS) noise level at dc. As a consequence of this relationship, a high-pass cutoff frequency at a very low frequency produced minimal noise in the bandwidth relevant to

neural action potentials, regardless of the high noise levels at low frequencies. The need to minimize noise in the bandwidth relevant to neural signals (30 Hz–3 kHz) placed an upper limit on the high pass cutoff frequency, so that the filtering was best suited for removing low-frequency electrochemical offsets from the electrode. The data acquisition system (DAQ) had to provide an additional high-pass filter to remove out-of-band signals.

Although frequency-shaping techniques may have reduced the noise contributions of the feedback amplifier in the neural bandwidth (Spang and Schultheiss, 1962), a simpler method was through optimization of the feedback capacitors, which reduced the noise contributions from the feedback amplifier. Increasing C_F required a proportional increase in g_{mf} to maintain the same cutoff frequency, which resulted in lower output noise, according to (3.5). Increases in C_I had to occur with any increase in C_F ; otherwise, the reduction in system gain would have resulted in higher input referred noise, as in (3.7). Alternatively, increasing C_I alone increased the gain of the system, which provided an area-efficient means of reducing the effect of the feedback noise.

3.1.2 Stimulation Buffer

After establishing the ability to record from the neural culture, the next step was to design electronics for stimulation, thus providing the means for two-way communication with the culture. Our design combined stimulation and recording circuitry, which gave experimenters the ability to switch any electrode between stimulation and recording functionality.

The stimulation buffer (Figure 3.5 and Table 3.2), used digital inputs to control the connection of stimulation voltages, provided by an external digital-to-analog converter (DAC), to the electrode. When the enable line was un-asserted, the combinational logic turned off M_{P10} and M_{N10} , so that the stimulation output was high-impedance. When enabled, the phase input controlled whether the pull-up or pull-down branch was active, and the output stage drove the electrode to $+V_{stim}$ or $-V_{stim}$.

The two transistors closest to the stimulation voltages, M_{P11} and M_{N11} , limited the current that the buffer was able to source or sink into the electrode. Because the transistors M_{P11} , M_{P12} , M_{N11} , and M_{N12} had their sources connected to the stimulation voltages, the

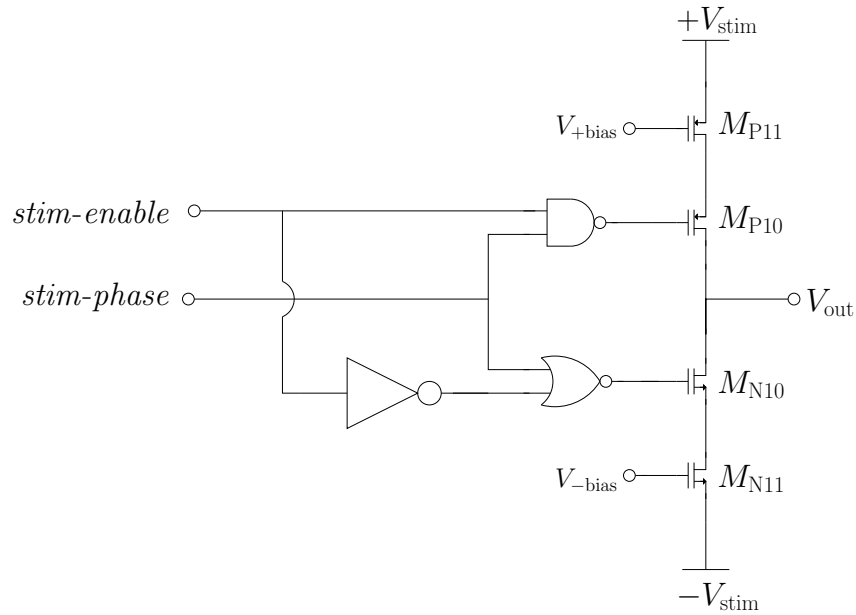


Figure 3.5: Schematic of stimulation buffer. When the enable signal was active, the output stage drove the electrode to the positive or negative stimulation voltage, depending on the phase input. The stimulation voltages were generated by a DAC external to the IC. When disabled, the output was high impedance.




Figure 3.6: Bias network that produced V_{-bias} and V_{+bias} from a single voltage input. The connection of the sources of M_{P12} and M_{N12} to the stimulation voltages forced the maximum stimulation currents to depend on the stimulation voltages.

Table 3.2: Transistor Sizes in the Stimulation Buffer and Bias Network

Transistor	Width (μm)	Length (μm)
$M_{P10}, M_{P11}, M_{P12}$	18.0	1.0
$p\text{FETs in Logic Gates}$	6.0	1.0
$M_{N10}, M_{N11}, M_{N12}$	6.0	1.0
$n\text{FETs in Logic Gates}$	2.0	1.0

limit currents were dependent on both the stimulation voltages and the gate bias voltages on those transistors (see Figure 3.6). This did place a minimum voltage requirement on stimulation, as low values of $\pm V_{\text{stim}}$ would have resulted in stimulation currents that were insufficient to evoke neuronal activity.

The current limits also governed the transition speed from positive to negative stimulation. For electrode capacitances on the order of 1 nF, the time constant $2c_e/g_{\text{m,out}}$, where $g_{\text{m,out}}$ is the transconductance of the output stage transistor (M_{P10} or M_{N10}), dominated over the relatively quick switching speeds of the logic gates.

3.1.3 Artifact Elimination

To maximize the experimental benefit of the capability of the stimulation buffer to activate at arbitrary electrodes, it was necessary that we include circuitry to actively eliminate the stimulation artifact. To eliminate the interference with recording after stimulation, commonly referred to as the stimulation artifact, we had to discharge the electrode back to the electrochemical offset voltage of the electrode. The feedback capacitors of the recording preamplifier provided the storage elements necessary to track the average electrode voltage. By tracking the average voltage, rather than instantaneous voltage, we minimized interference from neuronal activity that may occur immediately before stimulation on the stored voltage.

To reduce the alteration of the stored charge in the capacitors and to prevent saturation of the rest of the recording signal chain, we disabled the feed-forward amplifier during stimulation. We disabled the feed-forward amplifier by open-circuiting its tail current supply, a method that minimized charge injection at the output of the recording system.

Table 3.3: Transistor Sizes in the Discharge Amplifier

Transistor	Width (μm)	Length (μm)
M_{P14}, M_{P15}	30.0	2.0
M_{P16}, M_{P17}	30.0	2.0
M_{N14}, M_{N15}	3.0	2.0

After stimulation, we connected a discharge amplifier (Figure 3.7, Table 3.3) in a feedback loop with the recording preamplifier (Figure 3.8). Neglecting the slow effect of the feedback amplifier, the capacitors in the recording preamplifier stored the offset voltage of the electrode, such that the output of the recording preamplifier returned to ground when the electrode returned to its pre-stimulation voltage. The feedback loop controlled the current output of the discharge amplifier, ensuring that it acted to bring the electrode back to its pre-stimulation voltage. The tail current supply of the discharge amplifier limited the maximum discharge current, preventing unintended neural stimulation. This feedback method ensured that the stimulation and artifact elimination currents had no net effect on the electrode charge, although the use of charge balancing on the stimulus signal still improved artifact elimination duration by minimizing the necessary correction to the electrode voltage that the discharge amplifier had to make. For a summary of the activity of the different amplifiers during stimulation and artifact elimination, see Table 3.4.

During the discharge phase, the discharge amplifier and recording preamplifier functioned in a closed loop. Because there were multiple poles in the discharge path, we had to consider the possibility of unstable loop dynamics. Investigating stability required identification of the open-loop poles and zeros, as well as the loop gain. In addition to the poles and zeros from the recording preamplifier (3.1–3.3), the discharge amplifier and a polarizable electrode introduced a pole,

$$p_d = -\frac{1}{r_s C_I} \quad (3.10)$$

and a zero,

$$z_d = -\frac{1}{r_s c_e} \quad (3.11)$$

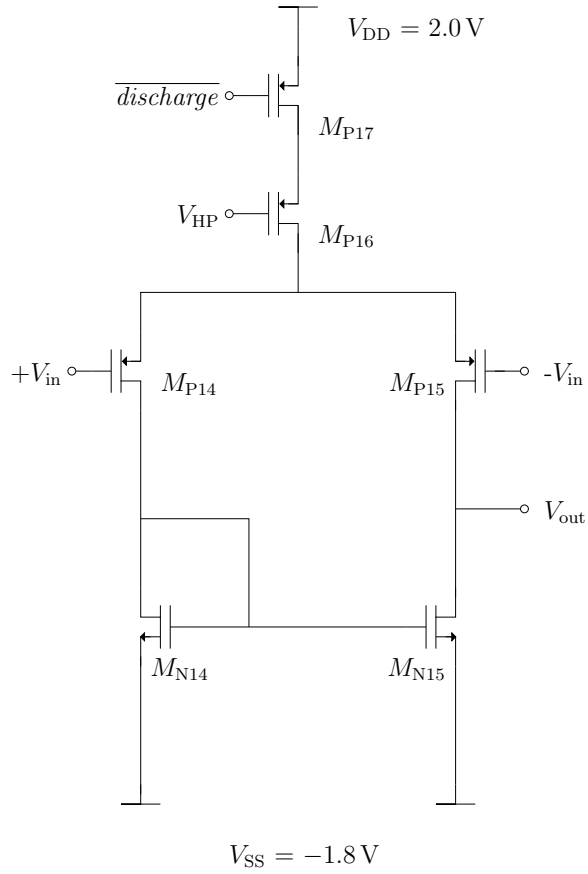


Figure 3.7: Discharge amplifier used in artifact elimination. The topology was that of an OTA, with an added switching transistor in the tail current supply, M_{P17} , to control amplifier activity.




Figure 3.8: Complete recording, stimulation, and artifact elimination system. During the discharge phase, the capacitors in the recording system stored the electrode offset voltage, V_{elec} .

Table 3.4: Amplifier Activity During Operation Modes

Amplifier	Recording	Stimulation	Artifact Elimination
Feed-forward	On	Off	On
Feedback	On	On	On
Stimulation	Off	On	Off
Discharge	Off	Off	On

where r_s is the small-signal spreading resistance of the electrode, and c_e is the small-signal electrode interface capacitance. Assuming typical electrode parameters of $r_s = 10 \text{ k}\Omega$ and $c_e = 3 \text{ nF}$, we had $|p_d| = 8 \text{ MHz}$ and $|z_d| = 2.6 \text{ kHz}$. The discharge system also introduced a pole at the origin that cancelled out the zero of the bandpass recording preamplifier.

The loop gain consisted of the product of the recording system gain and the gain of the discharge amplifier,

$$A_{\text{disch}} = g_{\text{md}} [r_{\text{out}} || (r_t + r_s)] \quad (3.12)$$

where r_{out} is the output resistance of the discharge amplifier and r_t is the small-signal charge transfer resistance.

Increasing the bias current of the discharge amplifier increased the loop gain, which affected the system in two ways. First, the slow pole from the high-pass filter moved towards the zero from the electrode and spreading resistance, which increased the speed of the system response. The time constant of that electrode zero set the limiting speed for discharge, so that less capacitive electrodes discharged faster. Second, the poles from the low pass recording response and the discharge system met, branched off from the real axis, and moved toward the RHP zero. For large loop gains, these complex poles crossed the $j\omega$ axis and entered the RHP, resulting in an unstable system. A loop gain of 100 000 yielded a 45° phase margin (see Figure 3.9). At this gain, the dominant pole, originally from the high-pass filter, moved almost to the electrode zero. Further increases to the loop gain increased the overshoot, which eventually led to instability, without a significantly improvement the linear behavior of the discharge. In practice, slew-rate limiting occurred during discharge, so that there was some performance improvement for larger discharge currents.

Separating the recording and discharge poles would have increased the stability margin.




Figure 3.9: Linear analysis of the discharge loop stability, assuming typical electrode parameters of $r_s = 10\text{ k}\Omega$ and $c_e = 3\text{ nF}$. The root locus plot (a) shows the closed loop poles for a 45° phase margin, which occurred for a closed loop gain of 100 000. The bode plots show the magnitude (b) and phase (c) response for the open loop. We used the Matlab SISO design tool was to generate these plots.

This required lowering the recording pole or increasing the discharge pole. The recording pole had to remain above 10 kHz, or attenuation of action potentials would have occurred. The remaining method to increase the stability margin was to speed up the discharge pole, which would have occurred through using a small value of C_1 . The stability requirements of moderate loop gain and small input capacitance were in direct opposition to those for minimizing the effect of noise from the feedback amplifier. The combination of our recording gain of 50 and the maximum discharge gain of 130 ($g_{\text{md}} \approx 85 \mu\text{S}$, $r_{\text{out}} \approx 1.6 \text{ M}\Omega$) was a loop gain of 6500, which left a large gain margin.

Using the Spectre circuit simulator and the transistor models from the North Carolina State University Cadence Design Kit, we are able to conduct simulations of the transient behavior of the IC (NCSU; Spectre). To model the electrode, we used a linear model with the same parameters as in the linear stability analysis. Figure 3.10 shows the simulated result of the artifact suppression system. After the stimulation and discharge, the recording system did not saturate. Although an artifact was still present, operation of the recording system in its linear region implied that post-processing of the recorded signal would have been able to filter out the remaining artifact.

3.1.4 Digital Control

The operation of each stimulation and recording channel was subject to three control signals that enabled the preamplifier, the stimulation buffer, and the discharge amplifier. Shift registers, which were composed of standard cells, provided the digital control signals for the switches in each channel. For each signal, there was a separate shift register. The three shift registers were all connected to a serial peripheral interface (SPI) bus, and digital inputs and combinational logic enabled the loading of one shift register at a time. The use of the SPI bus permitted easy interfacing to microcontrollers that had dedicated hardware for SPI. A potential disadvantage with this system was that the control speed was limited. The microcontroller had to update the entire 16-bit control word to control the switch in a single channel, and had to update multiple words to control the different switches.




Figure 3.10: Simulated operation of the artifact-elimination circuitry. This Spectre simulation used detailed transistor models and a linear electrode model. During the stimulation (shaded interval), the feed-forward amplifier was disabled and the output signal was due to capacitive feed-through. After the stimulation, the artifact elimination circuitry discharged the electrode capacitance. At 1 ms, the discharge amplifier shut off. After the end of discharge, there was a remaining transient signal; however, the magnitude of the transient was small enough to have permitted operation of the recording system in its linear range.




Figure 3.11: Die photograph showing the area of active circuitry for 16 stimulation and recording channels. The circuitry for a single electrode occupied $140\ \mu\text{m} \times 230\ \mu\text{m}$ (not including digital control).

3.2 *Experimental Characterization of the First-Generation Integrated Circuit*

We fabricated the design using the Taiwan Semiconductor Manufacturing Company Ltd. (TSMC) $0.35\ \mu\text{m}$ process (Figure 3.11), available through Metal Oxide Semiconductor Implementation Service (MOSIS) (MOSIS). The IC contained 16 stimulation and recording channels, analog biases, and digital control. The stimulation and recording circuitry for a single electrode occupied $140\ \mu\text{m} \times 230\ \mu\text{m}$ (not including digital control).

Our test setup included the IC, a microcontroller (Microchip PIC18LF452) to control the stimulation and artifact elimination timing, external DACs to set the stimulation bias voltages (Analog Devices DAC8420), and additional off-chip amplifiers (National Semiconductor LF347) that brought the total recording gain up to 1800 (65 dB).

3.2.1 Low-Noise Preamplifier

Evaluating the capabilities of the recording preamplifier consisted of measuring its frequency response (Figure 3.12) and input noise (Figure 3.13). We measured a mid-band gain of 48.1 (33.6 dB), which fell slightly short of the design goal of 50. The experiments confirmed that adjusting the bias voltage of the feedback amplifier effectively controlled the high pass pole over a range of $2\ \text{Hz} < f_{\text{HP}} < 300\ \text{Hz}$. For power supplies of $V_{\text{DD}} = 2.0\ \text{V}$ and $V_{\text{SS}} = -1.8\ \text{V}$, the power consumption was $100\ \mu\text{W}$.

The location of the high-pass pole strongly affected the overall system noise, as (3.7)




Figure 3.12: Frequency response of the recording system. The different high pass poles were the result of 50 mV increments in the feedback bias voltage. The equal frequency spacing was a consequence of the operation of the feedback amplifier in weak inversion (subthreshold).

predicted. Low values of the high-pass control voltage, corresponding to high cutoff frequencies, introduced large amounts of noise into the bandwidth of neural signals. The curves for high values of the control voltage did not reach a flat level at low frequencies, which indicated the presence of significant $1/f$ noise in addition to the thermal noise.

An important measure in evaluating the noise performance for an amplifier is its noise efficiency factor (NEF), which compares the noise level with that of a single bipolar transistor consuming the same power and having no excess or $1/f$ noise (Steyaert et al., 1987).

The definition of the NEF is

$$NEF = \sqrt{\frac{\overline{v_{ni}^2} \cdot 2I_{total}}{\pi \cdot U_T \cdot 4kT \cdot \Delta f}} \quad (3.13)$$

where I_{total} is the total current that the amplifier draws from the power supply. Ideally, the NEF is unity; all practical circuits have a higher value.

Table 3.5 summarizes the noise and NEF in the relevant neural bandwidth of 30 Hz–3 kHz. The most significant source of noise was that the feed-forward amplifier introduced more thermal noise than expected. This may have been due to the noise currents from the




Figure 3.13: Input-referred noise of the recording system. As the bias current of the feedback amplifier increased, the high-pass cutoff frequency and the amount of noise in the neural bandwidth of 3 kHz–30 kHz both increase.

Table 3.5: Summary of Input Referred Noise in the bandwidth 30 Hz–3 kHz ($C_I = 2$ pF, $C_F = 40$ fF, $I_{\text{bias}} = 13$ μ A)

f_{HP} (Hz)	Predicted Noise (μV_{rms})	Measured Noise (μV_{rms})	Measured NEF
2	1.67	4.77	17.6
4	2.21	4.84	17.9
12	3.50	5.36	20.0
36	5.21	7.32	27.4
116	6.98	12.3	46.1

current mirror transistors being too large to neglect, or it may have been noise present on the power supplies coupling onto the output. There were two other possible sources for the discrepancy between predicted and measured noise. First, the predictions neglected the contributions of $1/f$ noise, which dominated for very low cutoff frequencies. Second, the inaccuracies in the assumption of real poles, as in (3.1), became more significant for faster high-pass cutoffs.




Figure 3.14: Frequency–impedance plot of the electrode. The impedance plot corresponded approximately to $C_e = 8 \text{ nF}$ and $R_s = 10 \text{ k}\Omega$.

3.2.2 Stimulation Buffer

The choice of electrode played an important role in the performance of the stimulation and artifact-elimination circuitry. We chose a commercially available electrode array consisting of $30 \mu\text{m}$ diameter gold electrodes (Ayanda Biosystems), upon which we have deposited platinum black (Marrese, 1987; Ross et al., 2004). Before connecting the electrode to the stimulation and recording system, we characterized its impedance by comparing the voltage drops across the electrode and a known series resistor (Figure 3.14). At frequencies below 10 kHz , the electrode impedance was dominated by C_e . The measured impedance of $20 \text{ k}\Omega$ at 1 kHz was equivalent to an 8 nF capacitance. At high frequencies, the impedance was limited by R_s , which was approximately $10 \text{ k}\Omega$.

The electrode was immersed in a saline medium formulated to emulate *in vivo* conditions (Hank’s Balanced Salt Solution, Hyclone), although no neurons were present. This setup mimicked the electrical conditions present during electrophysiology experiments without the burden of maintaining environmental controls for cell survival.

Using the stimulation buffer, we applied a biphasic voltage pulse, consisting of a $250 \mu\text{s}$ positive voltage followed by a negative voltage of equal duration, to the electrode. As a

consequence of the design of the bias network of the stimulation buffer (see Figure 3.6), the maximum available stimulation current depended on the stimulation voltages. To generate enough source–gate voltage on M_{P12} to be able to provide approximately $10\ \mu\text{A}$, we required $+V_{\text{stim}} = V_{\text{DD}}$. To prevent electrochemical damage that this large voltage might cause, we limited the maximum voltage at the electrode with a 1N4148 diode external to the IC. The negative phase amplitude was $-V_{\text{stim}} = 0.5\ \text{V}$.

Measuring the voltage across a small resistor in series with the electrode allowed us to determine the current provided during stimulation. The additional components (diode, resistor, and an Texas Instruments INA129P instrumentation amplifier) necessary for this experiment did add parasitic components that may have been of the order of the electrode impedance, thus possibly affecting the results. The diode in parallel to the electrode also may have introduced asymmetry in the current response. The stimulation current (Figure 3.15) was expected to spike to an initial high value and then decay exponentially; the observed signal did not conform to this shape, indicating that the stimulation buffer was not able to provide enough current to fully charge the electrode capacitance. The peak current provided by the stimulation was approximately $9\ \mu\text{A}$, enough to evoke neuronal activity (Wagenaar et al., 2004).

3.2.3 Artifact Elimination

For testing the artifact-elimination circuitry, we used the same electrode and setup, including the external diode, as in Section 3.2.2. We observed the artifact behavior for discharge periods of 0.5, 0.7, and 1.0 ms (Figure 3.16). The stimulation parameters were similar to those used during the stimulus current characterization; however, the duration of the negative phase was reduced to $150\ \mu\text{s}$, resulting in improved charge balance. We set the high-pass cutoff frequency for the preamplifier to 21 Hz, minimizing the thermal noise contributions and the effect of filtering on the artifacts.

Under these stimulation conditions, the stimulation artifact saturated the recording system for over 15 ms. The use of the artifact-elimination circuitry was able to reduce the duration during which the recording system was unusable (Figure 3.16). The duration of




Figure 3.15: Current provided by the stimulation buffer. The stimulation buffer was able to provide approximately $9\ \mu\text{A}$ to the electrode.

the discharge played a critical role in the artifact-elimination. For the 0.5 ms discharge, the electrode had not returned to baseline before the end of discharge, resulting in large remaining artifacts. The size of the artifact remaining after discharge decreased with increasing discharge duration; however, we observed an artifact size for 1.0 ms that did not decrease significantly for longer durations.

To demonstrate the effectiveness of the recording system after stimulation, we repeated the stimulation and artifact-elimination trials while applying a $500\ \mu\text{V}_{\text{peak}}$, 1 kHz sine wave to the saline solution. For all the discharge durations tested, the amplified sine wave was visible within 2 ms of the end of stimulation. The ringing observed after the end of stimulation was likely due to an off-chip amplifier, as the duration of the ringing was independent of the duration of the discharge phase.

Additional reduction of the remaining artifact was possible using an external, first-order high-pass filter (Figure 3.17). A 100 Hz filter, commonly used for extracellular recordings, attenuated the remaining transient to less than $\pm 250\ \mu\text{V}$. (Figure 3.17). Although similar results were possible with the filter inherent to the preamplifier, such use would have resulted in a higher level of thermal noise, according to (3.7).




Figure 3.16: Effect of the artifact-elimination circuitry. The curves shown are for discharge durations of 0.5, 0.7, and 1.0 ms. Two sets of curves are shown: with and without a $500 \mu\text{V}_{\text{peak}}$, 1 kHz sine wave applied to the saline solution. During stimulation, the recording preamplifier was inactive, and the recording in this duration (shaded interval) was due to parasitic coupling through the preamplifier. After stimulation, the artifact-elimination circuitry allowed for observation of the sine wave within 2 ms. Without use of the artifact-elimination circuitry, the recording system was saturated for over 15 ms.




Figure 3.17: Effect of a 100 Hz high-pass filter on the artifact remaining after a 1.0 ms discharge. The preamplifier itself contributed a high-pass filter of 21 Hz. A $500 \mu\text{V}_{\text{peak}}$, 1 kHz sine wave was applied to the saline solution.

Table 3.6: Transistor Sizes in the Output Buffer

Transistor	Width (μm)	Length (μm)
p FETs	60.0	1.0
n FETs	20.0	1.0

3.3 Design of the Second-Generation Integrated Circuit

Although the IC described in the previous section (Section 3.2) significantly advanced the capability for neural stimulation and recording, it did suffer from some performance limitations. Most notable among the previous IC’s deficiencies were the noise levels and the limits on the stimulation current. Consideration of these deficiencies led to a revised design for the IC.

3.3.1 Preamplifier Noise Reduction

As suggested by (3.7), we were able to reduce the input-referred noise by adjusting the capacitor sizes and the transconductance of the feed-forward amplifier. Increasing the transconductance directly lowered the thermal noise floor at the cost of increased power dissipation.

A significant portion of the noise in the previous design was due to the feedback amplifier. In accordance with (3.7), we increased the size of C_I , which we expected to result in a reduction in the input-referred noise level. For the revised design, we chose $C_I = 16$ pF and $C_F = 80$ fF. Resizing the capacitors had the additional effect of increasing the gain of the preamplifier to 200, requiring a commensurate decrease in C_L to maintain the same bandwidth as the previous IC.

Because reducing C_L made the bandwidth more sensitive to off-chip capacitance, the addition of a buffer stage was desirable. We used a design that uses current-feedback to reduce the output resistance of the common-source stage (see Figure 3.18) (Manetakakis and Toumazou, 1996). This design was simple, self-biasing, and able to drive large capacitive loads. As a disadvantage, it introduced a signal loss of approximately 6 dB into the signal pathway. For the transistor sizes, refer to Table 3.6.

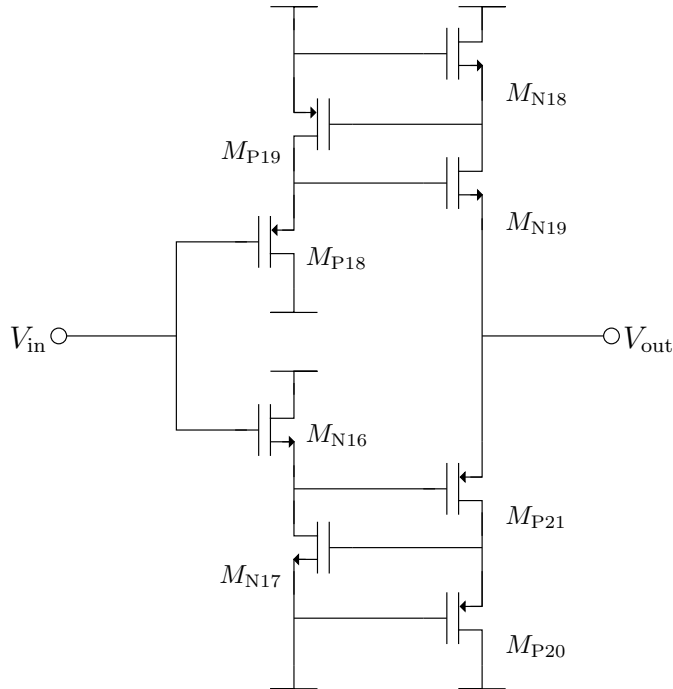


Figure 3.18: Low output impedance buffer, from Manetakis and Toumazou (1996). The amplifier design used feedback to lower the output impedance, which permitted the amplifier to drive large capacitive loads.

3.3.2 Increasing the Current Output of the Stimulation Buffer

The previous stimulation buffer suffered from limited current-drive capabilities. The bias network for the buffer imposed a relationship between the stimulation voltage and current that required nearly full-scale stimulation voltage for an appreciable current. To prevent damage to the electrodes while providing effective stimulation current, additional voltage-limiting discrete elements were necessary. These additional elements were cumbersome to include in the test setup, and they impacted the performance of the IC. Because of this limitation, we considered an alternate design for the stimulation buffer.

For our second-generation design, we needed to isolate the stimulation current supply from the stimulation buffer. To accomplish this requirement, we developed an improved design, based on a standard WRA topology, in which the tail current limited the stimulation current (Figure 3.19). We provided two branches capable of supplying the tail current for the stimulation buffer, one for each phase (positive or negative) of the stimulation. Two digitally controlled transistors, M_{N27} and M_{N28} , were included in the tail current branches

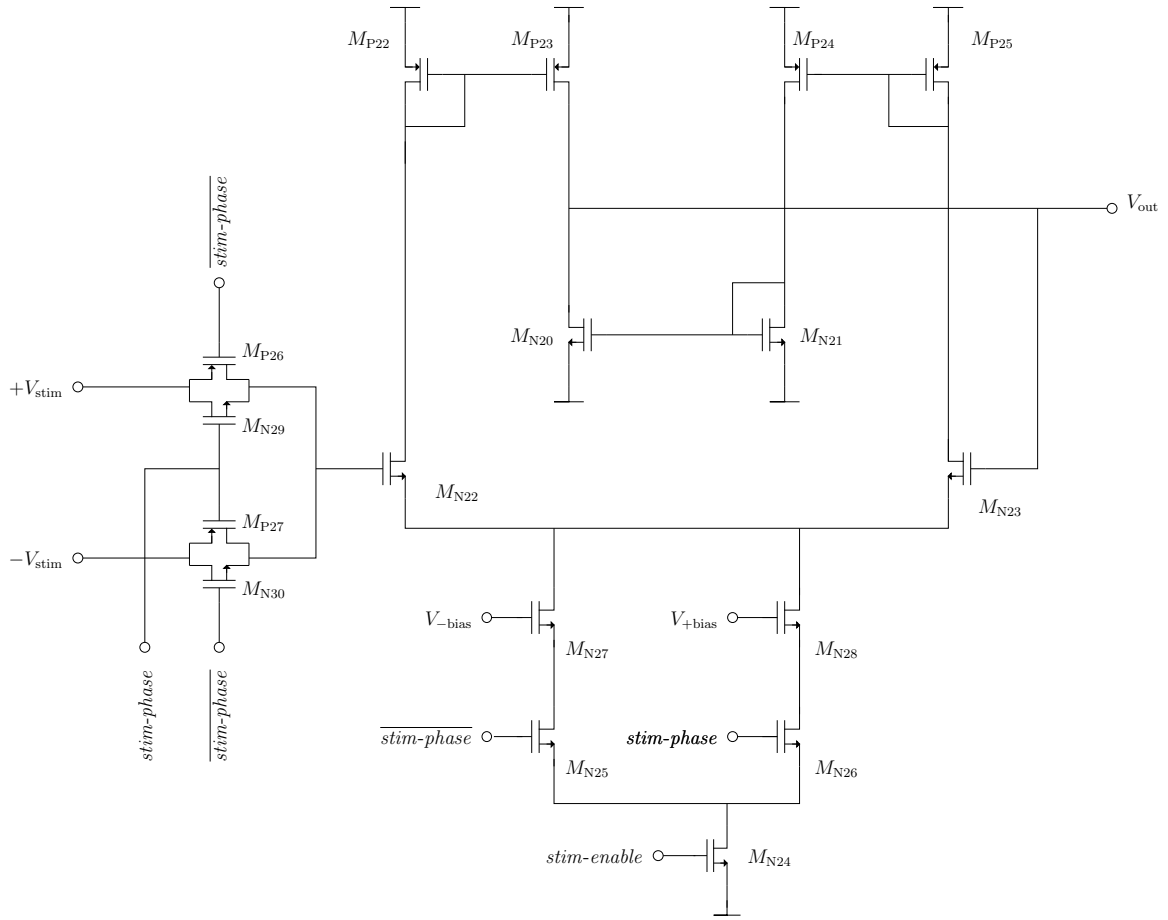


Figure 3.19: Revised stimulation buffer, based on a WRA topology. The digital control signals switched the input voltage between $+V_{stim}$ and $-V_{stim}$, and they also switched the tail current supply so that the current limits for positive and negative stimulation were permitted to differ.

to select one supply as active. The stimulation voltages were provided as inputs to the WRA. As with the tail current, we included transistors with digital control (M_{P26} , M_{P27} , M_{N29} , and M_{N30}) to select between the two input voltages.

We also modified the external control circuitry to select the stimulation voltage separately for each channel. This allowed for the stimulation of two different electrodes with opposite-phase signals. This increased flexibility in stimulation came at the cost of an additional shift register and more complex combinational logic to connect the shift registers to the SPI bus.

Table 3.7: Transistor Sizes in the Revised Stimulation Buffer

Transistor	Width (μm)	Length (μm)
M_{N20}, M_{N21}	8.0	1.5
M_{N22}, M_{N23}	10.0	1.5
$M_{N24}, M_{N25}, M_{N26}, M_{N27}, M_{N28}$	10.0	1.5
M_{N29}, M_{N30}	1.0	1.0
$M_{P22}, M_{P23}, M_{P24}, M_{P25}$	30.0	1.5
M_{P26}, M_{P27}	3.0	1.0

Table 3.8: Transistor Sizes in the Revised Feedback Amplifier

Transistor	Width (μm)	Length (μm)
M_{P28}, M_{P29}	2.0	1.0
M_{P30}	6.0	1.0

3.3.3 Addition of Pole Shifting to the Artifact Elimination Circuitry

Another improvement to the IC was the capability for *pole shifting*, or the adjustment of the high-pass pole frequency for the recording system. Increasing the high-pass pole has been shown to reduce the duration necessary for the recording system to recover from transients (DeMichele and Troyk, 2003). We observed a similar effect when we applied an external, high-pass filter to the IC during our experimental characterization of artifact elimination. By providing an additional, switchable tail current supply for the feedback amplifier, we were able to selectively shift the high-pass pole of the stimulation channel in order to recover from the post-discharge artifact in a shorter time (Brown et al.).

As shown in Figure 3.20, adding the capability for pole shifting required the addition of three transistors, M_{P10} , M_{P11} and M_{P12} , to the feedback amplifier design. These transistors added a second tail current supply to the feedback amplifier. When the control signal $\overline{pole-shift}$ was low (logically true), M_{P10} turns on, activating the additional tail current. M_{P11} and M_{P12} control the amount of additional current that the branch contributes. These changes did require the addition of another global bias voltage, $V_{\text{quick-bias}}$, and another shift register to provide $\overline{pole-shift}$.

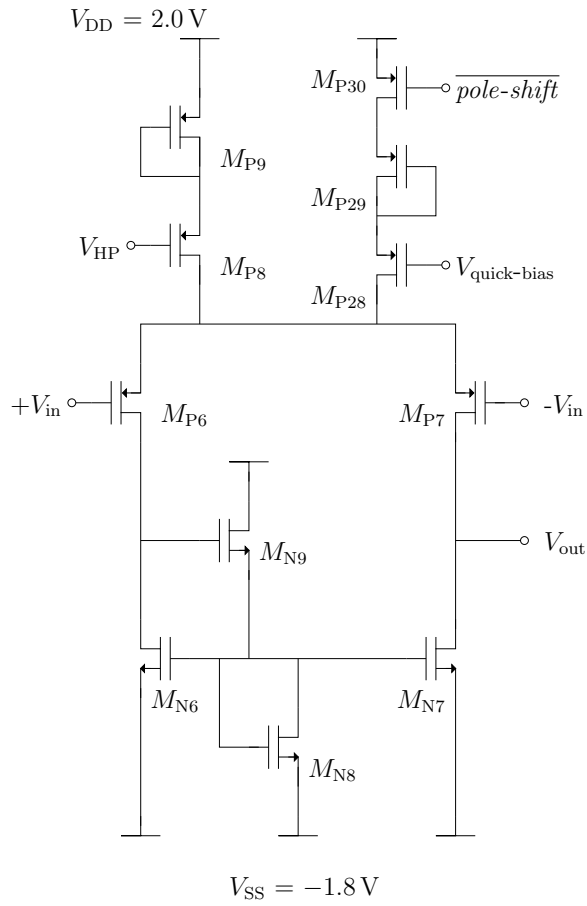


Figure 3.20: Revised feedback amplifier, including the capability for pole shifting. The addition of M_{P10} , M_{P11} , and M_{P12} provided a second tail-current pathway for the feedback amplifier. Activation of the second tail-current supply raised the high-pass cutoff frequency of the recording system, which enabled it to recover from the stimulation artifact faster.




Figure 3.21: Frequency response of the revised IC. As in the original IC, adjusting the high-pass bias control voltage in linear increments resulted in logarithmically spaced increments of the high-pass cutoff voltage. The midband voltage gain was 84, or 38.5 dB.

3.4 *Experimental Characterization of the Second-Generation Integrated Circuit*

The second-generation IC was designed to be compatible with the existing test system, described in Section 3.2. Revised firmware for the PIC18LF452 was necessary because of the addition of two control signals (stimulus phase and pole shift) and their shift registers.

3.4.1 **Low-Noise Preamplifier**

Using the swept-sine mode of the SR785 Signal Analyzer, we measured the frequency response of the revised IC. The preamplifier itself, by design, had a gain of $C_I/C_F = 200$, although the buffer stage attenuated this gain. The measured midband gain for the recording system (preamplifier and buffer) was 84, equivalent to 38.5 dB. Assuming that the preamplifier provided its designed gain, the buffer must have attenuated the signal by a factor of 2.4, or 7.5 dB to result in the measured voltage gain. As with the original IC, adjusting the bias current of the feedback amplifier controlled the high-pass cutoff frequency of the recording system. Figure 3.21 shows the effect of adjusting the feedback bias on the frequency response.




Figure 3.22: Input-referred spot noise of the new IC. As with the previous design, the noise from the feedback amplifier dominated at frequencies below 1 kHz. This noise was strongly dependent on the bias current of the feedback amplifier, which also controlled the high-pass cutoff frequency. Because of this dependency, reducing the high-pass cutoff frequency also reduced the noise of the recording system.

For the noise characterization, we used the FFT mode of the SR785. Figure 3.22 shows the input-referred spot noise of the revised IC. The input referred noise in the bandwidth of 30 Hz–3 kHz was as low as $3.5 \mu\text{V}_{\text{rms}}$ for a high-pass frequency of 16 Hz. Computation of the NEF at various high-pass frequencies (see Figure 3.23) showed that the revised IC performs significantly better in noise measurements than the original IC. The noise reduction compared to the original IC was more pronounced for faster high-pass cutoff frequencies. This effect was consistent with our modification of the sizes of C_1 and C_F in the new IC. As discussed in Section 3.1.1, the noise that the feedback amplifier introduced dominates as the high-pass cutoff frequency increased. According to (3.7), the noise from the feedback amplifier was proportional to C_F/C_1^2 , so that the modification to the capacitor sizing mostly affected the noise from the feedback amplifier. The increasing advantage of the revised IC at faster high-pass cutoff frequencies was due to the combination of these two effects. For a summary of the noise levels and NEF of the revised IC, refer to Table 3.9.




Figure 3.23: Comparison of the NEF in the original and revised ICs. The revised IC had a significantly lower NEF, which verified the importance of the capacitor sizing in the noise performance. The improvement in the NEF was more noticeable for faster high-pass cutoff frequencies; this was because the noise from the feedback amplifier dominated as the high-pass cutoff frequency increased, and it was this noise which is most affected by the capacitor sizing.

Table 3.9: Summary of Input Referred Noise in the Revised IC in the bandwidth 30 Hz–3 kHz

($C_I = 16 \text{ pF}$, $C_F = 80 \text{ fF}$, $I_{\text{bias}} = 16 \text{ }\mu\text{A}$)

f_{HP} (Hz)	Predicted Noise (μV_{rms})	Measured Noise (μV_{rms})	Measured NEF
16	1.07	3.50	14.5
34	1.24	3.68	15.3
75	1.38	4.12	17.1
166	1.47	5.38	22.2
363	1.42	7.37	30.5




Figure 3.24: Stimulation current provided by the revised buffer. The bias currents were set for an approximate maximum current amplitude of $15\ \mu\text{A}$. The stimulation was a $500\ \text{mV}$, $250\ \mu\text{s}$ positive pulse followed by a $-500\ \text{mV}$, $250\ \mu\text{s}$ negative pulse. The shape of the stimulation current was consistent with charging a capacitor.

3.4.2 Stimulation Buffer

We measured the stimulation current that the new buffer can deliver to an electrode by amplifying the voltage across a $100\ \Omega$ resistor in series with the electrode with an INA129P instrumentation amplifier. We set the current limit of the buffer to approximately $\pm 15\ \mu\text{A}$ to prevent any electrochemical damage to the electrode. We set the stimulation voltages to $\pm 500\ \text{mV}$, and each phase was $250\ \mu\text{s}$ long. The stimulation current peaked at the beginning of each stimulation phase, and it quickly decayed towards zero (see Figure 3.24). This observed behavior was consistent with the rapid charging of a capacitor to a voltage. Recent studies have suggested that exponentially decaying stimulation currents, such as those provided by the second-generation stimulation buffer, evoke neural activity more efficiently than constant-current stimulation (Sahin and Tie, 2007). By comparing the stimulation current to that of the first-generation IC (see Figure 3.15), we observed significant improvement between the first and second generations.

3.4.3 Application of Pole Shifting to Artifact Elimination

To verify the utility of pole shifting in reducing the post-discharge artifact, we conducted an experimental trial of artifact elimination in which we varied the duration of the pole shift. The stimulation was a ± 500 mV, biphasic voltage pulse. The initial, positive phase had a duration of 600 μ s, and the negative phase had a duration of 100 μ s. The asymmetric durations were chosen to minimize the size of the resulting artifact. The stimulation current limits were set at ± 15 μ A. After the stimulation, the artifact elimination circuitry discharged the electrode capacitance for 3 ms.

Without the use of the pole-shift technique, a large post-discharge artifact was present. Even without the additional current from the pole-shift circuitry, the feedback amplifier set $f_{HP} \approx 100$ Hz. The pole-shift bias was set to shift f_{HP} to approximately 1 kHz. We tested activating the pole-shift current for 3, 6, and 9 ms after the end of the discharge phase. As Figure 3.25 shows, the recording system recovered from the post-discharge artifact much faster when the pole-shift current remained active. Additionally, the recording amplifier was capable of amplifying a 1 kHz, 100 μ V sine wave applied to the saline medium while the pole-shift circuitry was active. The shutoff of the pole-shift current did introduce a step in the recorded waveform, although it may have been possible to eliminate this effect by modifying the firmware so that it would have slowly decreased the pole-shift current bias (raising $V_{\text{quick-bias}}$) before the un-assertion of control signal $\overline{\text{pole-shift}}$ shut off the current.

Even with the pole shift, the first-generation IC was able to bring the electrode and recording system back to their operational states in less time than the second-generation IC. We considered the possibility that the second-generation stimulus buffer generated artifacts that were more difficult to eliminate. To test if the stimulus buffer was the source of the increased duration of the artifact elimination, we conducted a set of artifact elimination trials for varying levels of the maximum stimulation current. The experimental data (Figure 3.26) showed that that the size and duration of the post-discharge artifact did not decrease for smaller stimulation current limits, indicating that the re-design of the stimulation buffer was not responsible for the change in artifact behavior.

Alternatively, it was possible that the increased gain of the recording system makes the




Figure 3.25: Effect of the pole-shifting technique, in which the bias current of the feedback amplifier for the stimulation channel increases briefly after stimulation. Without use of the pole shift, the recording system required over 20 ms to recover from saturation because of the stimulation artifact. With the pole-shift, the recording system recovered from saturation immediately, and the $100\ \mu\text{V}$, 1 kHz sine wave that was applied to the saline medium was clearly visible in the recordings within 6 ms of stimulation.




Figure 3.26: Artifacts for various stimulation current levels. Even when the maximum stimulation current limit was below $1\ \mu\text{A}$, significant post-discharge artifacts were present. These data plots indicated that the increased current drive of the second-generation stimulation buffer was not the source of the longer post-discharge artifacts.

artifact-elimination circuitry more sensitive to the disturbance to the electrode caused by the end of the discharge phase. In future revisions of the IC, it may be beneficial to increase the sizes of both C_I and C_F reducing the input-referred noise while maintaining the gain of the original IC, as the benefits in noise level and artifact elimination speed would more than justify the increase in die area.

3.5 Conclusions

Overall, experimental characterization of the stimulation and artifact-elimination IC validates its utility for neuroscience experiments. For recording, the measured noise levels were as low as $3.5\ \mu\text{V}$ —small enough to measure extracellular action potentials from vertebrate neurons. The stimulation buffers were able to provide peak currents over $10\ \mu\text{A}$, enough to have quickly charged and discharged a platinum black electrode with a capacitance of $8\ \text{nF}$. The artifact-elimination circuitry was able to bring the electrode and recording system back to a usable state within $5\ \text{ms}$ of the end of stimulation, which would have been enough to have permitted the recording of the neuronal response to stimulation.

CHAPTER 4

DIGITAL INTERFACES TO THE INTEGRATED CIRCUITS

In the previous chapters, we have studied the origins of the stimulation artifact and developed integrated circuits (ICs) to reduce its interference with neuroscience experiments. Although the ICs were a significant contribution to the field of neural interfacing, there were still areas in which we could improve on their performance. One of the major experimental difficulties that remained was the presence of a *post-discharge artifact*, a transient effect that occurred after the artifact-elimination circuitry had discharged the electrode to its pre-stimulation voltage.

Unlike the stimulation artifact that was considered in the design of the IC, the post-discharge artifact does not result in saturation of the recording system, and because the recording system does continue to operate in its linear range, it is possible to recover neural activity through use of filtering techniques. In this chapter, we present the design of a filtering system to remove the post-discharge artifact from the output of the artifact-elimination IC.

The filtering system used a multi-resolution analysis based on wavelet filters. Through the filtering, we were able to separate out the post-discharge artifact from neural action potentials. The filtering was also able to remove noise that had temporal characteristics that differed from action potentials. This filtering complemented the artifact elimination that our custom IC provided.

To develop our filtering system, we used a field-programmable gate array (FPGA), a reconfigurable digital system that is capable of high-speed operation. The FPGA was a natural counterpart to very large scale integration (VLSI), in that they both were well-suited to the design and implementation of highly parallel systems. Just as in VLSI, we designed simple building blocks, and then used multiple instances of those blocks to construct a complex system. Also, by converting the analog signals into a digital representation for

FPGA computation, we introduced the possibility of direct, digital interfacing between the FPGA and a data acquisition system.

In this chapter, we first present a custom analog-to-digital converter (ADC) (Section 4.1, starting with the overall choice of topology, and then proceeding to the details of its optimization. After we establish the ability to convert the analog output of the IC into a digital representation in the FPGA, we present our filtering algorithm (Section 4.2). We consider the design of the system, evaluate it using simulated signals, and finally connect it to the artifact elimination IC for experimental characterization.

4.1 Interfacing Analog Signals to the FPGA

The FPGA operates on digital signals, so interfacing it to the stimulation and recording IC, which generates analog signals, requires the use of an ADC. Although there are many commercial, off-the-shelf ADCs, using the computational resources of the FPGA to build a custom ADC does offer advantages. First, by minimizing the complexity of the external components, we can reduce the overall cost of the system. Second, the design of a custom converter can simplify the design of the rest of the FPGA system, because the digital value generated by the ADC is internal to the FPGA, eliminating the need for a digital interface that would either require a bus controller for a serial interface or a large number input pins for a parallel interface.

We based our design criteria on the analog properties of the signals from the IC. The noise inherent to the stimulation and recording IC limited the resolution for which all bits were meaningful. For the second-generation IC, the output-referred noise level was approximately $10\text{ mV}_{\text{rms}}$ and because of this noise, there was no benefit to making an ADC with a resolution less than 20 mV . For a full-scale value of 28 V (the voltage limit of the amplifiers external to the IC), the necessary resolution was $\log_2(28\text{ V}/20\text{ mV}) = 10.4$ bits; accordingly, we designed an 11 bit ADC. The low-pass filter external to the IC limited the signal to frequencies below 10 kHz ; accordingly, we chose a sample rate of 20 kHz . These design criteria for speed and resolution suggested that we use a medium-speed ADC design. We chose to use a *successive-approximation* converter, which uses a binary-search algorithm

to determine the value of an analog signal (Gray and Hodges, 1978).

4.1.1 Successive Approximation Analog-to-Digital Converter

A successive approximation ADC, shown in Figure 4.1, consists of a digital-to-analog converter (DAC), a comparator, and digital control logic. The digital control logic, often known as a *successive-approximation register (SAR)*, drives the DAC, which produces a reference voltage, and the comparator determines whether the reference voltage is greater than or less than the input voltage. The SAR uses the information in the output of the comparator to guide a binary search that determines the input voltages. The converter initially assumes a mid-range value, and sets a reference DAC to this value. A comparator determines if the input signal lies above or below the value of the reference DAC, and sets the most significant bit (MSB) of the conversion register accordingly. The control logic then sets the reference DAC to the middle of the half in which the input is known to reside, and the comparator again indicates whether the actual input voltage is above or below the reference voltage. This process continues, accumulating one additional bit of precision for each comparison cycle. In total, for a 2^N -bit converter, N comparison cycles are necessary.

The maximum resolution that the successive-approximation ADC can obtain is a function of the resolution of the DAC and comparator. If either the DAC or the comparator is not accurate to one least significant bit (LSB), the comparator output will not truly reflect the relationship between the value of the SAR and the input voltage. These components also limit the speed of the converter. The DAC and comparator must settle within one clock cycle, so the sampling rate of the ADC is limited to

$$f_s = N \cdot t_s \tag{4.1}$$

where f_s is the sampling rate, and t_s is the settling time of the DAC.

4.1.2 Sigma-Delta Digital-to-Analog Converters

Because the DAC that generates the reference voltage is the limiting factor in the sampling rate of the ADC, we must consider its performance.

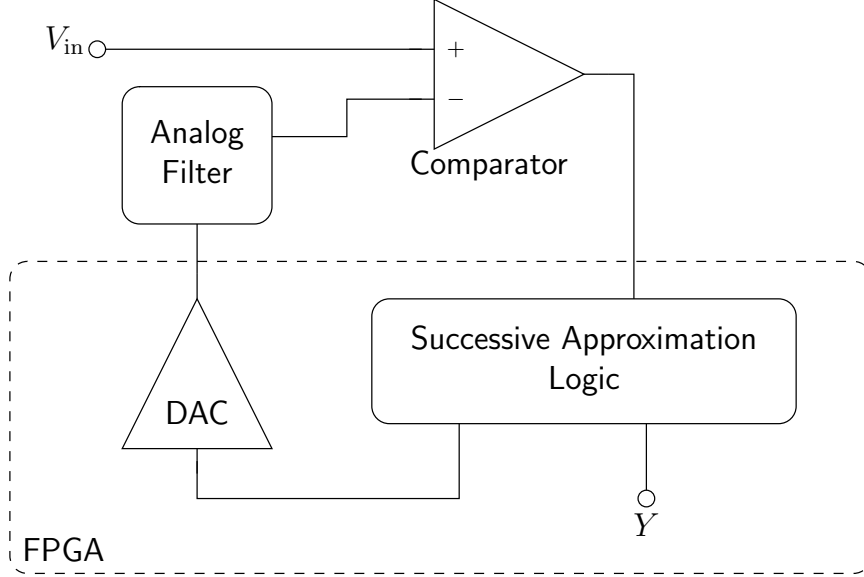


Figure 4.1: Successive-Approximation ADC. The ADC consists of a successive-approximation register (SAR), a DAC, and a comparator. The SAR controls the reference voltage that the DAC generates, and through comparing the reference voltage to the input voltage, the SAR determines the value of the input voltage. The analog filter is necessary for the operation of the DAC. A 2^N -bit ADC requires N clock cycles to complete one conversion.

Although the FPGA has digital outputs, it is possible to generate analog voltages filtering a digital output signal. The FPGA has the advantage that it can operate at much higher speeds than the Nyquist rate of biological signals. A quickly varying, digital output voltage can encode an analog voltage in its average value. An analog low-pass filter can recover the encoded analog voltage, although this requires that the digital switching operates fast enough so that it is in the stop-band of the filter. This method of extracting a slow, analog signal out of a fast, digital one is an example of an *oversampling converter*.

The digital output is only a single bit, so at any instance, there may be a large difference between the digital, instantaneous voltage and the desired analog value. This difference is known as *quantization error* (Bolton et al., 1999). For a single-bit quantizer with output levels at $\pm\Delta/2$, the probability distribution of the quantization error is uniformly distributed along the interval $[-\Delta/2, \Delta/2]$. The power of the quantization noise, S_Q , is equivalent to the variance of the quantization probability distribution,

$$S_Q = \frac{1}{\Delta} \int_{-\Delta/2}^{\Delta/2} x^2 dx = \frac{1}{12} \Delta^2 \quad (4.2)$$

(Widrow, 1956). In the frequency domain, the quantization noise has a uniform distribution along the interval $[0, f_S]$, and the noise power in the Nyquist bandwidth of the signal, S_B , is

$$S_B = S_Q \frac{f_N}{f_S}. \quad (4.3)$$

By increasing f_S relative to f_N the total quantization noise power in the signal bandwidth decreases.

$\Sigma\Delta$ converters are a class of oversampling converters that incorporate feedback into their operation (Figure 4.2) (Inose and Yasuda, 1963). The feedback loop has a single integrator, that tracks the error between the input and the output.

The integration loop follows the difference equation

$$y[n] = x[n-1] + q[n] - q[n-1] \quad (4.4)$$

where $y[n]$ is the output signal, $x[n]$ is the input signal, and $q[n]$ is the quantization noise. The z -transform of (4.4) is

$$Y(z) = \underbrace{z^{-1}X(z)}_{STF} + \underbrace{(1-z^{-1})Q(z)}_{NTF}, \quad (4.5)$$

which shows that the output is the sum of the filtered signal and noise. The signal transfer function (STF) is a delay, so the signal is almost unaffected by the integration loop. In contrast, the noise transfer function (NTF) has a high-pass characteristic, attenuating the noise in the signal bandwidth and amplifying the noise at high frequencies. An external, low-pass filter removes components of the output signal that are above the input data rate, removing a large portion of the quantization noise.

Computing the ratio of the signal power to the quantization noise power, from (4.5), gives the signal-to-noise ratio (SNR) of the $\Sigma\Delta$ DAC, and comparing it to the SNR of a Nyquist-rate DAC gives the expression for the resolution, in bits, of the $\Sigma\Delta$ DAC,

$$B = 1.5 \log_2 M - 0.85 \quad (4.6)$$

where M is the oversampling ratio, $M \equiv f_S/f_N$. Accordingly, for every octave increase in M the converter gains 1.5 bits of resolution.

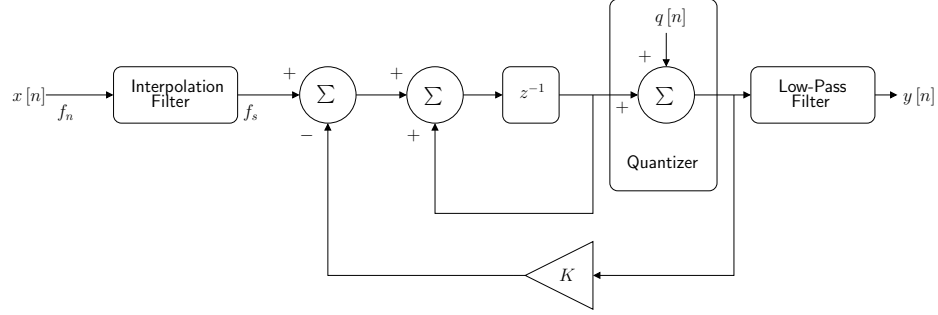


Figure 4.2: First-order $\Sigma\Delta$ loop. The integrator shapes the frequency content of the quantization noise, reducing the noise at frequencies much lower than the sampling rate.

Adding a second integrator to the $\Sigma\Delta$ converter enhances the ability of the feedback loop to shape the feedback noise (Candy, 1985). In this second-order converter,

$$Y(z) = \underbrace{z^{-2}}_{STF} X(z) + \underbrace{(1 - z^{-1})^2}_{NTF} Q(z), \quad (4.7)$$

This increases the order of the NTF, resulting in less noise at low frequencies, corresponding to the bandwidth of the signal. The STF is now a two-cycle delay, which does not affect the frequency content of the signal. Figure 4.4 compares the NTF for a first-order and second-order $\Sigma\Delta$ loop. Because the NTF of the second order has a lower magnitude in the bandwidth of the signal, the SNR is higher for a second order converter running at the same oversampling ratio. The resolution of the second-order $\Sigma\Delta$ converter is

$$B = 2.5 \log_2 M - 2.14 \quad (4.8)$$

which is an additional bit per octave compared to the resolution of the first-order converter (4.6).

4.1.3 Data Converter Implementation and Testing

We designed a successive approximation converter on the Virtex-4 FPGA, based on a reference design (Logue, 1999). The design of the DAC in the reference design limited the sampling rate to

$$f_s = \frac{f_{\text{clk}}}{(N + 1) 2^{N+1}} \quad (4.9)$$

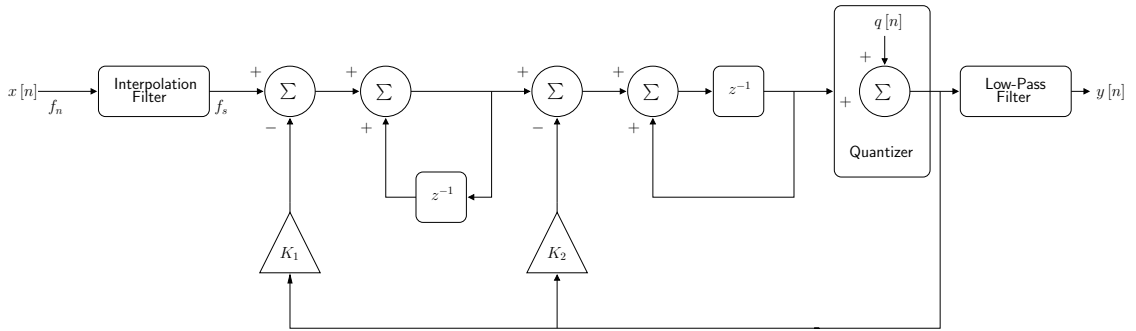


Figure 4.3: Second-order $\Sigma\Delta$ loop. The addition of a second integrator increases the order of the NTF, resulting in less quantization noise at low frequencies relative to the sampling rate.




Figure 4.4: Comparison between first and second-order NTFs. The second-order NTF has a smaller magnitude for frequencies much less than the sampling frequency. Because the output from the feedback loop passes through a low-pass filter, the increase in noise at high frequencies does not affect the overall output of the converter




Figure 4.5: Sampling frequency constraint of the original ADC. At the desired sampling frequency of 20 kHz, the ADC was able to operate at a resolution of 7 bits.

where f_{clk} is the FPGA clock speed (100 MHz). Based on our requirement of $f_S = 20$ kHz, the reference design was limited to 7 bits of resolution. Because this resolution was insufficient, we had to speed up the operation of the reference DAC.

To increase the conversion speed, we designed a second-order $\Sigma\Delta$ acDAC to provide the reference voltage. Because one of the integrators used in the second-order loop used a different topology than in the first-order converter, the converter was not able to run at the 100 MHz clock speed of the FPGA; instead, the converter operated at 50 MHz.

To operate the ADC at a 20 kHz sampling rate, the DAC that provides the reference voltage must operate at

$$f_S = \frac{50 \text{ MHz}}{M} = B \times 20 \text{ kHz} \quad (4.10)$$

combining (4.10) and (4.8) results in a transcendental equation,

$$\frac{B + 2.14}{2.5} = \log_2 \left[\frac{50 \text{ MHz}}{(B + 1) \times 20 \text{ kHz}} \right]. \quad (4.11)$$

Rounding the numerical solution down to an integer gives $B = 15$. Although a 15 bit converter would have been possible, the SNR of the artifact elimination IC was not high

enough to justify the design of an ADC with more than 11 bits.

We also designed an 11 bit DAC to provide analog outputs to the FPGA. Although the input signal to the reference DAC, by nature of the binary search algorithm used in the ADC, had discontinuous changes, the input to the DACs used as voltage outputs had to be interpolated prior to entering the $\Sigma\Delta$ integration loop. We designed an interpolation filter that had a $256\times$ linear interpolation cascaded with a repeating filter. Figure 4.6 shows the design of the linear interpolation section of the interpolation filter. The linear interpolation used a binary shift operation to compute $1/256$ of the difference between two consecutive samples, and repeatedly added that fraction of the difference at the higher sample rate. The frequency response of this filter was

$$H_{\text{int}}(e^{j\omega}) = \frac{1}{2304} \left[\frac{\sin(128\omega)}{\sin(\omega/2)} \right]^2 \left[\frac{\sin(9\omega/2)}{\sin(\omega/2)} \right] \quad (4.12)$$

which approximated the ideal interpolation response,

$$H_{\text{int}}(e^{j\omega}) = \begin{cases} 2304 & \omega \leq \frac{\pi}{2304}, \\ 0 & \omega > \frac{\pi}{2304}. \end{cases} \quad (4.13)$$

Figure 4.7 shows the frequency response of our interpolation filter. The filter that we designed provided a good approximation to an ideal interpolation filter, while having a simple implementation in the FPGA.

To test the ADC, we implemented an FPGA with the 11 bit ADC connected to the 11 bit DAC. We applied a full-scale input sine wave at 768 Hz to the input of the ADC, and we measured the signal at the output of the DAC (Figure 4.8). The combination had a noise floor of $1.4 \times 10^{-5} V_{\text{rms}}/\sqrt{\text{Hz}}$, or $1.5 \text{ mV}_{\text{rms}}$. The SNR of the converter was

$$\text{SNR}_{dB} = 20 \log \frac{V_{\text{full-scale}}}{V_{\text{noise}}} \quad (4.14)$$

where $V_{\text{full-scale}}$ is the root-mean-square (RMS) value of the full-scale output voltage. Using $V_{\text{full-scale}} = 2.3 V_{\text{rms}}$, the SNR was 64 dB, which corresponded to a 10.6 bit converter.

4.2 Signal Processing in FPGAs

To design separate out the post-discharge artifact from the action potentials, we must use a signal processing algorithm capable of differentiating between the two types of signals.

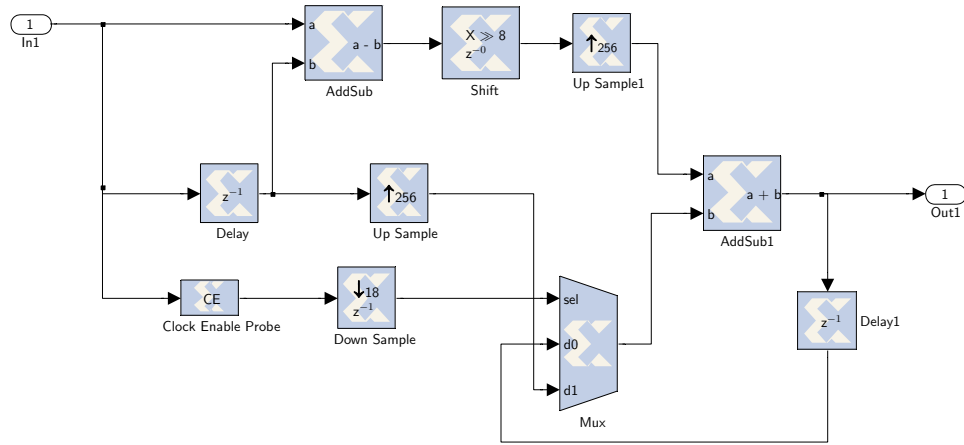


Figure 4.6: Linear interpolation filter. The filter takes the difference between two consecutive samples, uses a binary shift to compute $1/256$ of the difference, and steps the output by this fraction at every interpolated sample.




Figure 4.7: Interpolation filter for the $\Sigma\Delta$ DAC. Because the $\Sigma\Delta$ DAC oversamples its input, an interpolation filter is necessary to prevent the introduction of high-frequency signals. This filter is the cascade of a $256\times$ linear interpolation filter and a $9\times$ hold filter.




Figure 4.8: Power spectrum of the ADC–DAC cascade. For a constant, midscale input, the output noise was $1.5 \text{ mV}_{\text{RMS}}$. For the full-scale input, we applied a 768 Hz sine wave to the ADC. Comparing the full-scale input to the noise floor, we found that the SNR of the system was 64 dB, which corresponded to a 10.6 bit converter.

These signals have different timescales, with action potentials occurring over the span of a millisecond, and the post-discharge artifact extending for several milliseconds. Because of this distinction, the wavelet transform, which excels at distinguishing between transients of different timescales, is a natural choice for post-discharge artifact removal (see Appendix C for a review of wavelets and multi-resolution analysis). For an appropriate choice of wavelet family, the wavelet transform concentrates the energy of the action potentials into a small number of time–scale resolutions, while concentrating the post-discharge artifact into a different resolution and spreading out the noise over many resolutions. In the transform space, it is then possible to remove much of the noise and post-discharge artifact by applying a threshold to the resolutions containing action potentials and discarding other resolutions. Running the cleaned wavelet transform through reconstruction filters generates a time-domain signal that still contains recognizable action potentials.




Figure 4.9: Diagram of the wavelet filterbank for multi-resolution analysis. The analysis filters split the signal into detail and approximation signals. The approximation signal was then split into second level detail and approximation signals, and continued filtering generated four detail signals and a coarse, fourth-level approximation signal. The combined data rate of the four detail signals and fourth-level approximation signal was equal to that of the original signal.

4.2.1 Design of the Filtering System

We implemented the discrete wavelet transform using a filterbank of cascaded wavelet analysis units, each consisting of a high-band and a low-band filter. These two filters separated their input signal into a detail signal and an approximation signal. In terms of multi-resolution analysis, the approximation was the projection of the input signal onto the next coarser approximation space, while the detail signal was the difference between the original signal and the approximation (Mallat, 1989). The filters each decimated the signals by a factor of two, and the combined data rate of the detail and approximation signals was equivalent to that of the input signal. By connecting four analysis units in cascade, with the approximation signal generated by one unit used as the input signal for the next, we generated a multi-resolution representation of the input signal, consisting of four detail signals and a coarse approximation signal (Figure 4.9).

The filterbanks used the *symlet3* wavelet basis (Figure 4.10). Symlets were chosen because of their approximate linear phase response, compact support, and because the symlet wavelet family has been noted as being well suited to model neural signals (Oweiss,

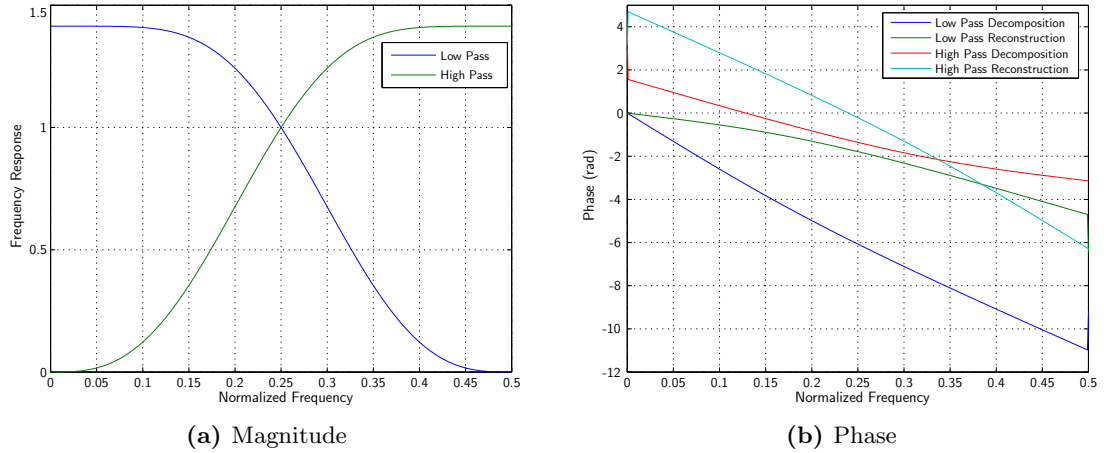


Figure 4.10: FFT of the symlet wavelet filters. Each filter had a pass band equal to half of the Nyquist frequency. The cascade of the decomposition and reconstruction filters had a linear phase response.

2006).

We took test signals for the algorithm using the first-generation IC, which we chose because it performed better at artifact elimination. We used an IOtech DAC488HR to apply waveforms containing simulated action potentials to the saline medium. We recorded waveforms with and without simulated action potentials, and with and without stimulation. Figure 4.11 shows the multi-resolution analysis of those signals.

For the signal without simulated action potentials, the artifact that remained after the discharge phase mostly appeared in the fourth-level approximation signal; this observation suggested that discarding the fourth-level approximation signal would remove most of the post-discharge artifact.

Examination of the analysis of the simulated action potentials in the absence of stimulation revealed that the first-level detail signal consisted mostly of noise, suggesting that this signal could also be discarded. There was still significant signal content in the fourth-level approximation signal; however, the addition of another level of resolution to further separate action potentials from artifacts would have significantly increased the delay associated with the algorithm. Even without the fourth-level approximation, the remaining three detail signals still contained enough information for reconstruction of the signal.

We designed a reconstruction filterbank in a similar manner to the analysis filterbank.

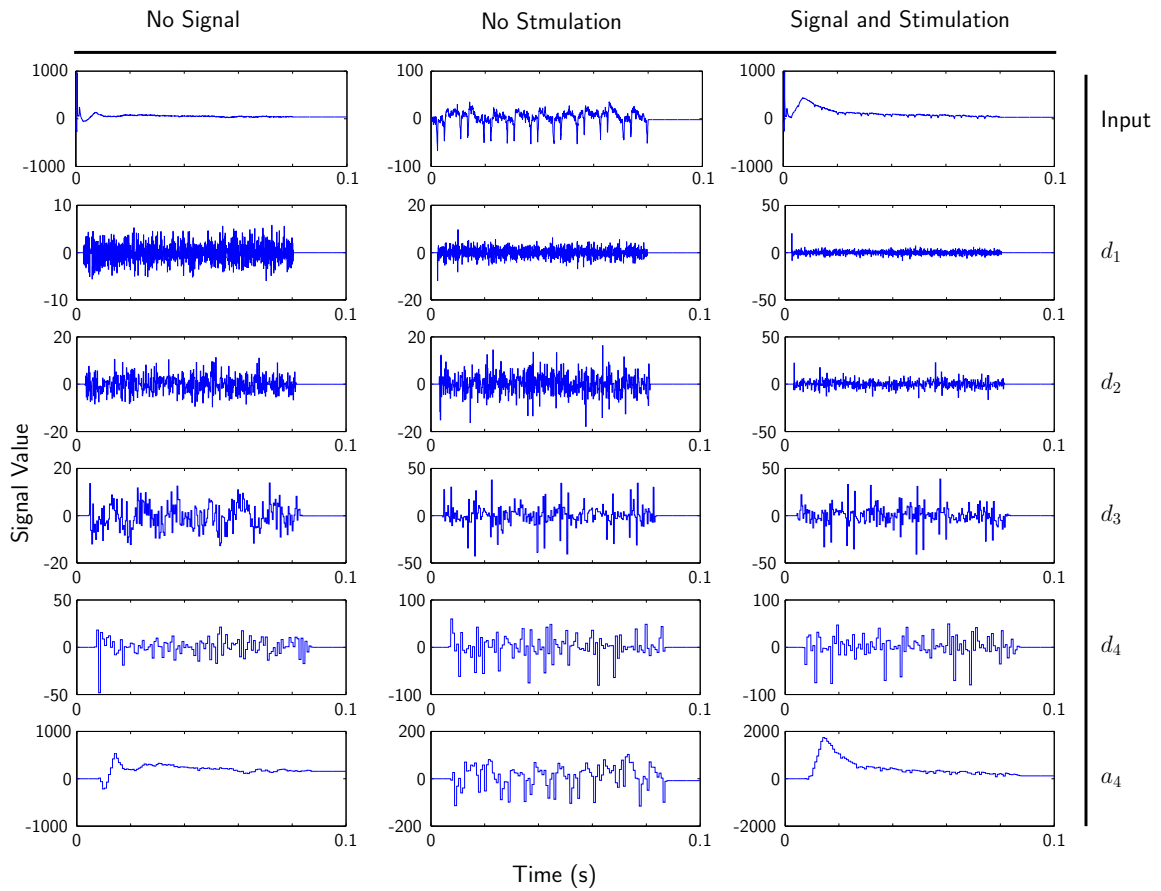


Figure 4.11: Multi-resolution analysis applied to signals recorded from the first-generation IC. The post-discharge artifact dominated the fourth-level approximation signal, and noise dominated the first-level detail signal, suggesting that we should discard both these resolutions.

Table 4.1: Error Normalized to Signal Power

		Threshold				
		1	4	16	64	256
Signal Maximum	9	0.7868	0.8566	1.0000	1.0000	1.0000
	17	0.6312	0.6239	0.7449	1.0000	1.0000
	30	0.5671	0.5775	0.7016	1.0000	1.0000
	52	0.5200	0.5255	0.5683	1.0000	1.0000
	93	0.5006	0.5009	0.5081	0.6441	1.0000

The first-level detail signal and the fourth-level approximation signal were excluded from the reconstruction, as suggested by our observation that noise and artifacts dominated those resolutions. We applied a threshold to the remaining resolutions before they were used as inputs to the reconstruction filterbank (Donoho, 1995).

To test the distortion introduced by the wavelet de-noising, we applied a test signal through the system at different amplitudes and for different thresholds applied to the wavelet coefficients. We observed that the output signals preserved many of the characteristics of the input signals, although for small signal amplitudes, large thresholds resulted in a total loss of the input signal (Figure 4.12). We computed the total squared error between the input and output signals, weighted by the total squared input signal, and these values are listed in Table 4.1.

A second simulation repeated the distortion analysis, but with white noise added to the input to the wavelet de-noising filter. We computed the weighted error between the original, noiseless signal and the output of the filterbank (Table 4.2). For small signals, the noise was larger than the signal, resulting in a large error. For larger signals and moderate thresholds; corresponding to column three, rows 2–5 of Figure 4.12; the filter was able to remove much of the noise, resulting in distortion levels similar to those resulting from filtering the noiseless signal.

One difficulty with this method is that the stimulation and artifact elimination generated transient signals at the same resolutions as the artifact. Blanking the signal for a short time (2–3 ms) from the onset of stimulation suppressed these transients. To avoid generating a transient at the end of blanking, the algorithm sampled the input signal at the end of

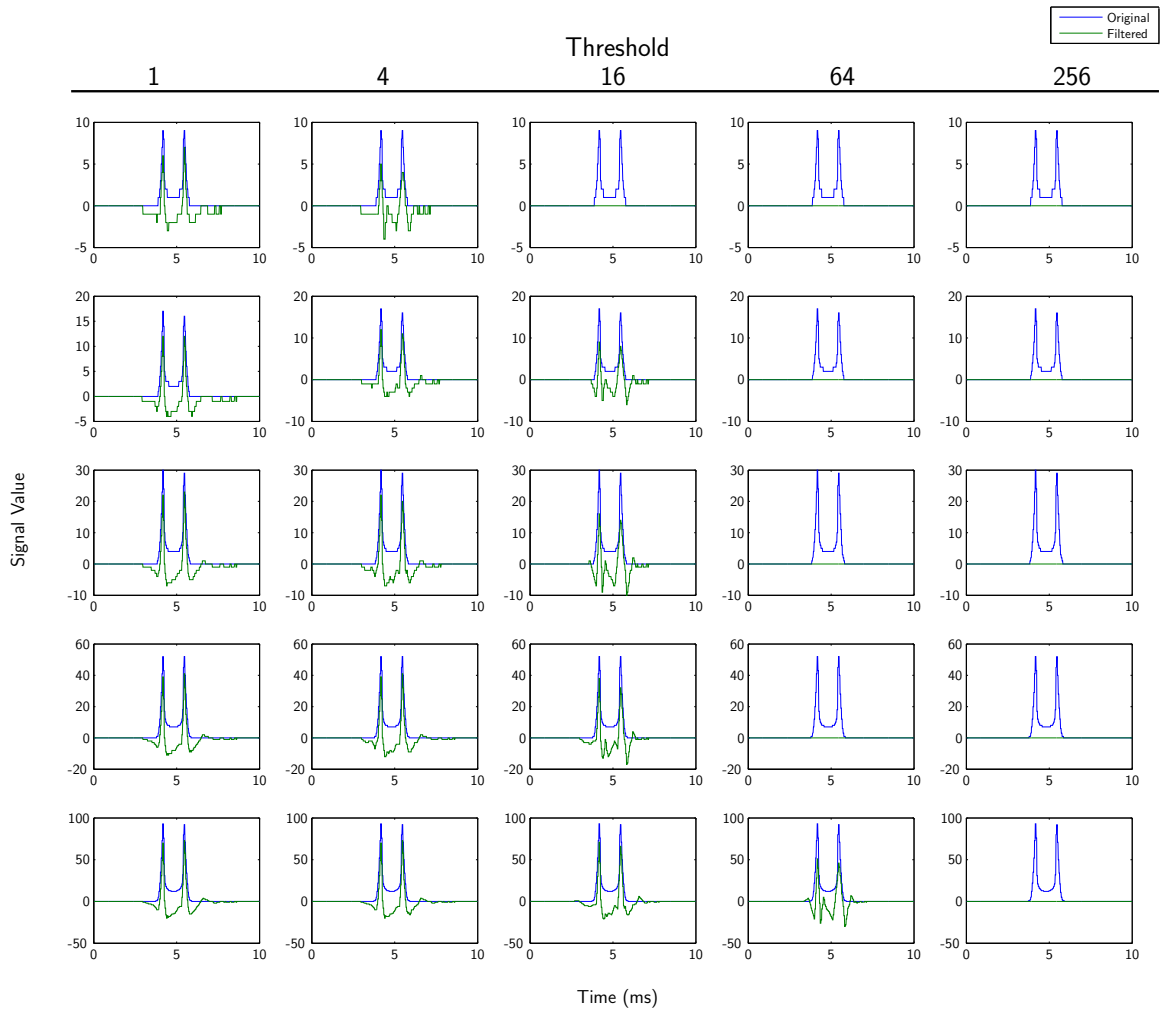


Figure 4.12: Effect of the wavelet filterbank, varying the signal amplitude and the threshold applied to the signal coefficients. For thresholds that were small relative to the signal amplitude, the output signal retained many of the characteristics of the input signal.

Table 4.2: Wavelet De-noising Error, Normalized to Signal Power, with White Noise

		Threshold				
		1	4	16	64	256
Signal Maximum	9	4.4698	4.4025	1.0000	1.0000	1.0000
	17	1.7076	1.5899	0.8682	1.0000	1.0000
	30	0.9215	0.8915	0.7619	1.0000	1.0000
	52	0.6233	0.6193	0.5738	1.0000	1.0000
	93	0.5328	0.5256	0.5197	0.7393	1.0000

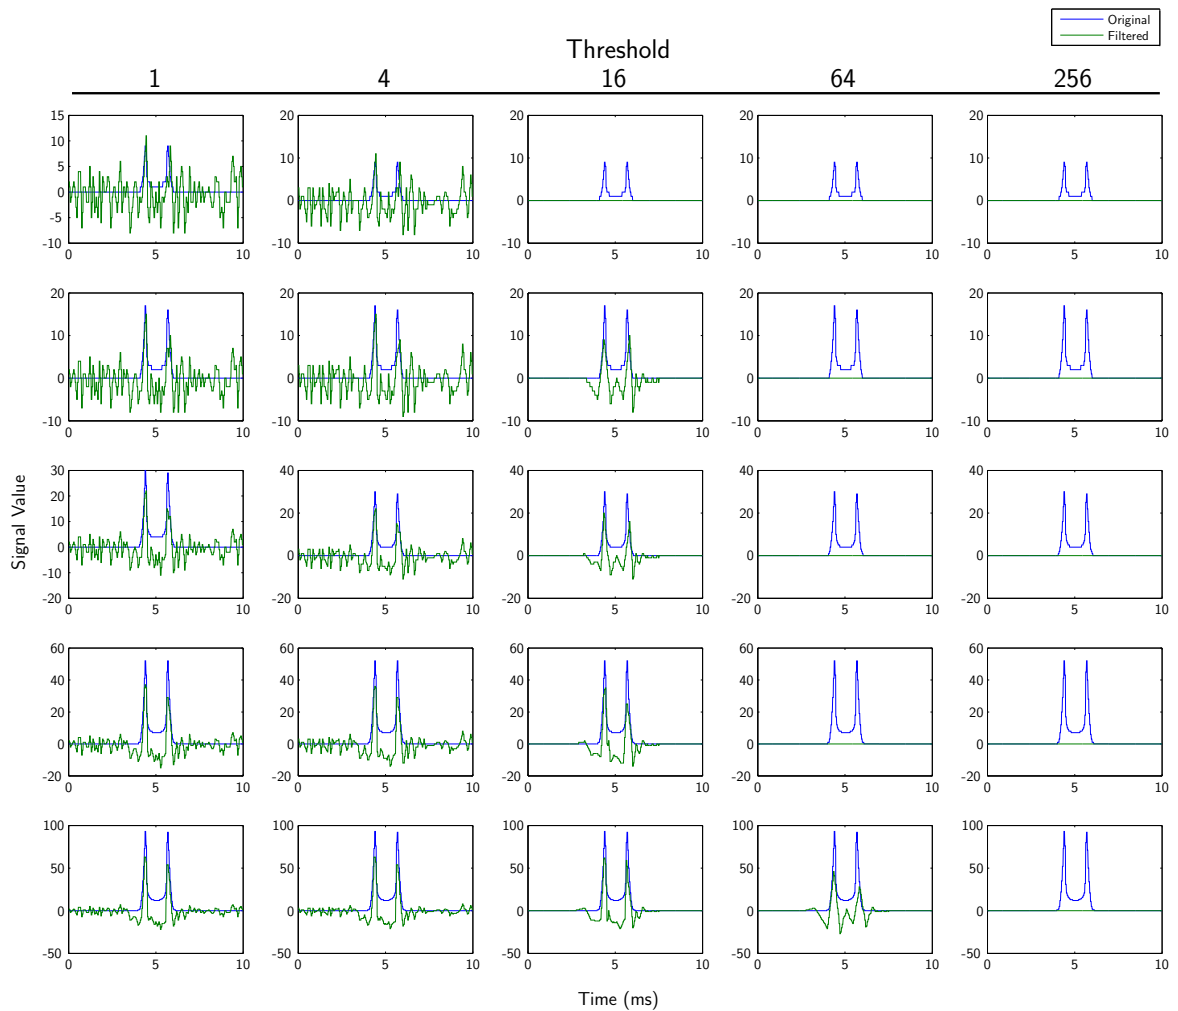


Figure 4.13: Effect of the wavelet filterbank, varying the signal amplitude and the threshold applied to the signal coefficients. White noise was added to the signals before wavelet filtering. For thresholds that were small relative to the signal amplitude, the output signal retained many of the characteristics of the input signal.




Figure 4.14: Block diagram of the FPGA filtering system. The blanking unit suppressed the initial effects of stimulation and artifact elimination, while the wavelet thresholding suppressed noise and the post-discharge artifact.

blanking and subtracted the initial value from the input signal before the wavelet transform. The subtractive correction linearly decayed to zero over a time scale sufficiently slow as to have minimally affected the resolutions used to reconstruct action potentials.

4.2.2 Experimental Characterization

To test the FPGA filtering system, we connected the first-generation IC to the FPGA through the ADC. We connected the DAC outputs to a National Instruments 6036E data acquisition system. We connected an Ayanda Biosystems multi-electrode array (MEA) to the inputs of the IC, and we applied test signals from the IOtech DAC488HR to the saline medium in the MEA through a bare wire electrode. We generated six test signals that random, simulated action potentials (Figure 4.15 shows one of the six signals). Some of the simulated action potentials were of the same form as the test signals in Figure 4.12, but most of them had only a single peak. The action potentials had a mean amplitude of $100\ \mu\text{V}$, with a standard deviation of $42\ \mu\text{V}$.

Using the data acquisition system, we captured waveforms of the IC and FPGA output voltages for 70 ms after a $500\ \mu\text{s}$ stimulation pulse. The artifact-elimination circuitry on the IC was able to prevent saturation of the recording system by discharging the stimulation electrode. The FPGA system was blanked for 2.8 ms, starting from the beginning of stimulation, to prevent the transient signals generated during stimulation and artifact elimination from interfering with the FPGA filtering. Even at a threshold value of zero, the FPGA algorithm was able to remove the post-discharge artifact (Figure 4.16). By comparison, a band-pass filter with cutoff frequencies that corresponded to the retained wavelet coefficients allowed a noticeable artifact to pass through. The band-pass filter did, however,

have a much shorter delay than the wavelet filtering. The delay associated with the wavelet filter approximately doubles with each additional level of decomposition, leading to a delay of 12.34 ms.

For each test signal, we repeated the recordings, adjusting the thresholds that we applied to the wavelet coefficients by logarithmically-spaced intervals. As we raised the wavelet coefficient threshold, the de-noising removed more of the noise that had been present in the signal (Figure 4.17). At low thresholds (Figure 4.17(a)–(d)), we observed no significant changes in the qualitative aspects of the waveform. At higher thresholds (Figure 4.17(f)–(g)), the filtering suppressed most of the noise in the original signal. At the highest threshold (Figure 4.17(h)), the filtering also significantly attenuated the action potentials.

We quantified the effects of the filtering by examining the noise present in the output of the filtering (for trials with no simulated action potentials) and by counting the number of simulated action potentials that we were able to observe. Coinciding with the increase in threshold was distortion of the simulated action potentials. We observed that the signal loss occurred at a higher threshold than the noise reduction (Figure 4.18), confirming our qualitative observation that the moderate threshold values reduced the noise while preserving the simulated action potentials.

4.3 Conclusions

In this chapter, we presented a filtering algorithm to remove noise and the post-discharge artifact from recordings made with the artifact-elimination IC. We implemented the algorithm using the Xilinx Virtex-4 FPGA. The implementation included data converters that used the computational resources of the FPGA, together with inexpensive external components, to provide analog input and output capabilities. The signal processing algorithm consisted of separating the signal out into multiple resolutions with the discrete wavelet transform. We applied a hard threshold to the resolutions that had the highest levels of signal compared to noise and artifact, and we used those resolutions to reconstruct the original signal without the post-discharge artifact. Combined with the ICs for artifact elimination, this filtering algorithm substantially reduced the effect of the stimulation artifact on neural




Figure 4.15: One of the six test signals used for characterizing the filtering algorithm. We randomly generated the sizes and times of the simulated action potentials.







Figure 4.16: Comparison of the original recording, in which the post-discharge artifact was clearly visible; the band-pass filtered signal, which still had a significant artifact; and the wavelet-filtered signal, which had very little remaining artifact.




(a) Threshold = 1




(b) Threshold = 2




(c) Threshold = 4




(d) Threshold = 8




(e) Threshold = 16



(f) Threshold = 32



(g) Threshold = 64



(h) Threshold = 128

Figure 4.17: Effect of varying the wavelet coefficient threshold.




Figure 4.18: Effect of the wavelet threshold on the noise levels and spike observations. As we increased the threshold that we applied to the wavelet coefficients, the input-referred noise of the filtered signal decreased. Our ability to observe action potentials, defined by signals that exceeded 2.5 standard deviations, also decreased as we increased the threshold; however, the loss in action potentials occurred at a higher threshold value.

interfacing.

In addition to the benefits for artifact elimination, the FPGA implementation of the filtering algorithm is important because it offers the possibility of digitally interfacing to data acquisition, analysis, and storage systems. Combined with the ability of the thresholded wavelet transform to generate a sparse representation of the action potential signals, this algorithm can also form the basis of a data-compression algorithm, reducing the data transfer requirements of large-scale systems. Future systems will be able to utilize digital data transfer links, which will ease the design of larger scale systems.

CHAPTER 5

DISCUSSIONS

In this dissertation we have presented a neural interfacing system designed for use in experimental studies of neural development and plasticity. The focus of our work was on the elimination of the stimulation artifact, which has limited the ability to observe the direct response of neural tissue to stimulation. Here, we review the design and performance of our system, discuss the novel contributions to the field of neural interfacing that the system entailed, and speculate on the impact and future development of the system.

5.1 Novel Contributions of this Work

Included in the content of this work were three important, novel contributions to the field of neural interfacing:

- A model of the stimulation artifact
- The use of feedback to eliminate the stimulation artifact
- The use of filtering to reduce the post-discharge artifact

Our first contribution was the development of a model for eliminating the stimulation artifact. Although researchers often mention the stimulation artifact as an experimental difficulty, there are very few investigations into its source (Mayer presents one of the few such studies (Mayer et al., 1992)). Even efforts at artifact-free recording fail to mention an underlying principle behind the approach (DeMichele and Troyk, 2003; Gnadt et al., 2003; Jimbo et al., 2003; Wagenaar and Potter, 2002). Through our model, we provide a tool for analyzing how the properties of the electrode affect the generation of the artifact. This is important, as it was through our model that we were able to understand the tradeoffs involved in the design of our system for artifact elimination.

Our second major contribution was our feedback-based circuit for artifact elimination. Unlike many approaches for filtering the stimulation artifact (Gnadt et al., 2003; Wageenaar and Potter, 2002), our circuitry dealt directly with the physical and electrochemical processes at the electrode that generate the artifact. Our design was also suitable for very large scale integration (VLSI) implementation, unlike the previous work in direct artifact-elimination (Jimbo et al., 2003). The combination of these two advantages makes our work uniquely suited for use in future, large-scale systems for interfacing to multi-electrode arrays (MEAs) with hundreds or even thousands of electrodes.

Our third contribution was a filtering system that augmented the artifact-elimination circuitry. The algorithm used filtering techniques for a novel purpose of removing the post-discharge artifact that remained after the use of the artifact-elimination circuitry. Through this filtering, we brought about the culmination of our goal of eliminating the stimulation artifact from neural recordings.

5.2 Summary of the Dissertation

The core of our system was an integrated circuit (IC) capable of multi-electrode stimulation and recording. The IC contained artifact elimination circuitry to minimize the interference with the recording circuitry that use of the stimulation circuitry generated. The development of our system was considered in three phases. First, we studied the physical and chemical mechanisms behind the stimulation artifact. Second, we applied knowledge of the stimulation artifact to the design of ICs. Finally, we considered additional filtering to remove what was left of the stimulation artifact.

5.2.1 Models of Stimulation Artifact Generation

First, we developed a circuit model of the generation of the stimulation artifact. The components in the model were a parallel resistor and capacitor, representing current flow in the electrode; a voltage source, representing the electrochemical half-cell potential of the electrode; a resistor, representing the medium; and a voltage source with a time-varying series resistance, representing a stimulation source. We considered models with linear and non-linear components for the electrode, and we used model parameters that approximated

gold, platinum, and platinum black electrodes.

We compared the artifacts that the models predicted to those that a physical system produced. Although the model did not fully agree with experimental observations, it did exhibit qualitatively similar behavior to physical electrodes. Most importantly, the model was suitable for incorporation into simulations of electronics, so that it served as a tool for the design of circuitry for the elimination of the stimulation artifact.

5.2.2 Integrated Circuits for Elimination of the Stimulation Artifact

Next, we used our model of stimulation artifact generation to develop a novel IC for stimulation and recording. The IC included artifact-elimination circuitry that countered the physical generation of the stimulation artifact.

Our design started with the choice of a low-noise amplifier topology. We based our design on that of Harrison, which we chose because of its power-efficient noise performance. To that base design, we made modifications to the dc feedback that were necessary for the artifact-elimination circuitry (Harrison and Charles, 2003). We analyzed the impact on the input-referred noise of our dc feedback circuitry, concluding that lowering the high-pass cutoff frequency introduced by the dc feedback resulted in lower noise levels.

In addition to the recording circuitry, we added a stimulation buffer and artifact-elimination circuitry. We designed the stimulation buffer to provide voltage pulses, and we used digital shift registers to control its activity. The artifact-elimination circuitry was necessary to prevent the stimulation pulses from saturating the recording amplifier. Our design for the artifact-elimination circuitry placed an amplifier in feedback with the recording amplifier that provided a discharge current to the electrode, returning it to its pre-stimulation voltage.

We presented an initial, first-generation design of the IC, and a revised, second-generation design. The first-generation design had input referred noise levels as low as $4.77 \mu\text{V}_{\text{rms}}$, which was slightly higher than desired. Also, the design of the stimulation circuitry limited the stimulation current to $9 \mu\text{A}$, which was not as large as desired. Our second-generation

design improved in these areas, reducing the noise to $3.50\ \mu\text{V}_{\text{rms}}$ and increasing the maximum stimulation current to over $50\ \mu\text{A}$.

5.2.3 Filtering to Augment the Integrated Circuits

Finally, we developed a filtering algorithm that augmented the artifact-elimination circuitry. The filtering reduced the post-discharge artifact that remained after the use of our IC. We used a field-programmable gate array (FPGA) to implement the filtering because of its advantages of reconfigurability and speed.

To connect the IC to the FPGA, we designed a successive approximation analog-to-digital converter (ADC) that used the FPGA as its control. For the reference voltage used in successive approximation, we used a second-order $\Sigma\Delta$ digital-to-analog converter (DAC) that used FPGA resources for its modulation loop. The converter required an external analog filter, comparator, and a switch. The levels of distortion in the ADC were equivalent to over 10 bits.

The filtering algorithm itself consisted of two parts. First, a blanking unit suppressed the input signal during the duration of stimulation and discharge. At the end of the blanking, we introduced a subtractive correction to the signal to prevent it from changing abruptly at the end of blanking. The correction decayed linearly to zero at a rate that was similar to that of the post-discharge artifact, so that had a minimal effect on the frequency bandwidth of neural signals. The second part of the filtering algorithm used the discrete wavelet transform to separate neural action potentials from the post-discharge artifact. We applied a threshold to the signal in the wavelet domain, which reduced the noise present in the signal.

We tested the filtering algorithm with simulated action potentials. We characterized the distortion generated by the filtering by computing the squared error between the original signal and the filtered signal. We found that, at a proper threshold level, the filtering algorithm was able to remove noise from input signals, such that the error between the filtered, noisy signal and the original, noiseless signal was comparable to the error between the filtered, noiseless signal and the original signal. We also were able to observe the simulated action potentials when they occurred during the post-discharge artifact.

5.3 Future Work

Possible improvements to the system presented include a revision of the artifact-elimination IC and additional functionality in the FPGA interface.

A future revision of the IC should adjust the sizes of the feedback capacitors, reducing the gain and input-referred noise levels at the expense of die area. The benefit to reducing the input-referred noise is straightforward, but the benefit to reducing the gain is not as readily apparent. Gain reduction will improve the stability of the artifact-elimination loop, and should reduce the effects of the post-discharge artifact. Additionally, alternate topologies should be considered for the output buffer, to reduce power dissipation and crosstalk. Potentially, the IC could include ADCs, which would help in connecting the IC to digital hardware, such as the FPGA.

Development on the FPGA should focus on a digital connection to a computer. Our work provides a start to a digital interface to a computer by converting the analog output of the IC into digital signals internal to the FPGA. The filtering algorithm we presented serves as an additional step toward reducing the data bandwidth necessary to transfer the neural recording to a computer, in that it effectively discriminates between neural activity, which must be transferred to the computer, and noise and artifacts, which should be discarded.

Perhaps the most significant future work lies not in the development of the system, but rather in its use. At the time of this document, neuroscientists have begun using the artifact-elimination IC for experimentation with neural tissue, and modest improvements in the functionality and usability of our system may lead to its wider adoption. It is our hope that the use of our system will enable new discoveries that would have been impossible with previous equipment.

APPENDIX A

SPICE MODELS FOR ELECTRODE SIMULATION

A.1 Linear Circuit Model

```
1  * Linear circuit model for the electrode
2  *
3  *****
4  * independent voltage sources
5  *****
6  Vstim stim 0 PWL(0 0 5ms 0 5.005ms 1 5.245ms 1 5.255ms -1
7  +      5.495ms -1 5.5m 0)
8  Vcont ctrl 0 PULSE(0 1 4.995ms 5us 5us 500us)
9
10 *****
11 * main circuits
12 * 3 different electrodes, each with a stim voltage
13 * and a switch
14 *****
15 EPtBlack PtB1 0 stim 0 1
16 SPtB PtB1 PtB2 ctrl 0 smodel1 off
17 XPtB PtB2 0 PtBlackElectrode
18
19 EAu Au1 0 stim 0 1
20 SAu Au1 Au2 ctrl 0 smodel1 off
21 XAu Au2 0 AuElectrode
22
23 EPt Pt1 0 stim 0 1
24 SPt Pt1 Pt2 ctrl 0 smodel1 off
25 XPt Pt2 0 PtElectrode
26
27 *****
28 * electrode subcircuits
29 *****
30 .SUBCKT PtBlackElectrode 1 2
31 Rt 1 3 3e7
32 C3 1 3 8n
33 Voc 3 4 0.05
34 Rs 4 2 10k
35 .ENDS
36
37 .SUBCKT AuElectrode 1 2
38 Rt 1 3 1.5e8
```

```

39 C3 1 3 300p
40 Voc 3 4 0.05
41 Rs 4 2 10k
42 .ENDS
43
44 .SUBCKT PtElectrode 1 2
45 Rt 1 3 7e7
46 C3 1 3 300p
47 Voc 3 4 0.05
48 Rs 4 2 10k
49 .ENDS
50
51 .MODEL smodel1 sw (vt=0.1 ron=100 roff=1e9)
52
53 *****
54 * anaylsis
55 *****
56 .OPTIONS CHGTOL=1e-19 RELTOL=1e-5
57 .TRAN 10u 30m 0 50n
58 .SAVE v(Pt2) v(Pt:3) v(PtB2) v(PtB:3) v(Au2) v(Au:3)
59 +      v(stim) v(ctrl)
60 .END

```

A.2 Circuit Model with Diodes

```

1  * Circuit model for the electrode, with diodes
2  *
3  *****
4  * independent voltage sources
5  *****
6  Vstim stim 0 PWL(0 0 5ms 0 5.005ms 1 5.245ms 1 5.255ms -1
7  +      5.495ms -1 5.5m 0)
8  Vcont ctrl 0 PULSE(0 1 4.995ms 5us 5us 500us)
9
10 *****
11 * main circuits
12 * 3 different electrodes, each with a stim voltage
13 * and a switch
14 *****
15 EPtBlack PtB1 0 stim 0 1
16 SPtB PtB1 PtB2 ctrl 0 smodel1 off
17 XPtB PtB2 0 PtBlackElectrode
18
19 EAu Au1 0 stim 0 1
20 SAu Au1 Au2 ctrl 0 smodel1 off
21 XAu Au2 0 AuElectrode

```

```

22
23 EPt Pt1 0 stim 0 1
24 SPt Pt1 Pt2 ctrl 0 smodel1 off
25 XPt Pt2 0 PtElectrode
26
27 *****
28 * electrode subcircuits
29 *****
30 .SUBCKT PtBlackElectrode 1 2
31 D1 1 3 diode1
32 D2 3 1 diode1
33 C3 1 3 8n
34 Voc 3 4 0.05
35 Rs 4 2 10k
36 .MODEL diode1 D(IS=800p)
37 .ENDS
38
39 .SUBCKT AuElectrode 1 2
40 D1 1 3 diode1
41 D2 3 1 diode1
42 Rt 1 3 1.5e8
43 C3 1 3 300p
44 Voc 3 4 0.05
45 Rs 4 2 10k
46 .MODEL diode1 D(IS=200p)
47 .ENDS
48
49 .SUBCKT PtElectrode 1 2
50 D1 1 3 diode1
51 D2 3 1 diode1
52 Rt 1 3 7e7
53 C3 1 3 300p
54 Voc 3 4 0.05
55 Rs 4 2 10k
56 .MODEL diode1 D(IS=400p)
57 .ENDS
58
59 .MODEL smodel1 sw (vt=0.1 ron=100 roff=1e9)
60
61 *****
62 * anaylsis
63 *****
64 .OPTIONS CHGTOL=1e-19 RELTOL=1e-5
65 .TRAN 10u 30m 0 50n
66 .SAVE v(Pt2) v(Pt:3) v(PtB2) v(PtB:3) v(Au2) v(Au:3)
67 +      v(stim) v(ctrl)
68 .END

```

APPENDIX B

REVIEW OF ELECTRONIC NOISE

Electronic circuits are subject to inherent noise mechanisms. Although these noise mechanisms are random processes, their statistics are subject to quantitative analysis with statistical methods. Noise sources, whether modeled as voltages or currents, are characterized according to their mean and variance, rather than their instantaneous values. The mean, or expected value, of a random variable is given by

$$\bar{x} = \int_{-\infty}^{\infty} x f(x) dx \quad (\text{B.1})$$

where $f(x)$ is the probability density function of the variable x . Noise processes are generally zero mean, as any mean value would be better described as a dc offset. The most important summary measure of a noise process is its variance,

$$\sigma_x^2 = \int_{-\infty}^{\infty} (x - \bar{x})^2 f(x) dx \quad (\text{B.2})$$

which is a measure of the power of the noise process. Because noise processes are zero mean, (B.2) simplifies to

$$\sigma_x^2 = \int_{-\infty}^{\infty} x^2 f(x) dx \quad (\text{B.3})$$

which is equivalent to the *mean square* value of x , also symbolized as $\overline{x^2}$. A related measure is the root-mean-square (RMS) value, $\sqrt{\overline{x^2}}$, which gives the dc value that results in the same power.

A noise model of a physical component consists of an ideal, noiseless component augmented by one or more noise sources. We can analyze the complete circuit, which consists of noiseless components and separate noise sources, using circuit-analysis techniques similar to those that apply to deterministic sources.

B.1 Noise in Resistors

In a resistor, thermal fluctuations in the charge carriers result in a noise voltage (or, by Norton equivalence, a noise current); an effect that was experimentally characterized by Johnson (Johnson, 1928). Nyquist derived (B.4) based on transmission line theory and the equipartition theorem of statistical mechanics (Nyquist, 1928).

The derivation considers the possible modes for power transfer between two resistors, R , connected by a lossless transmission line, having characteristic impedance R (see Figure B.1). If the ends of the transmission lines are short-circuited, it has characteristic modes of resonance at frequencies of $nv/2l$, where v is the propagation velocity along the line, l is the length of the line, and n is a positive integer. Each mode corresponds to two degrees of freedom—one electric, one magnetic—and by the equipartition theorem, there is an energy of $kT/2$ associated with each degree of freedom, where k is the Boltzmann constant (1.38×10^{-23} J/K) and T is the thermodynamic temperature. A given frequency bandwidth, Δf , contains $2l\Delta f/v$ resonance modes, corresponding to an amount of energy equal to $2lkT\Delta f/v$. Without the reflections at either end of the transmission line, the resistors provide the energy to sustain these oscillations during a time of l/v , which is equivalent to a noise power of $P_{\text{tn}} = 2kT\Delta f$. Each resistor provides half of the power, so the power per resistor is $P_{\text{tn}} = kT\Delta f$. We can model the source of this power as a voltage source in series with the resistor. The voltage source will create a noise current of $i_{\text{tn}} = v_{\text{tn}}/2R$, and it will deliver a power of $\overline{i_{\text{tn}}^2}R$ to the other resistor. Combining the above equations gives

$$\overline{v_{\text{tn}}^2} = 4kTR\Delta f \quad (\text{B.4})$$

or equivalently,

$$\overline{i_{\text{tn}}^2} = \frac{4kt}{R}\Delta f. \quad (\text{B.5})$$

Resistors also have an additional noise component, often called flicker noise. This noise source follows the equation

$$\overline{v_{\text{ex}}^2} = \frac{K_{\text{f}}V^2\Delta f}{f} \quad (\text{B.6})$$




Figure B.1: Model circuit for deriving the thermal noise in a resistor




Figure B.2: Models of noise in a resistor. The noise can be modeled as a current or voltage, according to Thévenin and Norton equivalent circuits.

where K_f is the flicker noise coefficient and V is the dc voltage across the resistor (Leach). Because of the inverse dependence on frequency, this type of noise is also known as $1/f$ noise.

The total noise voltage is the sum of the two noise voltages,

$$v_n = v_{tn} + v_{ex}. \quad (\text{B.7})$$

Because these two noise sources are uncorrelated, their mean square values also sum, according to

$$\overline{v_n^2} = \overline{v_{tn}^2} + \overline{v_{ex}^2}. \quad (\text{B.8})$$

The noise sources augment the noiseless resistor in a circuit, as Figure B.2 shows.

B.2 Thermal noise in Electrodes

Physical electrodes have some real impedance component, introducing noise according to

$$\overline{v_{noise}^2} = 4kT\Re(Z) \Delta f, \quad (\text{B.9})$$

where $\Re(Z)$ is the real part of the the complex electrode impedance, Z .

B.3 Noise in MOSFETs

There are two noise mechanisms that occur in a metal–oxide–semiconductor field-effect transistor (MOSFET) transistor: shot noise, which is associated with any dc current flow; and flicker noise, which stems from various device imperfections and other higher order effects.

Shot noise is a specific case of a general noise mechanism known as *thermal noise*, in which the random vibrations related to thermal energy are the underlying source of noise. At relatively low frequencies, thermal noise has a power spectrum that is independent of frequency; thus, thermal noise also goes by the term *white noise*. We can model shot noise as a noise current source connected from drain to source. The noise is proportional to the dc current in the transistor, according to

$$\overline{i_n^2} = 2qI_{DS}\Delta f. \quad (\text{B.10})$$

This noise can be shown to be equivalent to thermal noise, according to

$$\overline{i_n^2} = \frac{2kTg_m}{\kappa} \quad (\text{B.11})$$

for transistors operating under sub-threshold saturation (Liu et al., 2002).

The additional source of noise, flicker noise, has a power spectrum with a $1/f$ characteristic, according to

$$\overline{v_{\text{ex}}^2} = \frac{K}{WLf} \quad (\text{B.12})$$

where K is a constant associated with the fabrication process, and WL is the transistor gate area.

The noise sources in (B.11) and (B.12) combine through superposition to give a total noise model of the MOSFET, as shown in Figure B.3. Both noise sources are dependent on design choices. The thermal noise component is dependent on g_m , such that increasing the dc bias current of the transistor will result in lower input-referred noise. The $1/f$ noise component is dependent on the transistor geometry, such that increasing the gate area will reduce input-referred noise.




Figure B.3: Noise model of a MOSFET. There are two noise sources, a noise current modeling thermal noise, and a noise voltage modeling $1/f$ noise.

APPENDIX C

REVIEW OF THE WAVELET TRANSFORM

Traditional signal processing occurs in two domains: time and frequency. These domains are related by the Fourier transform of a signal $x(t)$,

$$X(\omega) = \int_{-\infty}^{\infty} x(t) e^{-j\omega t} dt \quad (\text{C.1})$$

or its discrete-time counterpart,

$$X(e^{j\omega}) = \sum_{n=-\infty}^{\infty} x[n] e^{-j\omega n}. \quad (\text{C.2})$$

where ω is the frequency variable. In many applications, the Fourier transform cannot be used because the infinite series will not converge. For these applications, the Laplace transform,

$$X(s) = \int_0^{\infty} x(t) e^{-st} dt \quad (\text{C.3})$$

where $s = \sigma + j\omega$ is the complex frequency; or its discrete-time counterpart the z -transform,

$$X(z) = \sum_{n=-\infty}^{\infty} x[n] z^{-n}, \quad (\text{C.4})$$

where z is the transform variable, are a powerful alternatives. These transforms are a powerful tool for the analysis of linear, time-invariant filters because they turn the difficult operation of convolution into the simpler one of multiplication.

The drawback of these transforms is that they impose a dichotomy between time and frequency. Many signals, however, exhibit time-varying frequency content. Analysis of these signals requires a mathematical technique that considers time and frequency at the same time.

C.1 Continuous Wavelet Transform

One popular transform that is capable of revealing the time-varying frequency content of signals is the wavelet transform, which divides a signal into bins of varying times and scales

								$\gamma_{4.1}$																$\gamma_{4.2}$							
				$\gamma_{3.1}$								$\gamma_{3.2}$								$\gamma_{3.3}$								$\gamma_{3.4}$			
$\gamma_{2.1}$		$\gamma_{2.2}$		$\gamma_{2.3}$		$\gamma_{2.4}$		$\gamma_{2.5}$		$\gamma_{2.6}$		$\gamma_{2.7}$		$\gamma_{2.8}$																	
$\gamma_{1.1}$	$\gamma_{1.2}$	$\gamma_{1.3}$	$\gamma_{1.4}$	$\gamma_{1.5}$	$\gamma_{1.6}$	$\gamma_{1.7}$	$\gamma_{1.8}$	$\gamma_{1.9}$	$\gamma_{1.10}$	$\gamma_{1.11}$	$\gamma_{1.12}$	$\gamma_{1.13}$	$\gamma_{1.14}$	$\gamma_{1.15}$	$\gamma_{1.16}$																
Time																															

Figure C.1: The wavelet transform divides a signal into time–scale bins, providing an intermediate representation between time and frequency domains.

(see Figure C.1). The wavelet transform is especially suited for finding regions in which the frequency content of a signal is changing, such as an edge, a discontinuity, or an action potential.

The wavelet transform projects a signal onto a basis of orthogonal wavelet functions,

$$\psi_{j,k}(t) = |j|^{1/2} \psi\left(\frac{t-k}{j}\right) \quad (\text{C.5})$$

which are translations by k and dilations by a factor j of a mother wavelet, $\psi(x)$ (Daubechies, 1990). Using these basis functions, we can calculate the projection a signal onto each function,

$$\gamma_{j,k} = \int_{-\infty}^{\infty} f(t) \psi_{j,k}^*(t) dt \quad (\text{C.6})$$

A requirement for ψ is that it must be square integrable (Grossmann et al., 1985), according to

$$\int_{-\infty}^{\infty} |\psi(t)| dt < \infty \quad (\text{C.7})$$

and

$$\int_{-\infty}^{\infty} |\psi(t)|^2 dt < \infty \quad (\text{C.8})$$

Although not necessary for the evaluation of the wavelet transform integral, there is an additional *admissibility condition*,

$$\int_{-\infty}^{\infty} \frac{|\Psi(\omega)|^2}{|\omega|} d\omega < \infty. \quad (\text{C.9})$$

An important property of a wavelet family is the number of vanishing moments, which relates to the effectiveness of the wavelet at decomposing polynomials. A wavelet with M

vanishing moments can model a polynomial of the same order. A wavelet has M vanishing moments when for all integers $m < M$,

$$\int_{-\infty}^{\infty} t^m \psi(t) dt = 0. \quad (\text{C.10})$$

C.2 Frequency Domain Analysis of Wavelets

It is useful to analyze a wavelet function in the frequency domain. The form of (C.6) is that of a convolution equation,

$$\gamma_k = (f * \psi_j)(k) \quad (\text{C.11})$$

suggesting that it is appropriate to represent a wavelet as a filter. We can determine from the admissibility condition (C.9) that the mean value of the wavelet is zero, which is equivalent to a band-pass filter. By the dilation property of the Fourier transform, the dilations of the mother wavelet correspond to narrower band-pass filters at lower frequencies. Successive dilation can cover most of the most of the frequency spectrum; however, Because of the band-pass nature of the wavelet transform, it is impossible to fully cover the frequency spectrum with a finite amount of wavelets. To overcome this difficulty, Mallat introduced the *scaling function*,

$$\phi(t) = \sum_{j,k} \gamma(j,k) \psi_{j,k}(t) \quad (\text{C.12})$$

which must adhere to the property

$$\int_{-\infty}^{\infty} \phi(t) dt = 1. \quad (\text{C.13})$$

In the frequency domain, the scaling function corresponds to a low-pass filter, complementing the band-pass nature of the wavelet function.

C.3 Discrete Wavelet Transform

As with the Fourier and Laplace transforms, there is a discrete-time analog to the continuous wavelet transform, known as the discrete wavelet transform. This transform is similar to that of (C.6); however, the mother wavelet, ψ , is a sampled function, and because of its sampled nature, we can only dilate and shift the mother wavelet by discrete intervals.

One important property of discrete wavelets is that it is possible to choose an orthonormal basis of wavelet functions, such that

$$\sum_n \psi_{a,b}[n] \psi_{j,k}^*[n] = \begin{cases} 1 & a = j, b = k, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{C.14})$$

There are a wide variety of discrete wavelet functions that follow (C.14). One of the simplest is the Haar wavelet,

$$f(x) = \begin{cases} 1 & 0 \leq x < 1/2, \\ -1 & 1/2 \leq x < 1, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{C.15})$$

Although this wavelet is easy to implement, it has only one vanishing moment. Daubechies proposed a family of discrete wavelets of support $2M$ and M vanishing moments, of which the Haar wavelet is the case $M = 1$ (Daubechies, 1992). The definition of these wavelets depends on the scaling function, and there is no closed-form expression of the wavelet itself. The symlets are a closely related wavelet family to the Daubechies wavelets, with a slight modification to increase their symmetry.

C.4 Multi-resolution Analysis and Wavelet Filterbanks

The wavelet transform provides a method for multi-resolution analysis of signals, in which we can examine the structure of the signal at different scales or resolutions. Low-resolution approximations retain the overall properties of the signal while representing the signal with a smaller number of samples. This concept was put forward by Mallat (Mallat, 1989), and we summarize his results here.

Central to the multi-resolution analysis is the concept of the approximation spaces. At a resolution of 2^{-j} , there is a vector space, \mathbf{V}_j , that contains all possible signal approximations. Approximations at finer resolutions include all the detail of a coarser approximation, such that

$$\mathbf{V}_j \subset \mathbf{V}_{j+1}. \quad (\text{C.16})$$

Finding the approximation signal requires projecting a signal onto the space of approximation signals. This projection is equivalent to applying a low-pass filter to the original signal. The filter is related to the scaling function by

$$h[n] = \left\langle \frac{1}{2} \phi \left[\frac{x}{2} \right], \phi[n] \right\rangle. \quad (\text{C.17})$$

Complementing the approximations are the detail signals, which are the differences between approximations at different scales. The detail signal is a projection onto a space orthogonal to the space of approximation signals. To generate the detail signal, we can filter with the *quadrature mirror filter* of the approximation filter,

$$G(\omega) = e^{-j\omega} H^*(\omega + \pi). \quad (\text{C.18})$$

This detail filter has a high-pass characteristic, and it is associated with the wavelet function, according to

$$\Psi(\omega) = G\left(\frac{\omega}{2}\right) \Phi\left(\frac{\omega}{2}\right). \quad (\text{C.19})$$

We can generate the approximation and detail signals at different scales by repeatedly applying the approximation and detail filters (Figure C.2). Because the outputs of the detail and approximation filters are inherently coarser than the input signal, we can downsample the outputs without any information loss. The effect of the decomposition (detail and approximation) filters is to split the original signal into two signals, each with half the data rate and information content of the original signal.

We can also reconstruct the original signal from the detail and approximation signals, using reconstruction filters. The reconstruction filters are time-reversed versions of the decomposition (detail and approximation) filters.




Figure C.2: Wavelet decomposition filterbank. Two filters split an approximation signal into a coarser approximation and a detail signal. The detail and approximation filters are quadrature mirror filters. The two coarser signals together contain all the information of the original signal. The filters can be cascaded to produce successively coarser approximations.




Figure C.3: Wavelet reconstruction filterbank. The reconstruction filters combine the detail and approximation filters to produce a finer approximation signal. The reconstruction filters have step responses that are time-reversed versions of those of the decomposition filters.

REFERENCES

- J. N. Y. Aziz, R. Genov, B. L. Bardakjian, M. Derchansky, and P. I. Carlen. Brain–silicon interface for high-resolution *in vitro* neural recording. *IEEE Trans. Biomed. Circuits Syst.*, 1(1):56–62, Mar. 2007.
- S. L. BeMent, K. D. Wise, D. J. Anderson, K. Najafi, and K. L. Drake. Solid-state electrodes for multichannel multiplexed intracortical neuronal recording. *IEEE Trans. Biomed. Eng.*, 33(2):230–241, Feb. 1986.
- W. Bialek, F. Rieke, R. R. de Ruyter van Steveninck, and D. Warland. Reading a neural code. *Science*, 252(5014):1854–1857, June 1991.
- R. A. Blum, J. D. Ross, S. K. Das, E. A. Brown, and S. P. DeWeerth. Models of stimulation artifacts applied to integrated circuit design. In *Proc. of the IEEE Engineering in Medicine and Biology Conference*, pages 4075–4078, San Fransisco, CA, Sept. 2004.
- R. A. Blum, J. D. Ross, E. A. Brown, and S. P. DeWeerth. An integrated system for simultaneous, multichannel neuronal stimulation and recording. *IEEE Trans. Circuits Syst. I*, 2007. In Press.
- M. Bolton, S. Nagheenanajar, and P. Gibbons. An application of roundoff error in financial calculations. TPS report, Initech, 1999.
- D. A. Borkholder. *Cell Based Sensors Using Microelectrodes*. PhD thesis, Stanford University, Nov. 1998.
- E. A. Brown, J. D. Ross, R. A. Blum, Y. Nam, B. C. Wheeler, and S. P. DeWeerth. Stimulation artifacts: a multi-electrode recording, stimulation, and artifact elimination system. In progress.
- J. C. Candy. A use of double integration in sigma delta modulation. *IEEE Trans. Commun.*, 33(3):249–258, Mar. 1985.
- D. L. Chapman. *Phil. Mag.*, 25(6):475, 1913.
- Y. Choi, M. A. McClain, M. C. LaPlaca, A. B. Frazier, and M. G. Allen. Three dimensional MEMS microfluidic perfusion system for thick brain slice cultures. *Biomed. Microdev.*, 9(1):7–13, Feb. 2007.
- E. Claverol-Tinture and J. Pine. Extracellular potentials in low-density dissociated neuronal cultures. *J. Neurosci. Meth.*, 117(1):13–21, May 2002.
- D. K. Cullen, J. Vukasinovic, A. Glezer, and M. C. LaPlaca. Microfluidic engineered high cell density three-dimensional neural cultures. *J. Neural Eng.*, 4(2):159–172, June 2007.
- W. Dabrowski, P. Grybos, and A. M. Litke. A low noise multichannel integrated circuit for recording neuronal signals using microelectrode arrays. *Biosens. Bioelectron.*, 19(7):749–761, Feb. 2004.
- L. I. Daikhin, A. A. Kornyshev, and M. Urbakh. Double-layer capacitance on a rough metal surface. *Phys. Rev. E*, 53(6), June 1996.

- I. Daubechies. *Ten Lectures on Wavelets*. SIAM, 1992.
- I. Daubechies. The wavelet transform, time-frequency localization and signal analysis. *IEEE Trans. Info. Theory*, 36(5):961–1005, Sept. 1990.
- R. de Levie. The influence of surface roughness of solid electrodes on electrochemical measurements. *Electrochim. Acta*, 10(2):113–130, Feb. 1965.
- T. Delbrück and C. A. Mead. Adaptive photoreceptor with wide dynamic range. In *Proc. of the IEEE Int. Symp. Circuits and Systems*, volume 4, pages 339–342, 1994.
- T. B. DeMarse, D. A. Wagenaar, A. W. Blau, and S. M. Potter. The neurally controlled animat: Biological brains acting with simulated bodies. *Auton. Robot.*, 11(3):305–310, Nov. 2001.
- G. A. DeMichele and P. R. Troyk. Stimulus-resistant neural recording amplifier. In *Proc. IEEE EMBS*, volume 4, pages 3329–3332, Cancun, Mexico, Sept. 2003.
- D. L. Donoho. De-noising by soft-thresholding. *IEEE Tran. Info. Theory*, 41(3):613–627, May 1995.
- B. Eversmann, M. Jenkner, F. Hofmann, C. Paulus, R. Brederlow, B. Holzapfl, P. Fromherz, M. Merz, M. Brenner, M. Schreiter, R. Gabl, K. Plehnert, M. Steinhauser, G. Eckstein, D. Schmitt-Landsiedel, and R. Thewes. A 128×128 CMOS biosensor array for extracellular recording of neural activity. *IEEE J. Solid-State Circuits*, 38(12):2306–2317, Dec. 2003.
- A. M. Feltham and M. Spiro. Platinized platinum electrodes. *Chemical Reviews*, 71(2):177–193, Apr. 1971.
- W. Franks, I. Schenker, P. Schmutz, and A. Hierlemann. Impedance characterization and modeling of electrodes for biomedical applications. *IEEE Trans. Biomed. Eng.*, 52(7):1295–1302, July 2005.
- H. Fricke. The theory of electrolytic polarization. *Philos. Mag.*, 14:310–318, 1932.
- L. A. Geddes, K. S. Foster, J. Reilly, W. D. Voorhees, J. D. Bourland, T. Ragheb, and N. E. Fearnot. The rectification propoerties of an electrode–electrolyte interface operated at high sinusoidal current density. *IEEE Trans. Biomed. Eng.*, 34(9):669–672, Sept. 1987.
- R. C. Gesteland, B. Howland, J. Lettvin, and W. H. Pitts. Comments on microelectrodes. *Proc. IRE*, 47:1856–1862, Nov. 1959.
- J. W. Gnadt, S. D. Echols, A. Yildirim, H. Zhang, and K. Paul. Spectral cancellation of microstimulation artifact for simultaneuos neural recording *In Situ*. *IEEE Trans. Biomed. Eng.*, 50(10):1129–1135, Oct. 2003.
- M. Gouy. Sur la constitution de la charge électrique a la surface d’un electrolyte. *J. Phys.*, 9:457–468, 1910.
- D. C. Grahame. The electrical double layer and the theory of electrocapillarity. *Chem. Rev.*, 41(3):441–501, Dec. 1947.

- P. R. Gray and D. A. Hodges. All-MOS analog–digital conversion techniques. *IEEE Trans. Circuits Syst.*, 25(7):482–489, July 1978.
- G. W. Gross. Simultaneous single unit recording *in vitro* with a photoetched laser deinsulated gold multimicroelectrode surface. *IEEE Trans. Biomed. Eng.*, 26(5):273–279, May 1979.
- A. Grossmann, J. Morlet, and T. Paul. Transforms associated to square integrable group representations. I. general results. *J. Math. Phys.*, 26(10):2473–2479, Oct. 1985.
- A. E. Grumet, J. L. Wyatt, Jr, and J. F. Rizzo, III. Multi-electrode stimulation and recording in the isolated retina. *J. Neurosci. Meth.*, 101(1):31–42, Aug. 2000.
- R. R. Harrison and C. Charles. A low-power, low-noise CMOS amplifier for neural recording applications. *IEEE J. Solid-State Circuits*, 38(6), June 2003.
- R. R. Harrison, P. T. Watkins, R. J. Kier, R. O. Lovejoy, D. J. Black, B. Greger, and F. Solzbacher. A low-power integrated circuit for a wireless 100-electrode neural recording system. *IEEE J. Solid-State Circuits*, 42(1):123–133, Jan. 2007.
- H. L. Helmholtz. Studien über electrische grenzsichten.
- A. L. Hodgkin and A. F. Huxley. The components of membrane conductance in the giant axon of *Loligo*. *J. Physiol.*, 116(4):473–496, 1952.
- A. C. Hoogerwerf and K. D. Wise. A three-dimensional microelectrode array for chronic neural recording. *IEEE Trans. Biomed. Eng.*, 41(12):1136–1146, Dec. 1994.
- H. Inose and Y. Yasuda. A unity bit coding method by negative feedback. *Proc. IEEE*, 51(11):1524–1535, Nov. 1963.
- J. Ji and K. D. Wise. An implantable CMOS circuit interface for multiplexed microelectrode recording arrays. *IEEE J. Solid-State Circuits*, 27(3):433–443, Mar. 1992.
- Y. Jimbo, T. Tateno, and H. Robinson. Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons. *Biophys. J.*, 76(2):670–678, Feb. 1999.
- Y. Jimbo, N. Kasai, K. Torimitsu, T. Tateno, and H. Robinson. A system for MEA-based multisite stimulation. *IEEE Trans. Biomed. Eng.*, 50(2):241–248, Feb. 2003.
- J. B. Johnson. Thermal agitation of electricity in conductors. *Phys. Rev.*, 32(1):97–109, July 1928.
- G. T. A. Kovacs. Introduction to the theory, design, and modeling of thin-film microelectrodes for neural interfaces. In D. A. Stenger and T. M. McKenna, editors, *Enabling Technologies for Cultured Neural Networks*, chapter 7, pages 121–166. Academic, San Diego, CA, 1994.
- W. M. Leach. Dr. Leach’s noise potpourri. URL <http://users.ece.gatech.edu/~mleach/ece6416/>. Accessed December, 2001.
- S.-C. Liu, J. Kramer, G. Indiveri, T. Delbrück, and R. Douglas. *Analog VLSI: Circuits and Principles*. MIT, 2002.

- J. Logue. Virtex analog to digital converter. Technical report, Xilinx, Sept. 1999.
- S. G. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Trans. Pattern Anal. Machine Intell.*, 11(7):674–693, July 1989.
- K. Manetakis and C. Toumazou. A new high-frequency very low output-impedance CMOS buffer. In *Proc. of the IEEE Int. Symp. Circuits and Systems*, volume 1, pages 485–487, Atlanta, GA, May 1996.
- C. A. Marrese. Preparation of strongly adherent platinum black coatings. *Anal. Chem.*, 59(1):217–218, Jan. 1987.
- S. Martinoia, L. Bonzano, M. Chiappalone, M. Tedesco, M. Marcoli, and G. Maura. In vitro cortical neuronal networks as a new high-sensitive system for biosensing applications. *Biosens. Bioelectron.*, 20(10):2071–2078, Apr. 2005.
- S. Mayer, L. A. Geddes, J. D. Bourland, and L. Ogborn. Electrode recovery potential. *Annal. Biomed. Eng.*, 20:385–394, 1992.
- E. T. McAdams, A. Lacknermeier, J. A. McLaughlin, D. Macken, and J. Jossinet. The linear and non-linear electrical properties of the electrode-electrolyte interface. *Biosens. Bioelectron.*, 10(1–2):67–74, 1995.
- C. Mead. *Analog VLSI and Neural Systems*. Addison-Wesley, 1989.
- H. Melville. *Moby-Dick; or, the Whale*. Harper, New York, 1851.
- MOSIS. MOSIS integrated circuit fabrication service. URL <http://www.mosis.org>. Accessed July 5, 2007.
- Multi Channel Systems. Multi Channel Systems. URL <http://www.multichannelsystems.com/index.html>. Accessed July 5, 2007.
- K. Najafi and K. D. Wise. An implantable multielectrode array with on-chip signal processing. *IEEE J. Solid-State Circuits*, 21(6):1035–1044, Dec. 1986.
- K. Najafi, K. D. Wise, and T. Mochizuki. A high-yield IC-compatible multichannel recording array. *IEEE Trans. Electron Devices*, 32(7):1206–1211, July 1977.
- NCSU. NCSU electronic design automation (EDA) wiki. URL http://www.eda.ncsu.edu/wiki/NCSU_EDA_Wiki. Accessed July 5, 2007.
- C. T. Nordhausen, E. M. Maynard, and R. A. Normann. Single unit recording capabilities of a 100 microelectrode array. *Brain Res.*, 726(1–2):129–140, Dec. 1996.
- J. L. Novak and B. C. Wheeler. Multisite hippocampal slice recording and stimulation using a 32 element microelectrode array. *J. Neurosci. Meth.*, 23(2):239–247, Mar. 1988.
- H. Nyquist. Thermal agitation of electric charge in conductors. *Phys. Rev.*, 32(1):110–113, July 1928.
- D. T. O’Keeffe, G. M. Lyons, A. E. Donnelly, and C. A. Byrne. Stimulus artifact removal using a software-based two-stage peak detection algorithm. *J. Neurosci. Meth.*, 109(2):137–145, Aug. 2001.

- R. H. Olsson, III, D. L. Buhl, A. M. Sirota, G. Buzsaki, and K. D. Wise. Band-tunable and multiplexed integrated circuits for simultaneous recording and stimulation with microelectrode arrays. *IEEE Trans. Biomed. Eng.*, 52(7):1303–1311, July 2005.
- K. G. Oweiss. A systems approach for data compression and latency reduction in cortically controlled brain machine interfaces. *IEEE Trans. Biomed. Eng.*, 53(7):1364–1377, July 2006.
- J. J. Pancrazio, P. P. Bey, Jr, A. Loloee, S. Manne, H.-C. Chao, L. L. Howard, W. M. Gosney, D. A. Borkholder, G. T. A. Kovacs, P. Manos, D. S. Cuttino, and D. A. Stenger. Description and demonstration of a CMOS amplifier-based-system with measurement and stimulation capability for bioelectrical signal transduction. *Biosens. Bioelectron.*, 13(9): 971–979, Oct. 1998.
- J. Pine. Studying mammalian neurons *in vitro* with multielectrode arrays. In *Proc. of the IEEE Engineering in Medicine and Biology Conference*, pages 3686–3689, Cancun, Mexico, 2003.
- J. Pine. Recording action potentials from cultured neurons with extracellular microcircuit electrodes. *J. Neurosci. Meth.*, 2(1):19–31, Feb. 1980.
- Plexon. Plexon neurotechnology research systems. URL <http://www.plexoninc.com/>. Accessed July 5, 2007.
- S. M. Potter and T. B. DeMarse. A new approach to neural cell culture for long-term studies. *J. Neurosci. Meth.*, 110(1–2):17–24, Sept. 2001.
- J. E. B. Randels. Kinetics of rapid electrode reactions. *Discuss. Faraday Soc.*, 1:11–19, 1947.
- J. D. Ross. Modeling the electrode/electrolyte interface: Discerning the cause of the stimulation artifact. Project Report, 2003.
- J. D. Ross, S. O’Connor, R. A. Blum, E. A. Brown, and S. P. DeWeerth. Multielectrode impedance tuning: reducing noise and improving stimulation efficacy. In *Proc. of the IEEE Engineering in Medicine and Biology Conference*, pages 4115–4117, San Francisco, CA, Sept. 2004.
- L. Rowe, M. Almasri, N. Fogleman, A. B. Frazier, and G. J. Brewer. An active microsccaffold for culturing 3-D neuronal networks. In *Transducers ’05*, volume 1, pages 948–951, Seoul, South Korea, 2005.
- M. E. Ruaro, P. Bonifazi, and V. Torre. Toward the neurocomputer: Image processing and pattern recognition with neuronal cultures. *IEEE Trans. Biomed. Eng.*, 52(3):371–383, Mar. 2005.
- M. Sahin and Y. Tie. Non-rectangular waveforms for neural stimulation with practical electrodes. *J. Neural Eng.*, 4(3):227–233, Sept. 2007.
- H. A. Spang, III and P. M. Schultheiss. Reduction of quantizing noise by use of feedback. *IRE Trans. Commun. Syst.*, 10(4):373–380, Dec. 1962.

- Spectre. Virtuoso Spectre circuit simulator. URL http://www.cadence.com/products/custom_ic/spectre/index.aspx. Accessed July 5, 2007.
- M. S. Steyaert, W. M. Sansen, and C. Zhongyuan. A micropower low-noise monolithic instrumentation amplifier for medical purposes. *IEEE J. Solid-State Circuits*, 22(6):1163–1168, Dec. 1987.
- C. A. Thomas, Jr, P. A. Springer, G. E. Loeb, Y. Berwald-Netter, and L. M. Okun. A miniature microelectrode array to monitor the bioelectric activity of cultured cells. *Exptl. Cell Res.*, 74(1):61–66, 1972.
- J. van Pelt, P. S. Wolters, M. A. Corner, W. L. C. Rutten, and G. J. A. Ramakers. Long-term characterization of firing dynamics of spontaneous bursts in cultured neural networks. *IEEE Trans. Biomed. Eng.*, 51(11):2051–2062, Nov. 2004.
- D. A. Wagenaar and S. M. Potter. A versatile all-channel stimulator for electrode arrays, with real-time control. *J. Neural Eng.*, 1(1):39–45, Mar. 2004.
- D. A. Wagenaar and S. M. Potter. Real-time multi-channel stimulus artifact suppression by local curve fitting. *J. Neurosci. Meth.*, 120(2):17–24, Oct. 2002.
- D. A. Wagenaar, J. Pine, and S. M. Potter. Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *J. Neurosci. Meth.*, 138:27–37, Sept. 2004.
- E. Warburg. Ueber das Verhalten sogenannter unpolarisierbarer Elektroden gegen Wechselstrom. *Annalen der Physik*, 303(3):493–499, 1899.
- B. Widrow. A study of rough amplitude quantization by means of Nyquist sampling theory. *IRE Trans. Circuit Theory*, 3(4):266–276, Dec. 1956.
- K. D. Wise, J. B. Angell, and A. Starr. An integrated-circuit approach to extracellular microelectrodes. *IEEE Trans. Biomed. Eng.*, 17(3):238–247, 1970.

VITA

Call me Richard. Some years ago—never mind how long precisely—having little or no money in my purse, and nothing particular to interest me in the job market, I thought I would study about a little and see the academic part of the world. It is a way I have of driving off the spleen, and regulating the circulation. Whenever I find myself growing grim about the mouth; whenever it is a damp, drizzly November in my soul; whenever I find myself involuntarily pausing before coffin warehouses, and bringing up the rear of every funeral I meet; and especially whenever my hypos get such an upper hand of me, that it requires a strong moral principle to prevent me from deliberately stepping into the street, and methodically knocking people's hats off—then, I account it high time to get to a lab bench as soon as I can. This is my substitute for pistol and ball. With a philosophical flourish Cato throws himself upon his sword; I quietly take to the oscilloscope. There is nothing surprising in this. If they but knew it, almost all men in their degree, some time or other, cherish very nearly the same feelings towards circuits with me (Melville, 1851).