

Long-Term, Subdermal Implantable EEG Recording and Seizure Detection

by

Bruno G. Do Valle

M.S. (ECE), Boston University (2008)

B.S. (ECE), Boston University (2006)

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

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2016

© Massachusetts Institute of Technology 2016. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
December 17, 2015

Certified by.....
Charles G. Sodini
Lebel Professor of Electrical Engineering
Department of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by
Professor Leslie A. Kolodziejcki
Chairman, Committee on Graduate Students
Department of Electrical Engineering and Computer Science

Long-Term, Subdermal Implantable EEG Recording and Seizure Detection

by

Bruno G. Do Valle

Submitted to the
Department of Electrical Engineering and Computer Science
on December 17, 2015, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Epilepsy is a common chronic neurological disorder that affects about 1% of the world population. Although electroencephalogram (EEG) has been the chief modality in the diagnosis and treatment of epileptic disorders for more than half a century, long-term recordings (more than a few days) can only be obtained in hospital settings. Many patients, however, have intermittent seizures occurring far less frequent. Patients cannot come into the hospital for weeks on end in order for a seizure to be captured on EEG — a necessary prerequisite for making a definitive diagnosis, tailoring therapy, or even establishing the true rate of seizures.

This work aims to address this need by proposing a subdermal implantable, eight-channel EEG recorder and seizure detector that has two modes of operation: diagnosis and seizure counting. In the diagnosis mode, EEG is continuously recorded with a 500 Hz bandwidth. This bandwidth is five times higher than typically done to record high-frequency oscillations, which, according to recently published research, may aid in the diagnosis of epilepsy. In the seizure counting mode, the system uses a novel low-power algorithm to track the number of seizures a patient has, providing doctors with a reliable count to help determine medication efficacy. This mode is especially important given the severity of antiepileptic drugs' side-effects.

An ASIC that implements the EEG recording and seizure detection algorithm was designed and fabricated in a 0.18 μm CMOS process. The ASIC includes eight EEG channels and was designed to minimize the system's power and size. The result is a power-efficient analog front end that requires 2.75 μW per channel in diagnosis mode and 0.84 μW per channel in seizure counting mode. Both modes have an input-referred noise of approximately 1.1 μVrms . The seizure detection algorithm has a sensitivity of 98.5%, a false alarm rate of 4.4 per hour, and a detection delay of 9.1 seconds. It consumes only 0.45 μW , which is over an order of magnitude less power than comparable algorithms.

Thesis Supervisor: Charles G. Sodini
Title: Lebel Professor of Electrical Engineering
Department of Electrical Engineering and Computer Science

Acknowledgments

I am very lucky to have met so many great people who had a significant impact in my life both academically and personally during my PhD at MIT. First, I would like to thank my advisor, Professor Charlie Sodini. He has pushed me to become a better engineer and I have learned a tremendous amount from him throughout these years. Professor Sodini also instilled in me an interest for baseball – we spent memorable afternoons at Fenway with him. I was very lucky to have him as an advisor.

This project wouldn't have been possible without Dr. Syd Cash's help and guidance. He has been so involved in all steps of the project that I consider him as a co-advisor. I am very thankful for his willingness to teach me so much of what I know today about EEG and epilepsy. I am also grateful for Professor Chandrakasan's willingness to serve on my committee and for all his help during my PhD years at MIT.

Next, I would like to thank all the student with whom I worked. They made my experience in the lab much more fruitful and fun. Sabino Pietrangelo, Maggie Delano, Joohyun Seo, Grant Anderson, and Kohei Onizuka, it has been a pleasure working with you guys. A special thanks goes to Daniel Kumar and Eric Winokur for all the design reviews throughout my project and during the IC design, and also for the coffee/hot chocolate breaks. Your support has been truly invaluable.

Several other people deserve to be acknowledged. Jason Yang was a great UROP and worked really hard on the external device. Marcus Yip provided the ADC that I used on my IC. Jason Naftulin and Lauren McClain showed me around MGH and helped me recruit patients for the project. Tom O'Dwyer generously offered to package my IC. Brian Brandt helped me during the IC design, especially with the anti-alias filter. I also made many friends outside the lab that had a big impact during my career, especially Lorenzo Turicchia, Scott Arfin, and Woradorn Wattanapanich.

Finally, I need to thank my family for supporting and guiding me throughout my life. I wouldn't have made this far if it wasn't for them. Last but not least, a special thanks goes to the most important person in my life: my wife Isabella. She is the best

thing that happened to me and she has been my biggest supporter. Thank you for being such a great person and understanding that a PhD takes more than a couple of years. This work is dedicated to you.

Contents

1	Introduction	23
1.1	Epilepsy	23
1.2	Diagnosing and Treating	24
1.3	Current Ambulatory Systems	25
1.4	Minimally Invasive System	27
1.5	Aims of Thesis Work	29
1.6	Thesis Organization	29
2	Electroencephalogram (EEG)	31
2.1	Introduction to EEG	31
2.2	10-20 Electrode Placement System	32
2.3	Typical EEG Plots	35
2.4	Artifacts	35
2.5	Seizures	38
2.6	Summary	38
3	Wearable EEG Recorder	41
3.1	First Prototype	42
3.1.1	System Testing	44
3.2	Second Prototype	46
3.2.1	System Testing	50
3.3	Summary	52

4	Seizure Detection Algorithm	53
4.1	Patient-Specific vs. Non-Specific Algorithms	53
4.2	Proposed Algorithm	56
4.2.1	Algorithm Steps	57
4.2.2	Training	59
4.2.3	Determining the Algorithm Parameters	62
4.3	Performance Measurement	71
4.3.1	Effect of Artifacts on Algorithm	73
4.3.2	Effect of Number of Channels on Algorithm	75
4.3.3	Effect of Number of Training Seizure Records	76
4.4	Summary	78
5	Low-Power, Eight-Channel EEG Recorder and Seizure Detector ASIC	79
5.1	Modes of Operation	79
5.2	Analog Front End Circuit Design	80
5.2.1	Low-Noise Amplifier	82
5.2.2	Programmable Gain Amplifier	88
5.2.3	Anti-Alias Filter	91
5.2.4	ADC	97
5.3	Digital Core	98
5.3.1	Data Decimation	98
5.3.2	Registers for AFE	99
5.4	Seizure Detection	99
5.5	Peripheral Circuits	103
5.5.1	Clock Generation	103
5.5.2	Current Reference	104
5.5.3	Voltage Reference	105
5.6	Physical Layout Considerations	105
5.7	Measurement Results	107
5.7.1	Frequency Response	109

5.7.2	Noise Measurement	111
5.7.3	Power Measurement	113
5.7.4	AFE Comparison	113
5.7.5	Seizure Detection	115
5.8	Summary	116
6	Subdermal Implantable System	119
6.1	System Overview	119
6.2	Inductive Link	119
6.2.1	Power Transfer	120
6.2.2	Programming the ASIC	122
6.2.3	Transmitting Data to External Device	125
6.2.4	Frequency of Operation	126
6.3	Implant	128
6.3.1	Power Management	128
6.3.2	Battery Charger	129
6.3.3	PCB	131
6.4	External Device	132
6.4.1	PCB	132
6.4.2	Power Consumption	132
6.5	External Device Housing	134
6.6	Animal Testing	136
6.7	Summary	145
7	Conclusion	147
7.1	Summary of Contributions	147
7.2	Future Work	148

List of Figures

1-1	Patient with scalp electrodes to record EEG	25
1-2	Two examples of ambulatory EEG (left) and (right).	26
1-3	Example of an invasive ambulatory system	27
1-4	Subdermal EEG recorder and seizure detector implant	28
2-1	Potential difference across the neuron's membrane	32
2-2	Measuring distance from nasion to inion and dividing the skull into areas according to the 10-20 system. Each letter represents a different area	33
2-3	Measuring distance between preauricular points and dividing the skull into areas according to the 10-20 system. Each letter represents a different area	34
2-4	Electrode label according to its location on the scalp	34
2-5	Example of ten-second EEG plot	35
2-6	Ten-second EEG segment showing eye blinks in the Fp channels . . .	36
2-7	Ten-second EEG segment showing muscle artifact in the T4 channel .	36
2-8	Ten-second EEG segment showing electrode artifact in FP1-F7 and F7-T3 channels	37
2-9	Ten-second EEG segment showing ECG artifact indicated by horizontal lines in F7-T3 channel	37
2-10	Ten-second EEG segment showing spike-waves during an absence seizure	38
2-11	Ten-second EEG segment showing a tonic-clonic seizure	39

3-1	Simplified system block diagram.	43
3-2	Front of the first prototype PCB.	43
3-3	Back of the first prototype PCB.	43
3-4	Picture of the 3D package created for the prototype.	44
3-5	Ten-second segment measured from a patient at MGH. Top graph is from prototype and bottom graph is from MGH's EEG recording machine.	45
3-6	Silver/silver chloride electrode used in the first prototype.	45
3-7	Better option for silver/silver chloride electrode.	46
3-8	Pushing electrodes to see how much artifact is produced. Top graph shows artifacts using old electrode. Bottom graph shows artifacts with new electrode.	47
3-9	Block diagram of the second prototype.	48
3-10	Ten-second segment of EEG shown in the iPod Touch app.	48
3-11	Package used to house prototype.	49
3-12	Device and electrodes placement shown in a test subject.	50
3-13	Front of the second prototype PCB.	50
3-14	Ten-second segment of normal EEG recorded from a patient at MGH.	51
3-15	Ten-second segment of EEG recorded during a seizure from one of the patients using our device (top) and hospital recorder (bottom).	51
4-1	Ten-second segment of EEG data showing a seizure from Patient A. Seizure starts at 3369 seconds. The scale on the bottom right corner indicates the height of 300 μV	54
4-2	Ten-second segment of EEG data showing a seizure from Patient B. Seizure starts at 2996 seconds. The scale on the bottom right corner indicates the height of 300 μV	54
4-3	Another seizure from Patient B. Seizure starts at 1732 seconds. The scale on the bottom right corner indicates the height of 300 μV	55

4-4	Breaking down the seizure detection into two blocks: simple algorithm in IC and complex one in computer.	56
4-5	Top-level block diagram of proposed seizure detection algorithm.	58
4-6	Tonic-clonic seizure with spikes marked by red lines	58
4-7	Results of simple probability example. Each letter represents the energy measured during a two-second segment.	60
4-8	Probability of having a seizure versus time. Seizure starts at 1467 seconds and ends at 1494 seconds.	61
4-9	Percentage of seizures detected for different number of energy bins.	64
4-10	Average false alarm rate for different number of energy bins.	64
4-11	Average detection delay for different number of energy bins.	65
4-12	Percentage of seizures detected for different number of LUT rows.	66
4-13	Average false alarm rate for different number of LUT rows.	66
4-14	Average detection delay for different number of LUT rows.	67
4-15	Probability value output from algorithm without moving average. Red line represents threshold value.	68
4-16	Probability value output from algorithm with moving average. Red line represents threshold value.	68
4-17	Percentage of seizures detected for different moving average window lengths.	69
4-18	Average false alarm rate for different moving average window lengths.	70
4-19	Average detection delay for different moving average window lengths.	70
4-20	Percentage of seizures detected versus threshold value.	71
4-21	False alarms rate versus threshold value.	72
4-22	Detection delay versus threshold value.	72
4-23	False alarm rate for each patient in the PhysioNet dataset.	73
4-24	Ten-second segment of artifact-free EEG from patient 13. The scale on the bottom right corner indicates the height of 100 μ V.	74
4-25	Ten-second segment of patient 13 EEG data. The scale on the bottom right corner indicates the height of 150 μ V.	74

4-26	Effect of the number of channels on percentage of seizure detection and false alarm rate.	75
4-27	Effect of the number of training seizure records on percentage of seizure detected.	76
4-28	Effect of the number of training seizure records on average false alarm rate.	77
4-29	Effect of the number of training seizure records on seizure detection delay.	77
5-1	Analog front end block diagram.	82
5-2	Top-level schematic of LNA.	83
5-3	PMOS transistor with parasitic diodes that affect pseudo-resistor. . .	84
5-4	Measured resistance of the pseudo-resistor as the difference between input and output voltages are varied	84
5-5	LNA schematic.	86
5-6	CMFB schematic.	88
5-7	PGA top-level schematic.	90
5-8	PGA schematic.	91
5-9	Schematic of a second-order Gm-C filter	93
5-10	Schematic of a second-order sinc filter	93
5-11	Schematic of Rauch filter.	94
5-12	Implementing a resistor using switched capacitor.	95
5-13	Rauch filter schematic using switched capacitors to replace resistors. Clk1 and Clk2 are non-overlapping clocks.	96
5-14	Magnitude response of the decimation filter in Mode 1. The difference between the quantized and reference is less than 1 dB.	99
5-15	Schematic of current gain block.	100
5-16	Block diagram of the seizure detection implementation.	100
5-17	Magnitude response of the reference (red) and quantized (blue) Beta filter (13 Hz - 30 Hz).	101

5-18	Magnitude response of the reference (red) and quantized (blue) Alpha filter (8 Hz - 13 Hz).	101
5-19	Magnitude response of the reference (red) and quantized (blue) Theta filter (4 Hz - 8 Hz).	101
5-20	Magnitude response of the reference (red) and quantized (blue) Delta filter (0.5 Hz - 4 Hz).	102
5-21	Schematic of clock circuit	104
5-22	Current reference schematic.	105
5-23	Voltage reference schematic.	106
5-24	Top-level LNA schematic with mismatched parasitics at input terminals of the amplifier.	106
5-25	Die micrograph of the ASIC with circuit blocks labeled.	108
5-26	LNA frequency response in Mode 1 (blue) and Mode 2 (red).	109
5-27	PGA frequency response in Mode 1 (blue) and Mode 2 (red).	109
5-28	Anti-alias filter frequency response in Mode 1 (blue) and Mode 2 (red). Less than 1 dB difference from simulation.	110
5-29	AFE (from LNA to Filter) frequency response in Mode 1 (blue) and Mode 2 (red).	110
5-30	Output noise power spectral density in Modes 1 (blue) and 2 (red) measured with Agilent 35670A Dynamic Signal Analyzer.	111
5-31	Frequency response measured at the output of the PGA with its gain set to maximum in Modes 1 (blue) and 2 (red).	112
5-32	Input-referred noise power spectral density in Modes 1 (blue) and 2 (red).	112
5-33	Pie chart indicating the power breakdown of each block in Mode 2.	114
5-34	Comparison of the LNA with published results in literature.	114
6-1	Block diagram of the system.	120
6-2	Schematic of a class-E power amplifier.	121
6-3	Schematic of circuit used to rectify the power from the coil.	121

6-4	Example of a typical OOK.	123
6-5	Proposed data encoding using OOK.	124
6-6	Schematic of circuit used to receive data from external device.	125
6-7	Example of bits 1 and 0 being transmitted to program the ASIC. Top and bottom graphs are the clock that controls the class-E power amplifier and the output from the Schmitt trigger, respectively.	125
6-8	Schematic of circuit used to transmit data to external device.	126
6-9	Top and bottom graphs are the voltage at the external coil and the output of the envelope detector, respectively	127
6-10	Block diagram of the implant.	128
6-11	Two options considered for the power management.	130
6-12	3 mAh implantable medical battery from Quallion used in the implant.	130
6-13	Front and back of the implant PCB.	131
6-14	Block diagram of the external device.	132
6-15	Front and back of the external PCB.	133
6-16	External device housing with the lid opened.	134
6-17	External device housing with the lid closed.	135
6-18	Location of the external device and coil on the swine. The distance between the external housing and coil is approximately 10 cm	135
6-19	Pre-surgical planning.	136
6-20	Device being implanted in swine.	137
6-21	Device is implanted in swine.	137
6-22	Implant coated with epoxy.	138
6-23	Device implanted in swine and external device glued to ear tag.	139
6-24	Ten-second EEG segment recorded from the swine.	140
6-25	FFT of the ten-second EEG segment recorded from the swine. First peak at 3 Hz and second one at 8 Hz	140
6-26	External coil with holes for suture.	141
6-27	Implant being placed on the swine.	142
6-28	Coil sutured to the swine's head.	142

6-29	Swine wearing the jacket that holds external device.	143
6-30	Ten-second EEG segment recorded from swine. Data is plotted in Sequential montage.	144
6-31	FFT of the channel 'ch7-ch8'. Main peak is located at 1.6 Hz.	144
7-1	Circuit to bypass battery charger.	150

List of Tables

1.1	Types of generalized seizures and corresponding symptoms	23
1.2	Types of partial seizures and corresponding symptoms	24
2.1	EEG frequency bands	32
3.1	American Clinical Neurophysiology Society guideline for long-term epilepsy monitoring	41
4.1	Results from of whether seizure waveforms change over time	55
4.2	EEG frequency bands	57
4.3	Probability results from simple example	60
5.1	Differences between Modes 1 and 2	81
5.2	American Clinical Neurophysiology Society guideline for long-term epilepsy monitoring	81
5.3	Width and length of transistors used in the LNA (0.18 μm process) .	87
5.4	Simulated results of the LNA in Modes 1 and 2	89
5.5	Simulated results of the LNA in Modes 1 and 2 compared to the American Clinical Neurophysiology Society guideline	89
5.6	Possible values of C_1 and the corresponding PGA gain	90
5.7	Simulated results of the PGA in Modes 1 and 2	92
5.8	Capacitor values used in the anti-alias filter	96
5.9	Simulated results of the anti-alias amplifier in Modes 1 and 2	97
5.10	Estimated power consumption of individual blocks in seizure detection	103
5.11	Measured versus simulated power of main blocks in the ASIC	113

5.12 Our AFE compared with published literature	115
5.13 Seizure detection results compared with published literature	116
6.1 Power consumption of main blocks in the external device	133

Abbreviations and Acronyms

ACNS	American Clinical Neurophysiology Society
ADC	Analog-to-digital converter
AFE	Analog front end
ASIC	Application specific integrated circuit
BW	Bandwidth
CMFB	Common-mode feedback
CMRR	Common-mode rejection ratio
ECG	Electrocardiogram
EEG	Electroencephalogram
ENOB	Effective number of bits
EOV	Electrode offset voltage
FDA	Food and Drug Administration
FIR	Finite impulse response
HFO	High-frequency oscillation
IA	Instrumentation amplifier
IC	Integrated circuit
IIR	Infinite impulse response
LNA	Low-noise amplifier
LPF	Lowpass filter
LSB	Least significant bit
LUT	Lookup table
MGH	Massachusetts General Hospital
MSB	Most significant bit

NEF	Noise efficiency factor
OOK	On-off keying
PCB	Printed circuit board
PGA	Programmable gain amplifier
PSD	Power spectral density
RMS	Root mean squared
SAR	Successive approximation register
SNDR	Signal to noise distortion ratio
SPI	Serial peripheral interface
UART	Universal asynchronous receiver/transmitter
USB	Universal serial bus

Chapter 1

Introduction

1.1 Epilepsy

Epilepsy is a common neurological disorder that affects about 50 million people worldwide [1]. It is characterized by repeated seizures, which are caused by abnormal neuronal firing in the affected brain area [2]. If the abnormal neuronal activity is localized in a specific part of the brain, it is called a focal seizure; seizures that involve large portions of the brain at their outset are named generalized. Symptoms vary depending on the part of the brain affected, for example, a focal seizure on the auditory cortex will probably cause the patient to experience hearing sensations. Tables 1.1 and 1.2 list the types of generalized and focal seizures and their corresponding symptoms [3].

Table 1.1: Types of generalized seizures and corresponding symptoms

Generalized Seizures	Symptoms
"Grand Mal" or Generalized tonic-clonic	Unconsciousness, convulsions, muscle rigidity
Absence	Brief loss of consciousness
Myoclonic	Sporadic (isolated), jerking movements
Clonic	Repetitive, jerking movements
Tonic	Muscle stiffness, rigidity
Atonic	Loss of muscle tone

Table 1.2: Types of partial seizures and corresponding symptoms

Partial Seizures	Symptoms
Simple (awareness is retained)	Jerking, muscle rigidity, spasms, head-turning. Unusual sensations affecting either vision, hearing, smell, taste, or touch. Memory or emotional disturbances.
Complex (impairment of awareness)	Automatisms such as lip smacking, chewing, fidgeting, walking and other repetitive, involuntary but coordinated movements. Sporadic (isolated), jerking movements.
Partial seizure with secondary generalization	Symptoms that are initially associated with a preservation of consciousness that then evolves into a loss of consciousness and convulsions.

1.2 Diagnosing and Treating

In order to give a definite diagnosis of epilepsy, a neurologist must have the patient's electroencephalogram (EEG), which is the recording of the electrical activity along the scalp, during a seizure attack [2]. To obtain this recording, the patient usually stays at the hospital for a couple of weeks with electrodes attached to his or her scalp as shown in Figure 1-1. Limited outpatient EEG can be obtained for no more than 2-3 days. This becomes challenging because many patients have intermittent seizures occurring very infrequently — from once a week to once every few months [4] — and cannot come into the hospital for weeks in order for an event to be captured on EEG, making an accurate and definitive diagnosis very difficult.

Once the diagnosis is made, neurologists face another major challenge: to obtain the true rate of seizures of the patient to be able to determine if the prescribed medication is working correctly. Currently, doctors rely on patients to track the number of seizures they have even though it has been shown that patient reports tend to be very inaccurate [6]. However, because neurologists have no other means of getting that information, they have to rely on an inaccurate number when determining whether the dose should be increased/decreased or if the medication should be switched.



Figure 1-1: Patient with scalp electrodes to record EEG [5].

1.3 Current Ambulatory Systems

Currently, EEG systems can record the neuronal activity in the following ways: on the scalp (EEG), under the skull but on top of the brain (electrocorticography, ECOG), or in the brain with implanted electrodes (single-unit recordings). Figure 1-2 shows two examples of ambulatory EEG systems that record from the scalp. This type of system has many drawbacks: they only last up to 3 days, the patient must wear a backpack that contains the recording electronics all the time, the patient cannot shower since water interferes with the electrodes [7], and many avoid daily activities such as work and school while wearing the system given its aesthetics [8].

Some work has been done in trying to reduce the form factor of ambulatory systems by reducing the size of the electronics [9] and [11]. Although several researches have created different versions of non-invasive EEG monitors, all of them share the same issues [12]:

- 24-7 monitoring cannot be obtained because the patient must remove the device for certain activities (e.g. shower)
- Cannot ensure the electrodes remain attached for the duration of the recording
- Electrodes require wired connections to a board, which might limit patient



Figure 1-2: Two examples of ambulatory EEG [9] (left) and [10] (right).

movements

- Most systems are very visible so patients do not want to wear the device in public
- Long-term recordings generate large amounts of data, which requires significant amount of power to transmit and/or record

Wearable systems must place the electrodes on the scalp, so it is very hard, if not impossible, to design a multi-channel device that solves all of the problems listed above. One way that has been attempted is through invasive systems.

The invasive solution of placing the electrodes on the brain or implanting them also present challenges. The major one is the requirement of a craniotomy — surgical operation in which a hole is drilled through the skull in order to access the brain. Since this type of surgery is risky and invasive, implantable systems are only used in very specific cases, for example, an attempt to determine the exact focal point of a seizure before surgical procedure to remove that part of the brain. One example of an invasive solution is Neuropace's RNS system, illustrated in Figure 1-3. The idea is to detect a seizure as soon as possible from the recorded brain activity and stimulate the brain in order to stop the seizure. The electronics package is implanted in the skull, the recording electrodes are placed on top of the brain, and the stimulating electrodes

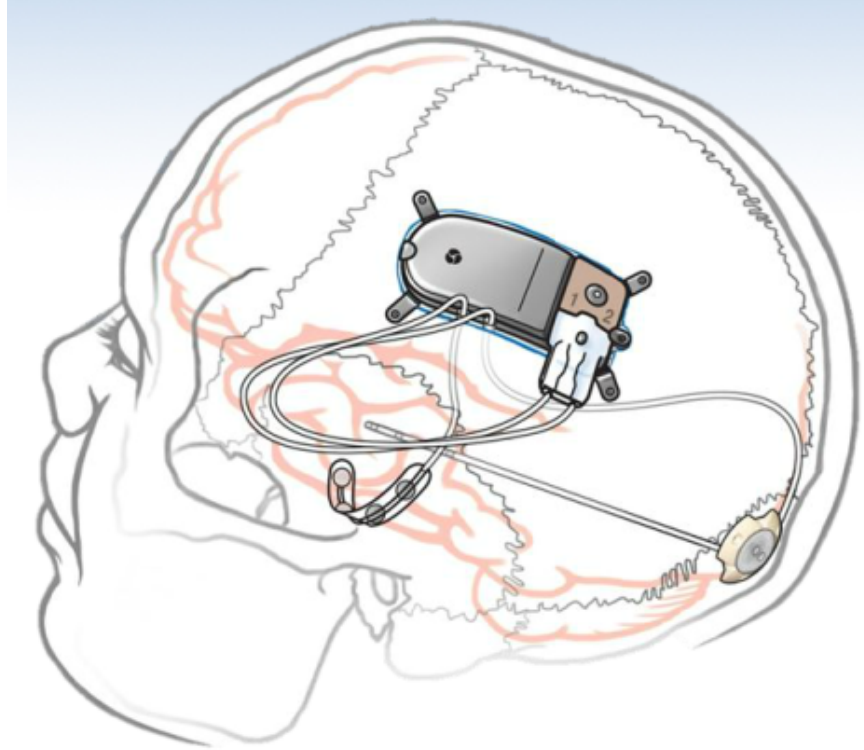


Figure 1-3: Example of an invasive ambulatory system [14].

are implanted in the brain. The system is very invasive since a part of the skull has to be removed and some electrodes are implanted into the brain. According to [13], Neuropace's system can cause severe complications, such as infection, intracranial hemorrhage, and paralysis.

Since non-invasive systems cannot guarantee continuous monitoring or patient compliance and invasive systems require major surgery, we propose a different approach: a minimally invasive system where the device and recording electrodes are implanted subdermally, between the scalp and skull.

1.4 Minimally Invasive System

This work presents an eight-channel, subdermal implantable EEG recorder and seizure detector. The system is implanted behind the patient's right or left ear and the electrodes are tunneled underneath the scalp into the desired locations, as shown in Figure 1-4.




Figure 1-4: Subdermal EEG recorder and seizure detector implant [15].

This approach has several advantages over current EEG systems:

- Does not require a craniotomy; the surgical procedure to place the device behind the ear is minimal and very similar to a cochlear implant
- 24-7 monitoring is obtained
- Long-term monitoring is possible with daily battery recharge
- System is small and hidden underneath the skin, so patient compliance is not an issue
- Does not limit patient movements

Our system has two modes of operation: diagnosis and seizure counting. In diagnosis, the EEG is continuously recorded until at least one seizure is recorded. In seizure counting, the system uses a low-power algorithm to track the number of seizures a patient has. Our system is not intended to stop seizures, as a result, detection delay is not a major concern. On the other hand, the invasive systems presented in [16]

and [17] need to detect seizures as fast as possible since they are trying to stop them, thus detection delay is a major concern. This difference enables us to significantly reduce the power in our detection algorithm.

1.5 Aims of Thesis Work

The research presented in this thesis has the following five aims:

- Design a wearable EEG recorder for long-term monitoring of patients with epilepsy, test it on patients in the hospital, and compare the acquired data with the hospital's recording.
- Develop a power-efficient seizure detection algorithm to be implemented in an application-specific integrated circuit (ASIC).
- Design a low-power, eight-channel ASIC to detect seizures and record EEG up to five times the typical bandwidth, possibly leading to improved diagnosis.
- Create a subdermal implantable system using the ASIC. Power and data are transmitted to/from the ASIC through an inductive link to an external device.
- Perform animal testing using a swine model in a hospital.

1.6 Thesis Organization

This thesis is organized into the following chapters:

- Chapter 2 explains how EEG is originated and describes the characteristics of a typical recording on the scalp, including how the presence of artifacts can impact recordings. We then present examples of seizure recordings and how it differs from a normal one.
- Chapter 3 describes the design of a one-channel ambulatory system. Our first prototype stores all the data into a flash memory. The second prototype transmits all the data to an iPod Touch and to the cloud, enabling real time access to

the data. The device is tested on patients with epilepsy at MGH and recorded data is compared to the hospital's recordings.

- Chapter 4 presents a novel power-efficient seizure detection algorithm to be implemented in an ambulatory system. The algorithm uses the energy measured in specific frequency bands to detect a seizure. The chapter concludes with a comparison between our algorithm and current state of the art.
- Chapter 5 describes the design of a low-power, eight-channel EEG recorder and seizure detector ASIC. The goal of the ASIC is to minimize power and enable the design of a smaller device when compared to other systems using off-the-shelf components. In this chapter, the low-power circuit design of each block and the electrical characterization are also discussed.
- Chapter 6 explains the design of a subdermal implantable EEG recorder and seizure detector containing the ASIC presented in Chapter 5. The external device, which is responsible for data communication and power transfer through an inductive link, is also described. The chapter ends with the results of an animal test on a swine model at MGH.
- Chapter 7 draws the thesis conclusions and presents future research directions.

Chapter 2

Electroencephalogram (EEG)

2.1 Introduction to EEG

Electroencephalogram (EEG) is the recording on the scalp of electrical activity produced by neurons. Each EEG channel requires two electrodes: one placed close to the area of interest on the brain and the other (reference electrode) on a location of almost no neuronal activity (e.g. mastoid).

In resting state, there is a potential difference of -65 mV between inside of a neuron and the fluid that surrounds it [18]. This potential difference, illustrated in Figure 2-1, is caused by an uneven distribution of ions across the neuron's membrane. When a neuron becomes depolarized, sodium and calcium ions (positive ions) enter the cell, leaving the extracellular fluid around this neuron more negative. On the other hand, during inhibitory activity, positive ions (primarily potassium) leave the neuron, causing the extracellular fluid to be relatively more positive. As a result, we are able to detect a voltage difference between the electrode close to this neuron and the reference electrode. In reality EEG is the recording of activity between thousands to millions of neurons and not a single one [2].

EEG signals recorded at the scalp are usually on the order of 10 μV to 100 μV and can reach about 1 mV during a seizure. A detailed analysis of EEG is typically done with a frequency range from 0.1 Hz to 70 Hz [2]; however, recent research has shown a link between oscillations around 250 Hz and seizure onset area [19]. Historically, EEG

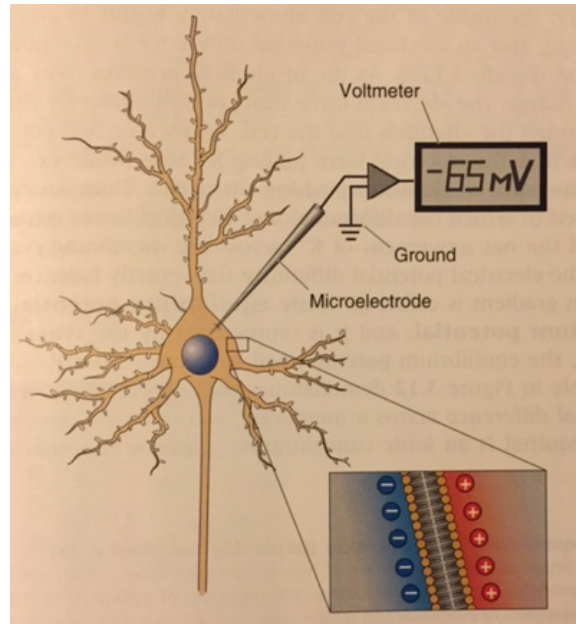


Figure 2-1: Potential difference across the neuron's membrane [18].

Table 2.1: EEG frequency bands

Band Name	Frequency Range
Delta	0.5 - 4 Hz
Theta	4 - 8 Hz
Alpha	8 - 14 Hz
Beta	14 - 30 Hz

has been divided into four different frequency bands, shown in Table 2.1. Waveforms in these frequency bands are found under different conditions and also at different ages. Delta waves are usually found in deep sleep while Theta and Beta in light sleep. In adults the background activity usually falls in the Alpha band while in children in the Theta band. The background activity is an indication of the brain's excitability. Its frequency increases with age and decreases during sleep [20].

2.2 10-20 Electrode Placement System

The 10-20 system was created to standardize the placement of electrodes used in EEG to ensure reproducibility. The distance between the nasion (intersection of the frontal bone and two nasal bones) and inion (most prominent projection of the occipital bone)




Figure 2-2: Measuring distance from nasion to inion and dividing the skull into areas according to the 10-20 system. Each letter represents a different area [21].

is measured and the head is divided into five areas. The first area is located at 10% of the measured distance and the remaining four at 20% intervals, as indicated by Figure 2-2 [21]. Then, the distance between left preauricular point through the C vertex mark to the right preauricular point is measured according to Figure 2-3 and the head is divided again into five areas [22]. The result of this mapping is shown in Figure 2-4, where each letter represents a different area. The letters Fp, F, C, T and O represent frontal polar, frontal, central, temporal, and occipital areas, respectively. The number after the letter represents if the electrode is placed on the left or right hemisphere. Even numbers are for left hemisphere and odd for right. The letter 'z' indicates placement in the midline.

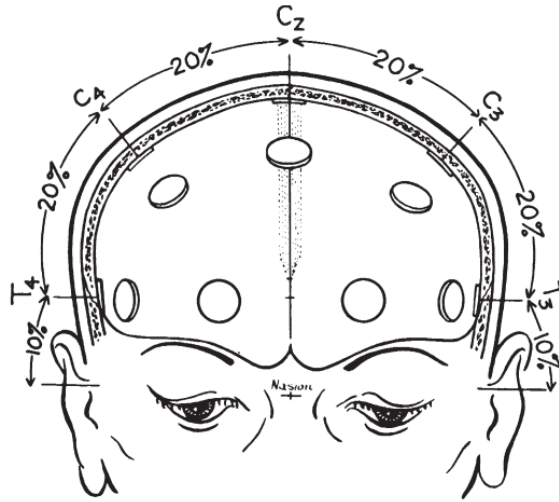


Figure 2-3: Measuring distance between preauricular points and dividing the skull into areas according to the 10-20 system. Each letter represents a different area [21].




Figure 2-4: Electrode label according to its location on the scalp [22].




Figure 2-5: Example of ten-second EEG plot [17].

2.3 Typical EEG Plots

A typical EEG plot is shown in Figure 2-5, where each line represents the difference between two electrodes on the scalp. For example, the first line plots the difference between the FP1 and F7 electrodes. The y-axis represents voltage and its scale is usually indicated in the bottom right corner of the figure; in Figure 2-5 the height of the scale represents $150 \mu\text{V}$. The x-axis represents time in seconds with a total time span of ten seconds. This work will follow this traditional way of plotting EEG.

2.4 Artifacts

EEG amplitudes are very small — between $10 \mu\text{V}$ to 1mV — so the recordings are very sensitive to artifacts. The most common types of artifacts are eye blinks, muscle artifact, electrode movement, and electrocardiogram (ECG) contamination. Figure 2-6 shows a ten-second EEG recording with eye blinks; they are the big voltage drops in the Fp channels. The frontal area is the one most susceptible to this type of artifact due to its close proximity to the eyes. Figure 2-7 shows an example of muscle artifact in an EEG recording. It consists of high-frequency, high-amplitude noise in the T4 electrode. The frequency content of muscle artifact can vary between 20 Hz to 300 Hz, which overlaps with EEG frequency range [23]. Another type of artifact is electrode movement that usually resembles Figure 2-8. Electrode movement causes




Figure 2-6: Ten-second EEG segment showing eye blinks in the Fp channels [2].

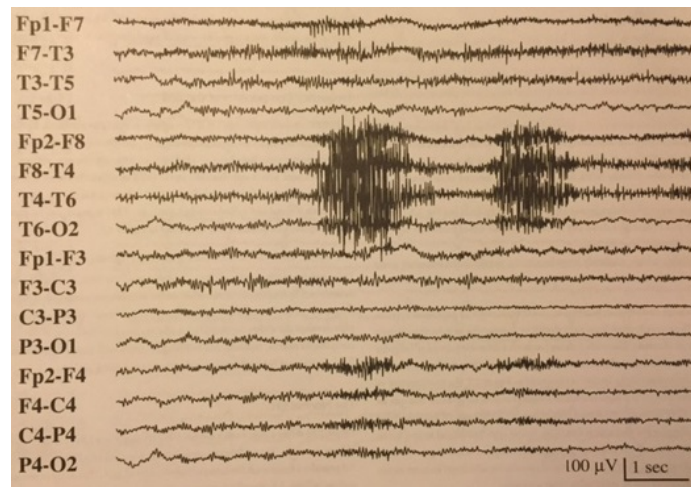


Figure 2-7: Ten-second EEG segment showing muscle artifact in the T4 channel [2].

large voltage swings that can saturate the channel. ECG signals can also show up in EEG, as illustrated by Figure 2-9, where the ECG is indicated by horizontal lines in F7-T3 channel. All the artifacts described in this section are common and appear with varying degree. Electrode movement and muscle artifact appear even more frequently in ambulatory systems since the patient is free to move around. Therefore, it is important to try minimize their impact in the recording system so that the actual EEG data can still be analyzed.




Figure 2-8: Ten-second EEG segment showing electrode artifact in FP1-F7 and F7-T3 channels [2].




Figure 2-9: Ten-second EEG segment showing ECG artifact indicated by horizontal lines in F7-T3 channel [2].




Figure 2-10: Ten-second EEG segment showing spike-waves during an absence seizure [2].

2.5 Seizures

Seizures are caused by an abnormal neuronal firing rate at the affected brain area, usually resulting in typical patterns on the EEG [2]. For example, absence (petit mal) seizures typically have spike-waves happening three to four times a second, as shown in Figure 2-10. Tonic-clonic usually has spikes and then a lot of muscle artifact, as shown in Figure 2-11. There are several different types of seizures and each one tends to have its own pattern. Doctors look for these patterns when making a diagnosis [20].

2.6 Summary

During the neuronal activity, ions are exchanged between the neuron and extracellular fluid, causing potential voltages to appear. EEG records these potential voltages through electrodes that are placed on the scalp. Since the amplitudes are typically less than 1 mV, EEG recordings are very sensitive to artifacts. During a seizure, neurons start firing synchronously, producing patterns on the EEG. It is by recognizing these patterns that doctors are able to diagnose patients.




Figure 2-11: Ten-second EEG segment showing a tonic-clonic seizure [2].

Chapter 3

Wearable EEG Recorder

Our first step towards creating an implantable EEG monitor was to understand the challenges in recording EEG. To achieve that we created prototypes using off-the-shelf components. For a system to be clinically useful in the diagnosis of epilepsy, the channel count must be at least four [24]. Since the main objective of our wearable devices is not to diagnose epilepsy, but instead to track the number of seizures a patient has, our prototypes did not need more than one channel.

Our prototypes follow the American Clinical Neurophysiology Society (ACNS) guideline of recommended specifications for a system recording long term EEG for epilepsy, which is listed in Table 3.1 [25]. We will now go into detail on the design of each one of our prototypes.

Table 3.1: American Clinical Neurophysiology Society guideline for long-term epilepsy monitoring

Parameter	Recommended Value
Low-Frequency Response	0.5 Hz or lower
High-Frequency Response	70 Hz or higher
Noise Level	$< 1 \mu\text{V}_{\text{rms}}$
Input Impedance	$> 1 \text{ M}\Omega$
Common Mode Rejection	$> 60 \text{ dB}$
Dynamic Range	$> 40 \text{ dB}$

3.1 First Prototype

One of our main objectives when designing our first prototype was to minimize the device's form factor. This was accomplished by reducing battery size. Usually the highest power consumption block of a system is transmitting the data wirelessly or storing it to memory. Since the amount of data is proportional to sample rate, to minimize power consumption, we decided to reduce the sample rate while still exceeding ACNS's recommendations listed in Table 3.1. As a result, the EEG was sampled at 256 Hz with a low and high-frequency response of 0.5 Hz and 100 Hz, respectively.

When an electrode is placed on the skin, a potential voltage, commonly referred as electrode offset voltage (EOV), is generated. This voltage is dependent on the electrode material, electrolyte composition and temperature [26]. EOV can be as high as 60 mV for a silver/silver chloride electrode and its frequency content is typically at DC [27]. Since EEG amplitudes can range from a few tenths of microvolts up to around a millivolt during a seizure, our system must reject the EOV while providing the necessary gain to the EEG signal [28].

The low-noise amplifier chosen for this prototype was the INA333 from TI because it was one of the lowest power and noise amplifiers available at the time. It uses an auto-calibration technique to remove flicker noise. The INA333 input-referred voltage noise is 50 nV/sqrt(Hz), which yields a total input noise of 800 nV_{rms} for a 256 Hz bandwidth. The INA333 also has an input impedance greater than 1 G Ω and a CMRR of 100 dB. Thus, all specifications are met according to Table 3.1.

Since the EOV can be as high as 60 mV, the gain of the INA333 was set to 20 in order to tolerate offset voltages as high as 100 mV for a 2 V supply. A bandpass filter was included after the INA333 to remove any component below 0.5 Hz and above 100 Hz, thus also removing the EOV before going to the second amplifier. Figure 3-1 shows a simplified block diagram of the system. EEG amplitudes are typically less than 1 mV, so the second amplifier provides the required gain such that a 1 mV input EEG corresponds to the maximum supply voltage.



Figure 3-1: Simplified system block diagram.




Figure 3-2: Front of the first prototype PCB.

The data is sampled at 2,048 Hz to simplify the design of the anti-alias filter and digitized using a 12 bit ADC. A lower sampling rate could have been used, but the anti-alias filter order would have to be increased in order to provide the same attenuation. A MSP430 microcontroller decimates the data from 2,048 Hz to 256 Hz and stores it into an 8 GB microSD card. Figures 3-2 and 3-3 show pictures of the front and back of the PCB, respectively. The front of the PCB contains all the chips while the back is devoted to the battery. The PCB size is 3.3 cm by 3.7 cm.

The total power consumption of the device is approximately 1.6 mW. About 51% of the power dissipation is used by the microcontroller. Storing the data into a flash (SD card) and the analog front end take up about 34% and 15%, respectively. The system runs from a lithium-ion coin cell battery CR2430, which has a capacity of 280 mAh,




Figure 3-3: Back of the first prototype PCB.



Figure 3-4: Picture of the 3D package created for the prototype.

and lasts approximately 21 days. One way to reduce the system's power consumption is by using polyphase filters in the microcontroller to decimate the data. This way the computation requirements is minimized and a power savings of approximately 20% can be achieved [29].

The package to house the system was designed in SolidWorks and created using a 3D printer. The package, shown in Figure 3-4, measures 7.5 cm by 4.3 cm and sits on top of the patient's ear with two electrodes attached. The package was made symmetrically so that it can be used either on the left or right ear. The circular shape is semi flexible to better conform to the ear and help secure the device.

3.1.1 System Testing

This device was tested on three patients with epilepsy at Massachusetts General Hospital (MGH). One electrode was placed at the patient's temporal region and the other at the mastoid. Figure 3-5 shows a ten-second segment recorded with this device (top graph) compared to the hospital's EEG recording (bottom plot). The two plots are not identical because the electrodes were placed at slightly different locations and the electrodes used in our first prototype did not adhere properly to the skin, producing artifacts. Figure 3-6 shows the front and back of the silver/silver chloride electrodes used, which measure 3.3 cm x 2.2 cm.




Figure 3-5: Ten-second segment measured from a patient at MGH. Top graph is from prototype and bottom graph is from MGH's EEG recording machine.




Figure 3-6: Silver/silver chloride electrode used in the first prototype.




Figure 3-7: Better option for silver/silver chloride electrode.

Testing with prototype one allowed us to identify two main issues: first, the data contained a lot of artifacts due mainly to the poor quality of the electrodes used; second, because all the data was stored to a memory card and not transmitted in real time, we were not able to uncover and address problems that were happening during recordings, for example, electrodes falling off. Fixing those two issues was our main objective when designing the second prototype.

3.2 Second Prototype

Before starting the design of our second prototype, we did a series of tests to identify an electrode that would work well for our purposes. The one that looked the most promising was the silver/silver chloride Huggable electrode from MVAP [30]. It measures one inch in diameter and is shown in Figure 3-7. The new electrode adheres a lot better to the skin and produces less artifacts, as can be seen in Figure 3-8, where we tested the old and new electrodes by slightly pushing them at the center.

With the new prototype, it was crucial to get EEG data in real time to allow us to fix any problems as they were happening. Similar to the first prototype, in the new version, the EEG signal is sampled at 2,048 Hz and digitized using a 12 bit ADC. A MSP430 microcontroller decimates the data from 2,048 Hz down to 256 Hz, but now transfers it to a CC2540 Bluetooth chipset that transmits the data through Low Energy Bluetooth to an iPod Touch. A custom app, which works with iPod Touch or iPhone, was developed and operates continuously receiving data from the CC2540 and uploading it in real time through WiFi to a server located at the hospital, as illustrated in Figure 3-9. The data is transferred using SSH (Secure Shell) and no

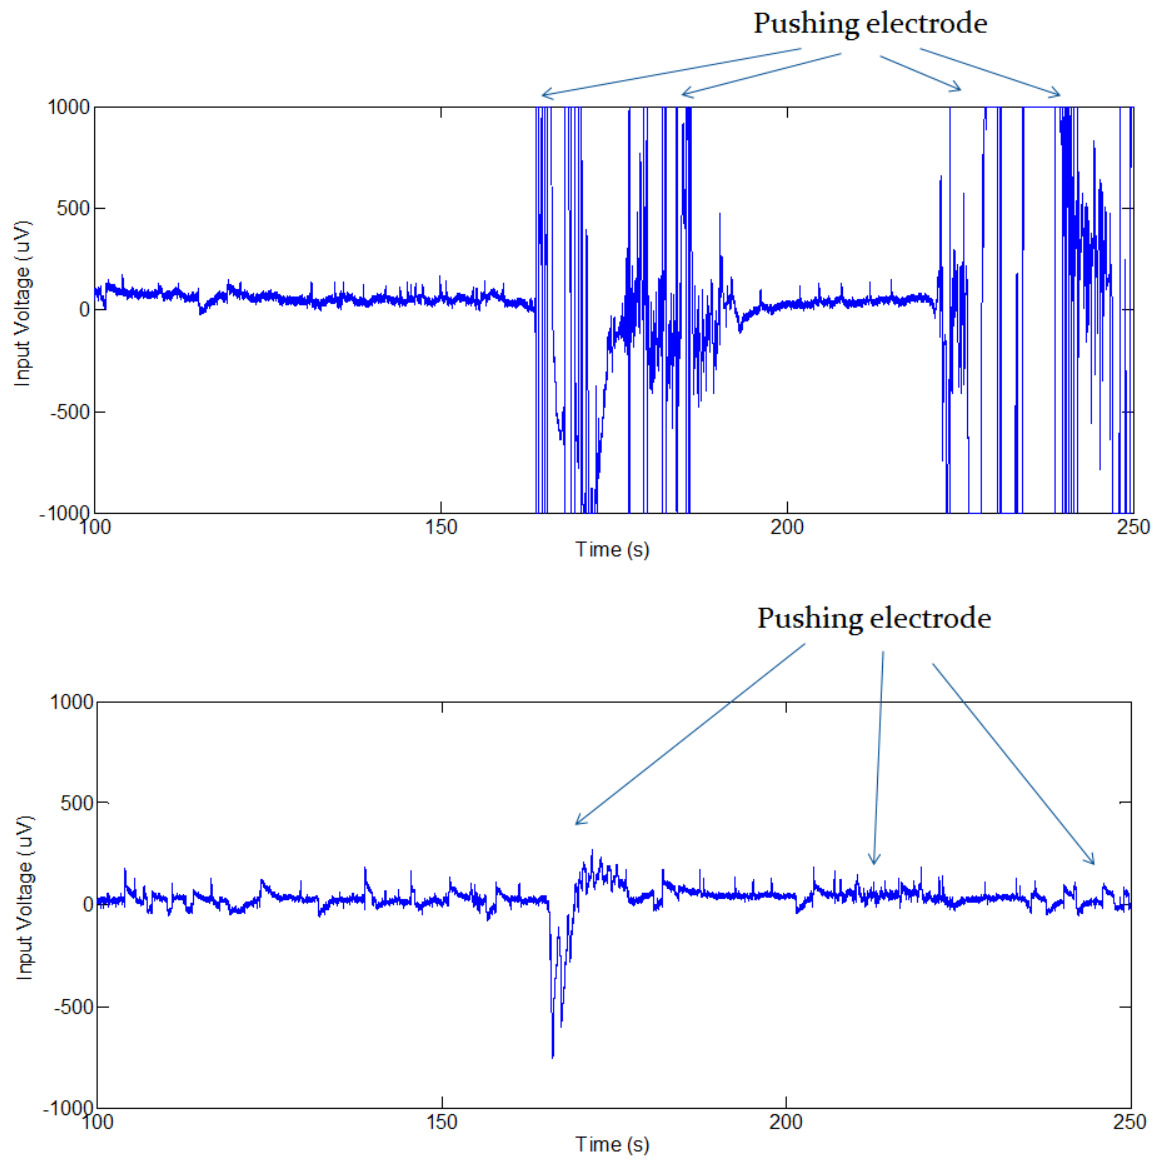


Figure 3-8: Pushing electrodes to see how much artifact is produced. Top graph shows artifacts using old electrode. Bottom graph shows artifacts with new electrode.




Figure 3-9: Block diagram of the second prototype.




Figure 3-10: Ten-second segment of EEG shown in the iPod Touch app.

patient identifier is included to comply with HIPPA (Health Insurance Portability and Accountability Act) policy. The app also plots the EEG in real time, as shown in Figure 3-10.

The MVAP Huggable electrodes have a wet gel in the center that dries in approximately 24 hours. When dry, its impedance increases significantly and degrades the recording. Since ultimately patients will be sent home with this device, they will have to replace the electrodes every day. However, one of the system’s advantages is that, when replacing them, the patient can use the app to see if the electrodes are replaced correctly. If placed incorrectly (e.g. not properly attached to the skin), the app will not display any data, indicating placement needs to be fixed.

Batteries typically consume a significant portion of a system’s area, so power consumption was minimized in order to reduce battery size, which in turn reduced the device’s area. One way to reduce power consumption is to reduce the sample rate; another is to minimize the supply voltage. In this prototype, the supply voltage was limited to 2 V by the Low Energy Bluetooth chipset (CC2540), which has a minimum recommended operating voltage of 2 V. The total power consumption of the device is 4.4 mW: approximately 77% of the power is used to transmit the data through Low Energy Bluetooth, the microcontroller consumes 18%, and the remaining 5% is



Figure 3-11: Package used to house prototype.

used by the analog front end (amplifiers, filter, and ADC). The system runs from a zinc-air battery A675P, which has a capacity of 550 mAh, and lasts for approximately 6.8 days. A zinc-air battery was chosen because it has the highest energy density due to the fact that the cathode reactant is air and not included inside the battery [31]. Therefore, for the same capacity, a zinc-air battery has a smaller form factor than other battery types. The end result is a smaller battery and a smaller device.

The package to house the system was again designed in SolidWorks and created using a 3D printer. The package, illustrated in Figure 3-11, measures 6.0 cm by 6.4 cm and sits behind the patient's ear as shown in Figure 3-12. The package can be worn on the left or right ear and is semi flexible to conform better and help secure the device. The PCB, seen in Figure 3-13, was also made in a circular shape.

The prototype requires two electrodes in order to measure EEG. One electrode is placed at the temporal lobe or the area of interest and the other is placed at the mastoid (reference electrode). An actual ground electrode is not needed because the package itself serves that purpose. The bottom part of the package was coated with silver conductive epoxy (MG Chemicals 8331) so that when it makes physical contact with the ear, it sets the body to our voltage reference.




Figure 3-12: Device and electrodes placement shown in a test subject.




Figure 3-13: Front of the second prototype PCB.

3.2.1 System Testing

The device was tested on two patients at MGH and over 100 hours of EEG were recorded. Figure 3-14 presents a ten-second EEG segment from a patient with electrodes placed at T4 (right temporal) and right mastoid. This EEG has normal background rhythms with some ECG artifacts. In contrast, Figure 3-15 shows EEG from the same patient taken during a seizure. The top and bottom graphs represent the data from our device and the hospital's EEG recording machine, respectively. Both sets clearly show high-amplitude ($100 \mu\text{V}$) 7 Hz waves that are an indication of a seizure. The distinctions between our data and the hospital's are explained by differences in the positioning of the reference electrode: ours was placed at the right mastoid while the hospital's at A2 (right ear).




Figure 3-14: Ten-second segment of normal EEG recorded from a patient at MGH.




Figure 3-15: Ten-second segment of EEG recorded during a seizure from one of the patients using our device (top) and hospital recorder (bottom).

3.3 Summary

This chapter presented a behind-the-ear EEG recording device that uses an iPhone or iPod Touch to continuously upload the patient's data to a hospital server, giving doctors real time access to the data. This device is one step forward in the design of smaller and better ambulatory EEG systems. The electrodes are easily replaceable, enabling the patient to remove and reapply the device for daily activities, such as taking a shower. This new system also offers patients suffering from infrequent seizures a more convenient way of monitoring their seizures, and gives doctors more accurate information for the diagnosis and treatment options.

Chapter 4

Seizure Detection Algorithm

In this chapter we present a novel low-power seizure detection algorithm that was designed to be implemented in an ASIC. The algorithm is patient-specific, which means at least one previously annotated seizure record is required during training phase.

4.1 Patient-Specific vs. Non-Specific Algorithms

One of the main difficulties in building a very accurate patient non-specific algorithm — one that can work for any person — is that seizure waveforms vary greatly from one patient to another. This point is illustrated by the EEG segments shown in Figures 4-1 and 4-2. The seizure in Figure 4-1 consists of approximately 300 μV amplitude waves at 6 Hz in the Parietal lobe while the one in Figure 4-2 starts as 50 μV waves at 9 Hz in the Temporal lobe and changes into 300 μV waves at 3 Hz. However, despite the fact that seizures from two different people might differ significantly, a second seizure from the same patient often resemble the first one, as illustrated by Figures 4-2 and 4-3. Patient-specific algorithms based on this premise can have high accuracy by training on historical records from a patient before it is employed to detect seizures in real time. During training phase, the algorithm extracts specific information that is later used to classify the segment. Consequently, patient-specific algorithms tend to yield better detection results than non-specific ones [17].




Figure 4-1: Ten-second segment of EEG data showing a seizure from Patient A. Seizure starts at 3369 seconds. The scale on the bottom right corner indicates the height of 300 μV .




Figure 4-2: Ten-second segment of EEG data showing a seizure from Patient B. Seizure starts at 2996 seconds. The scale on the bottom right corner indicates the height of 300 μV .

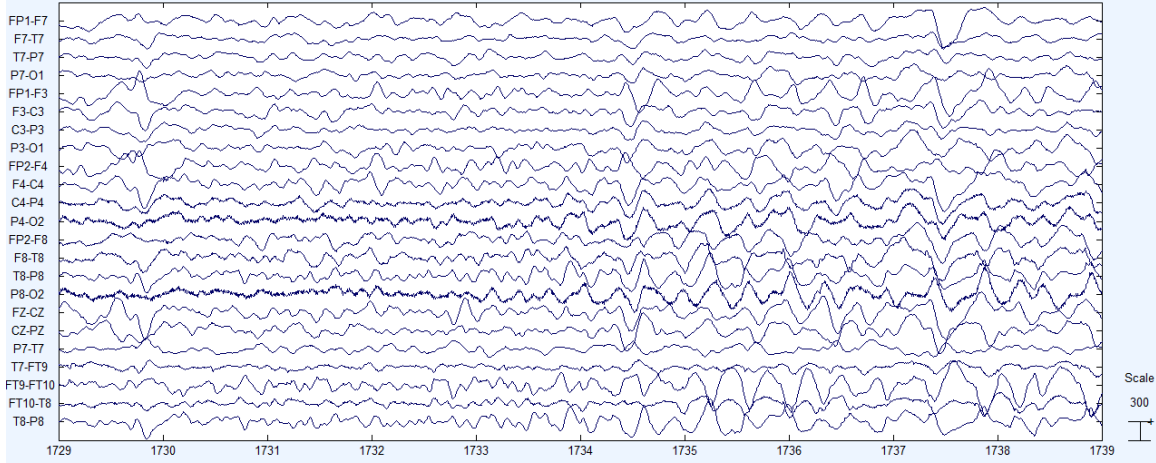


Figure 4-3: Another seizure from Patient B. Seizure starts at 1732 seconds. The scale on the bottom right corner indicates the height of $300 \mu\text{V}$.

Table 4.1: Results from [32] of whether seizure waveforms change over time
Results from research of whether seizure waveforms change over time

EEG Changes	Percentage of Patients
No Changes (EEG waveforms looked similar)	80%
Unilateral -> Bilateral	6%
Bilateral -> Unilateral or no seizures	14%

Although the sample recording in Figures 4-2 and 4-3 were only made a couple of days apart, research shows that seizures would resemble each other even after years. In one research, 35 adults with epilepsy were monitored during ten years in order to prove this fact [32]. The results listed in Table 4.1 show that in 80% of the cases, the seizure waveforms remained similar after ten years. In 6% of the cases seizures that initially only happened in one hemisphere of the brain started to happen on both. The remaining 14% of the patients either stopped having seizures or went from having them in both hemispheres to just one. This study proves that patient-specific algorithms should still work even years after the training phase is performed. It is important to point out that patients with progressive brain disorders, such as brain tumors, were excluded from this study since it is known that the seizure waveforms in these patients change over time.

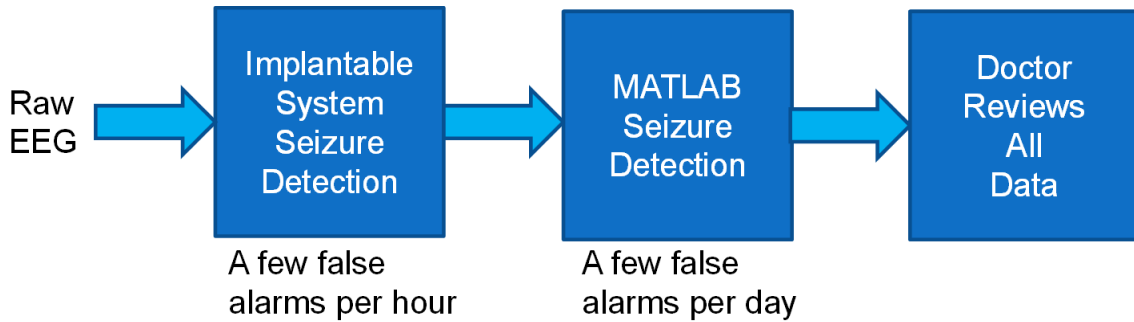


Figure 4-4: Breaking down the seizure detection into two blocks: simple algorithm in IC and complex one in computer.

4.2 Proposed Algorithm

Most of the published detection algorithms try to minimize false alarms while still detecting the highest number of seizures by applying complex algorithms such as machine learning [33], [34], and [35]. The challenge with these algorithms is that with complexity comes higher computation requirements and power dissipation. For an ambulatory system, power efficiency is extremely important in order to reduce battery size, which in turn reduces system size.

We propose a method, illustrated by Figure 4-4, where instead of having one power-hungry algorithm that detects all seizures and produces very few false alarms, we have a power-efficient, simple algorithm running in the ambulatory system. The proposed algorithm is intended to count the number of seizures and not to detect its onset; therefore, latency is not a major concern. The simple algorithm produces a higher number of false alarms than the complex systems. However, all the collected data can be stored and later analyzed in a computer using any complex algorithm, which will help reduce the count of false alarms, making the dataset manageable for a doctor to review. This thesis describes only the design of the simple algorithm since the complex one can be any computer-based algorithm, such as the machine learning example published in [33].

Table 4.2: EEG frequency bands

Band Name	Frequency Range
Delta	0.5 - 4 Hz
Theta	4 - 8 Hz
Alpha	8 - 13 Hz
Beta	13 - 30 Hz

4.2.1 Algorithm Steps

The key feature of the proposed algorithm is that it uses energy in physiological-meaningful frequency bands to generate a simple probability value of the patient having a seizure. Table 4.2 lists the frequency bands of interest used in the algorithm. Figure 4-5 shows a top-level block diagram illustration of the proposed four-step algorithm.

The first step in the algorithm is to break down the EEG into time-domain segments. There were specific reasons why we selected a two-second segment. The lowest frequency of interest is 0.5 Hz, so the time-domain segments must be at least two seconds in length. During a seizure, the EEG's spectral content varies over time; for example, spikes sometimes appear in the beginning of a seizure but disappear after a few seconds, as illustrated in Figure 4-6. Using long EEG segments might smear some spectral signatures of a seizure, which could hurt sensitivity. Also, using a long segment would increase detection delay because the time needed to guarantee that the segment only has seizure data equals the segment length; for example, with a ten-second segment we would need at least ten seconds of EEG from the start of the seizure. For these reasons, we have decided to break down the EEG segments into two-second segments. Each segment overlaps the next by 50% or one second. So, segment one consists of the first two seconds of EEG data, segment two consists of seconds two and three, and so on. The 50% overlap was included because it improved specificity by about 18% with only a 7% power penalty.

In the second step, the algorithm measures the energy of each two-second segment in the four frequency bands listed in Table 4.2. The energy is calculated by summing

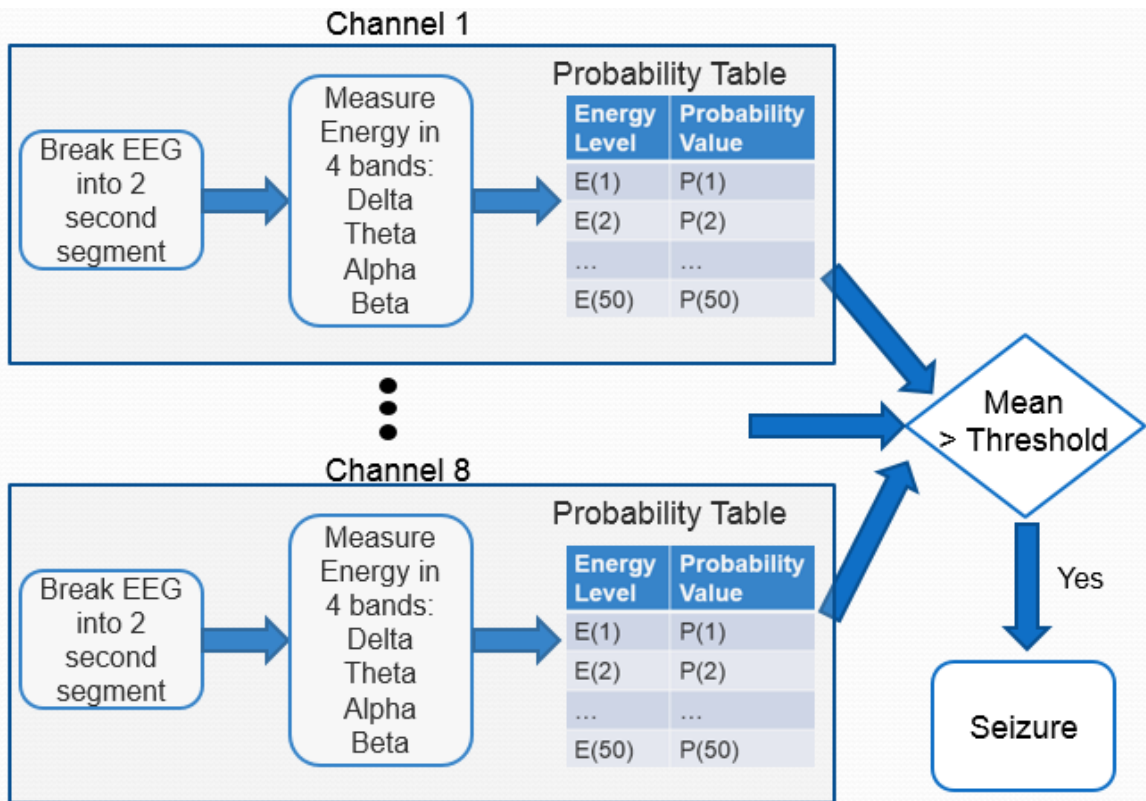


Figure 4-5: Top-level block diagram of proposed seizure detection algorithm.




Figure 4-6: Tonic-clonic seizure with spikes marked by red lines [2].

the square of the voltage values (see Equation 4.1, derived from Parseval’s theorem). In the third step, we use the measured energies in a lookup table (generated during patient-specific training) to get a probability value of a seizure taking place. Each channel has its own probability table so that channels closer to the seizure area have higher weight. This happens because a channel far away from the seizure area should have a probability value very close to zero, since the measured energy should not differ significantly between seizure and normal state. On the other hand, the energy measured from a channel very close to the seizure area should differ significantly between seizure and normal state, resulting in a probability value close to one.

In the last step, we take the mean of the probability output of all channels. The mean value goes through a ten-second moving average filter and is compared with a threshold determined during training. The moving average was added to reduce the number of false alarms. More details on how the length of the moving average affects the algorithm can be found in Section 4.2.3. If the moving average output is higher than the threshold, the algorithm flags a seizure and records the data for ten seconds for later analysis.

$$Energy = \sum |x[n]|^2 \tag{4.1}$$

4.2.2 Training

During the training phase, we gather data to generate a probability lookup table and determine the threshold value. At least one EEG record from the patient, with the seizure’s start and end time annotated, is required for this process.

Generating Probability Lookup Table

With the training data we follow the four-step procedure described in the previous section. To calculate the probability value of having a seizure given a measured energy value ‘ $Energy_n$ ’ we use Equation 4.2, dividing how often the measured energy value ‘ $Energy_n$ ’ occurs during the seizure by how often it occurs in the entire record.




Figure 4-7: Results of simple probability example. Each letter represents the energy measured during a two-second segment.

Table 4.3: Probability results from simple example

Probability of having a seizure given an energy value	Probability Value
$P(\text{Seizure} \text{Blue})$	0.25
$P(\text{Seizure} \text{Green})$	0.33
$P(\text{Seizure} \text{Orange})$	0.5
$P(\text{Seizure} \text{Red})$	0

$$P(\text{Seizure}|\text{Energy}_n) = \frac{P(\text{Seizure and Energy}_n)}{P(\text{Energy}_n)} \tag{4.2}$$

The simplified example described next should help understand this calculation. For simplicity, let’s assume we have only one EEG channel and we are only measuring the energy at one frequency band. Let’s also assume we only measure the following four color-coded energy levels: blue (B), green (G), orange (O), and red (R). We are given an annotated seizure record and after following the procedure described in the previous section, we get the results shown in Figure 4-7, where each letter represents the energy measured during a two-second segment.

Since the record has been annotated before, we know when the seizure starts and ends. As a result, we can calculate the probability of having a seizure given a measured energy value. For example, in this case the color red doesn’t appear during a seizure, so if we are monitoring the patient and we measure the color red, we can infer that the probability of him having a seizure is zero. Table 4.3 lists the probability values for each energy level.




Figure 4-8: Probability of having a seizure versus time. Seizure starts at 1467 seconds and ends at 1494 seconds.

Determining Threshold Value

The threshold value is a percentage of the peak probability value during the seizure in the training record. It is determined by running the training record through our algorithm, which will yield a probability value of having a seizure every second, as illustrated in Figure 4-8. Since we used this record to train, it's no surprise that there is a spike in probability value during the seizure between 1467 and 1494 seconds.

The threshold value can be set at different levels depending on the accuracy needed. There is an important trade-off when setting it: the higher the value, the lower the number of false alarms and the lower the number of seizures detected. Section 4.3 will present the results of the proposed algorithm as we vary the threshold value.

4.2.3 Determining the Algorithm Parameters

The EEG database used to measure the performance of the algorithm is the CHB-MIT Scalp EEG Database. The data was collected from 22 patients (5 males, ages 3-22; and 17 females, ages 1.5-19) at Children’s Hospital Boston [17].

To estimate the performance of our algorithm on data recorded from a patient, we used the leave-one-record-out validation scheme. In this scheme, we use all seizure records, except for one, to train and test on the record that wasn’t used. The algorithm performance will be determined by the percentage of seizures detected, the average number of false alarms per hour, and the detection delay, which is how long the algorithm takes to detect a seizure after it started.

Sensitivity and specificity are usual statistical performance measures of a binary classification algorithm and are described in Equations 4.3 and 4.4. Instead of using sensitivity and specificity, we will use percentage of seizures detected and average number of false alarms per hour. Percentage of seizures detected and sensitivity are identical, so they can be used interchangeably. However, specificity and average number of false alarms are different. Specificity measures the percentage of times the algorithm correctly classified the segment as non-seizure. For example, a specificity of 99% means that the algorithm correctly classified 99% of the non-seizure segments. We have chosen to use false alarm per hour instead of specificity because the former provides an easier way to understand the total number of records generated by the system. For example, saying that the average false alarm rate reduces from once every thirty minutes to once every hour is clearer than saying specificity increases from 99.94% to 99.97%.

$$Sensitivity = \frac{\textit{number of true positives}}{\textit{number of positives}} \quad (4.3)$$

$$Specificity = \frac{\textit{number of true negatives}}{\textit{number of negatives}} \quad (4.4)$$

The main parameters that need to be determined for the algorithm are: size of the lookup table and the length of the moving average. The next sections will go over

each one of them.

Size of the Probability Lookup Table

The probability lookup table (LUT) contains a probability value of having a seizure for every single possible energy combination from the output of the four bandpass filters. Since this algorithm will be implemented in an ASIC, the size of the LUT had to be considered to minimize die area and also power consumption. There are two ways to reduce the size of the LUT: bin the energy value to decrease the number of possible combinations, or only store the probability values of a few combinations. Both methods were used in our algorithm.

To bin the energy, we need to determine the minimum and maximum values and the number of bins to use. Logarithmic binning is used because EEG amplitudes can vary by three orders of magnitude: from 1 μV , typically the noise floor of recording systems, to 1 mV during a seizure. The lowest bin was set to the equivalent energy generated by a 1 μV sine wave during two seconds. Similarly, the highest bin was set to a 1 mV sine wave energy.

The effects of changing the number of energy bins are illustrated in Figures 4-9, 4-10, and 4-11. According to these figures, our algorithm has a peak performance at eight bins. Increasing the number of bins decreases false alarm rate because it reduces the range of energies that the algorithm will flag as a seizure. At the same time, increasing the number of bins decreases percentage of seizures detected because if there is any slight change in the seizure energy from the training set, the algorithm won't detect it.

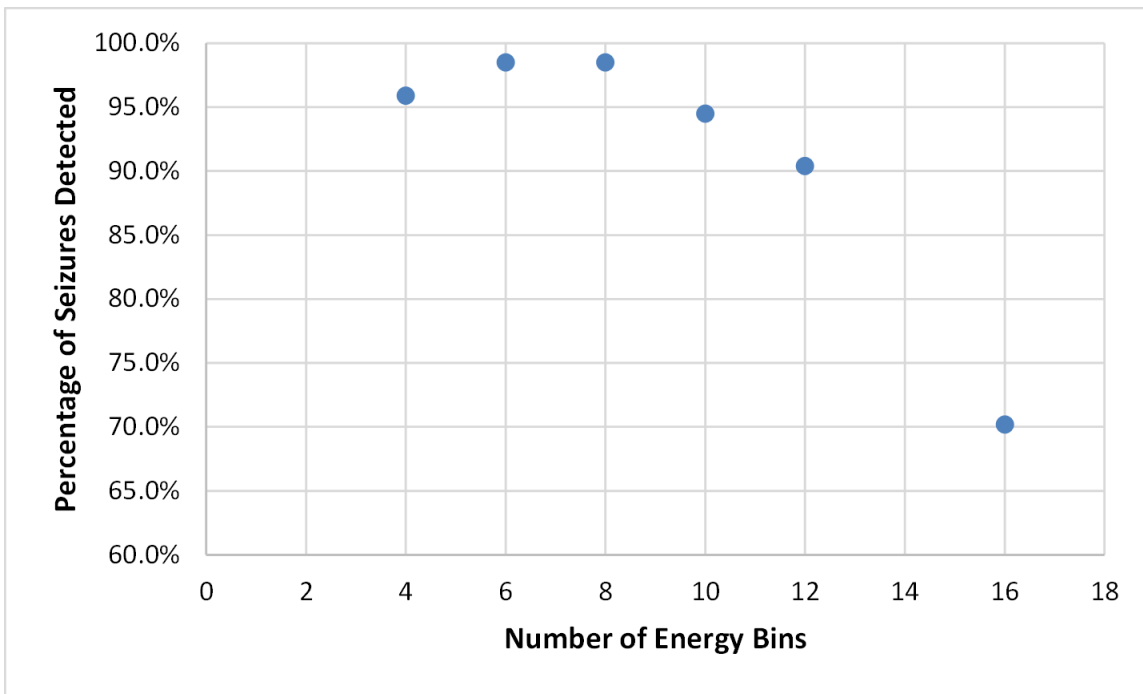


Figure 4-9: Percentage of seizures detected for different number of energy bins.




Figure 4-10: Average false alarm rate for different number of energy bins.




Figure 4-11: Average detection delay for different number of energy bins.

Each channel has its own probability LUT. In our case, we have four frequency bands and the energy is discretized into eight bins. As a result, in order to have one entry for every possible combination, we would need $8^4 = 4,096$ rows in the LUT for each channel. Most of the energy combinations are not recorded during a seizure, so there is no need to have one entry for every possible combination. Consequently, we limit the LUT to only fifty rows. After we measure the probability for every possible combination, we only keep the fifty highest probability values, all others are assumed to yield a value of zero. The effects of changing the number of LUT rows on the algorithm are shown in Figures 4-12, 4-13, and 4-14. The percentage of seizures detected reach the maximum value of 98.5% for fifty rows and stay constant for higher numbers. However, the average number of false alarms increase slightly beyond fifty. Therefore, we decided to limit the number of rows in the LUT to fifty.




Figure 4-12: Percentage of seizures detected for different number of LUT rows.




Figure 4-13: Average false alarm rate for different number of LUT rows.




Figure 4-14: Average detection delay for different number of LUT rows.

Length of Moving Average

The effect of the moving average can be seen in Figures 4-15 and 4-16. Without the moving average, several false alarms are detected in addition to the seizure that happens between 1982 seconds and 2029 seconds. By adding the moving average, the number of false alarms is reduced to only one while the actual seizure is still flagged.




Figure 4-15: Probability value output from algorithm without moving average. Red line represents threshold value.




Figure 4-16: Probability value output from algorithm with moving average. Red line represents threshold value.




Figure 4-17: Percentage of seizures detected for different moving average window lengths.

The effects of changing the moving average window length on the algorithm results are presented in Figures 4-17, 4-18, and 4-19. The effect on the percentage of seizures detected is very minimal; going from a four-second to a sixteen-second window only decreases the percentage of seizures detected by 0.5%. However, it has a major impact in the number of false alarms and detection delay. As we increase the length of the window, the number of false alarms starts to decrease until it plateaus around ten-second window length. On the other hand, the detection delay keeps increasing for higher lengths. As a result, we decided to go with a ten-second window, minimizing the number of false alarms without increasing detection delay too much.




Figure 4-18: Average false alarm rate for different moving average window lengths.




Figure 4-19: Average detection delay for different moving average window lengths.




Figure 4-20: Percentage of seizures detected versus threshold value.

4.3 Performance Measurement

The threshold controls the trade-off between false alarm rate and percentage of seizures detected. Depending on the experiment's accuracy needs, the threshold value can be changed to satisfy the requirements. Figures 4-20, 4-21, and 4-22 show respectively how the percentage of seizures detected, false alarm rate, and average detection delay vary depending on the chosen threshold. Since the threshold is stored in a register in the IC, it can even be changed during the monitoring if the number of false alarms for a specific patient is deemed too high.

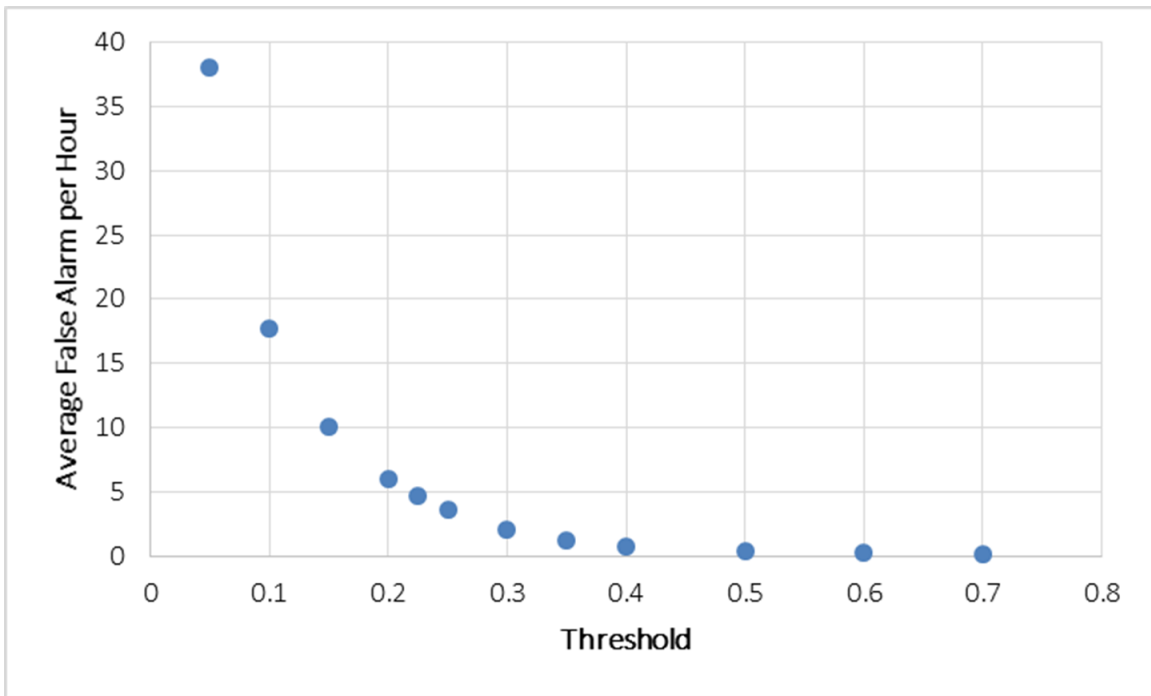


Figure 4-21: False alarms rate versus threshold value.




Figure 4-22: Detection delay versus threshold value.




Figure 4-23: False alarm rate for each patient in the PhysioNet dataset.

4.3.1 Effect of Artifacts on Algorithm

The presence of artifacts and noise in the recording can severely affect the algorithm results. Figure 4-23 shows the average false alarm rate for each one of the 24 patients in the PhysioNet dataset. According to the figure, patient 13 has over five times more false alarms than any other patient. The presence of artifacts in the training record for that patient explains this anomaly. As a comparison, Figures 4-24 and 4-25 show ten-second segments of an artifact-free EEG recording and patient 13 data, respectively. If the EEG recordings used for training contain artifacts with amplitudes significantly higher than the seizure waveforms, then the algorithm will give a high probability of having a seizure any time there are artifacts. This was the case for patient 13, explaining the high number of false alarms.




Figure 4-24: Ten-second segment of artifact-free EEG from patient 13. The scale on the bottom right corner indicates the height of $100 \mu\text{V}$.




Figure 4-25: Ten-second segment of patient 13 EEG data. The scale on the bottom right corner indicates the height of $150 \mu\text{V}$.




Figure 4-26: Effect of the number of channels on percentage of seizure detection and false alarm rate.

4.3.2 Effect of Number of Channels on Algorithm

Figure 4-26 shows the results of the algorithm for using different number of channels: from one to 23, as available in the PhysioNet dataset. As depicted by the chart, if only one channel is available, the algorithm can still detect approximately 94% of the seizures. This is very encouraging, given the fact that sometimes electrodes might be placed on areas of the brain barely affected by seizures. Our IC can record up to eight channels of EEG, and if the location of the seizure is known and all eight channels are correctly placed, 98.5% of seizures are detected.




Figure 4-27: Effect of the number of training seizure records on percentage of seizure detected.

4.3.3 Effect of Number of Training Seizure Records

The algorithm requires seizure records for training, so one question that arises is how many records are necessary and how they impact the performance of the algorithm. Figures 4-27, 4-28, and 4-29 help us evaluate the impact of the number of training records. By using only one seizure record during training, the algorithm can still detect 87.4% of seizures; however, the false alarm rate is very high, 20.79 per hour. By using two records, there is a huge improvement in false alarm rate (from 20.79 to 6.35 per hour) and some improvement in percentage of seizures detected (87.4% to 93.9%). After two records, the improvement of adding a third starts to diminish until it plateaus at about four. Therefore, to achieve optimal performance, four records are recommended. If four records are not available, then one option is to train on the number of records available and then use the recorded seizure data as new training sets.




Figure 4-28: Effect of the number of training seizure records on average false alarm rate.




Figure 4-29: Effect of the number of training seizure records on seizure detection delay.

4.4 Summary

This chapter presented a novel power-efficient seizure detection algorithm ideal for an ambulatory system where the goal is to count the number of seizures a patient has. With an accurate seizure count, doctors know precisely if the medication is working and whether the dose should be increased or decreased. This system has the potential to minimize antiepileptic medication side-effects by adjusting the dosage amount, thus improving patient's quality of life.

Chapter 5

Low-Power, Eight-Channel EEG Recorder and Seizure Detector ASIC

This chapter presents the design and measurement results of a low-power, eight-channel EEG recorder and seizure detector application-specific integrated circuit (ASIC), designed to be used in an implanted system. The ASIC includes four key components: an analog front end (AFE) to amplify and filter the EEG signals, analog-to-digital converters (ADC), a seizure detector, and a digital core responsible for decimation and data communication. The system has two modes of operation: high data rate used for diagnosis, and low data rate used for seizure detection. The specifics on the two modes of operation and each circuit block are discussed in detail in this chapter. Measured IC results are also presented.

5.1 Modes of Operation

Most EEGs recorded in the hospital are sampled at 256 Hz since the clinical analysis usually focuses on activities below 50 Hz [36]. However, one of the main drawbacks of this sample rate is the lack of high-frequency content (higher than 128 Hz). High-frequency oscillations (HFO), which lie between 100 Hz - 500 Hz and typically have amplitudes between $5 \mu\text{V}$ - $30 \mu\text{V}$, were shown to be present on some patients at the seizure onset region before the beginning of a seizure [37] and [38]. Since a sample

rate of only 256 Hz would not be able to record HFO, it was important that our system sampled the EEG at a frequency higher than 1,000 Hz.

Our system has two modes of operation, which will be referred to as Mode 1 and Mode 2. Mode 1 is used for diagnosis. In this mode the EEG is sampled at 1,280 Hz, enabling the recording of HFO, and all the data is transmitted continuously to an external device that is described in Chapter 6. Since the external device must be worn continuously by the patient in order to receive all the data, power can also be transmitted continuously and is not a major concern. Mode 1 is required in order to record a few seizures from the patient to create the necessary training data set for the seizure detection algorithm described in Chapter 4. After a few seizures are recorded, the system can be used in Mode 2: seizure detection mode. In this mode, the system only records ten seconds of EEG after a seizure is detected by our algorithm. Since data is not continuously transferred in this mode, the external device only has to be worn for approximately 15 minutes every day to recharge an implanted battery and to retrieve all the seizure data recorded during the day. The system must last at least one entire day between battery recharges, so low-power consumption is important. To minimize power consumption in Mode 2, the EEG is sampled at a typical clinical setting of 256 Hz and only recorded after our algorithm detects a seizure. The drawback of this approach is that the beginning of a seizure may be missed due to the algorithm's detection delay. However, as mentioned previously, it is not an issue since the use model is to only count the number of seizures. Also, the other approach of buffering the EEG to not miss the beginning of a seizure would consume too much power. Table 5.1 summarizes the main differences between Modes 1 and 2.

5.2 Analog Front End Circuit Design

As mentioned in Chapter 3, the American Clinical Neurophysiology Society (ACNS) published a guideline on the recommended specifications for a system recording long-term EEG for epilepsy — the guidelines are presented again in Table 5.2 for conve-

Table 5.1: Differences between Modes 1 and 2

Parameter	Mode 1	Mode 2
Application	Diagnosis	Counting Seizures
External Device Usage	Continuously	Once a day for 15 minutes
Power	Not a major concern	Major concern
EEG recording	Continuously	Ten seconds after seizure
Sample Rate	1,280 Hz	256 Hz

Table 5.2: American Clinical Neurophysiology Society guideline for long-term epilepsy monitoring

Parameter	Recommended Value
Low-Frequency Response	0.5 Hz or lower
High-Frequency Response	70 Hz or higher
Noise Level	$< 1 \mu\text{V}_{\text{rms}}$
Input Impedance	$> 1 \text{ M}\Omega$
Common Mode Rejection	$> 60 \text{ dB}$
Dynamic Range	$> 40 \text{ dB}$

nience [25]. This guideline was used in the design of our analog front end (AFE).

The block diagram of the AFE is illustrated in Figure 5-1. It is composed of a low-noise amplifier (LNA), a programmable gain amplifier (PGA), an anti-alias filter, and an ADC. The ASIC has eight channels. There were two architecture design options: a fully parallel with eight ADCs and one with only one ADC. We decided to go with a fully parallel architecture to minimize cross talk between the channels. There was no power penalty between the two options since the power of having one ADC running eight times faster is basically the same as having eight ADCs running eight times slower. The only drawback is that the fully parallel option requires more area due to the larger number of ADCs.

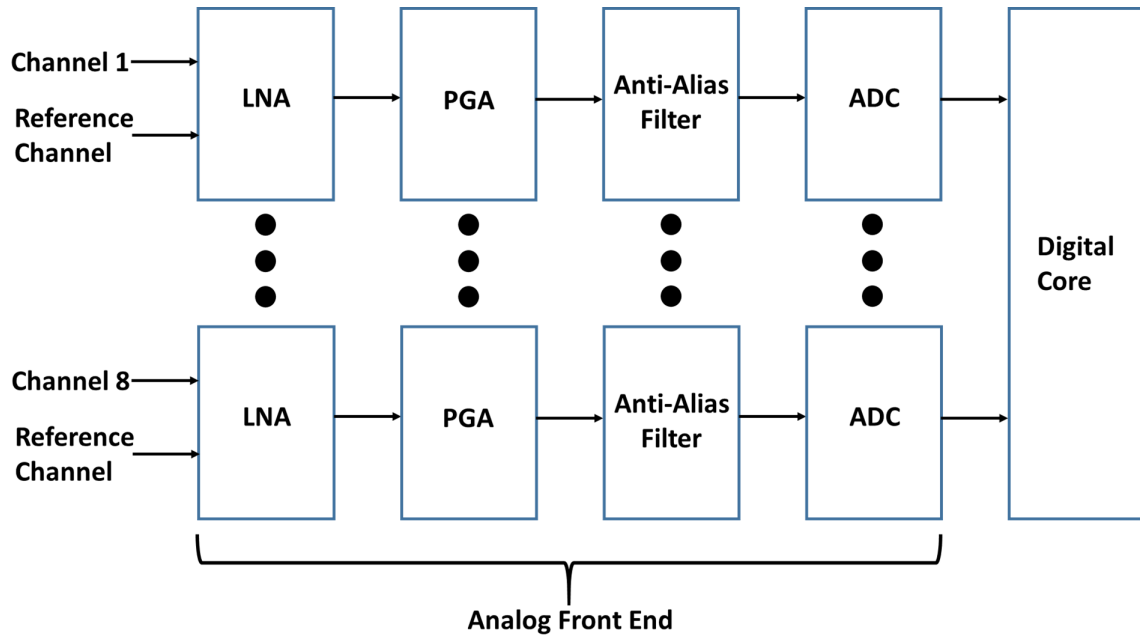


Figure 5-1: Analog front end block diagram.

5.2.1 Low-Noise Amplifier

The function of the low-noise amplifier (LNA) is to amplify the EEG coming from the electrodes with minimal noise and power consumption. Figure 5-2 shows the top-level circuit schematic of the LNA . The gain is set by the ratio of the input capacitor (24 pF) to feedback capacitor (200 fF), which yields a gain of 120. The 14 pF output capacitor was included to create a pole at approximately 700 Hz to prevent aliasing in the sampling stage of the anti-alias filter; more details in Section 5.2.3.

The input impedance of the LNA is dominated by the 200 fF feedback capacitors between the input and output terminals. Since our highest frequency of interest is 500 Hz, the lowest input impedance will be approximately $1.6 \text{ G}\Omega$, which is significantly higher than the required $1 \text{ M}\Omega$ value according to the ACNS's guideline in Table 5.2.

As mentioned in Chapter 3, when an electrode is placed on the skin, a potential voltage, commonly referred as electrode offset voltage (EOV), is generated [26]. To filter the EOV, thick-oxide PMOS transistors were placed in parallel with the feedback capacitors. By connecting the gate to the drain and source to the body, impedances in the $\text{T}\Omega$ range are realizable [39]. Figure 5-3 helps explain the operation of the




Figure 5-2: Top-level schematic of LNA.

pseudo-resistor. The diodes D1, D2, and D3 are junction diodes. If the voltage at the input is higher than the output, then diode D1 starts conducting and the impedance is reduced. If the voltage at the output is higher, then the parasitic PNP formed by source, well, and drain starts conducting and the impedance is reduced. When the voltage at the input and output are approximately equal, the diode and the parasitic PNP are off and a high impedance value is obtained. Simulations in the $0.18 \mu\text{m}$ process showed that the leakage currents through the diodes are on the order of 0.2 fA , thus not degrading the $\text{T}\Omega$ impedance. Figure 5-4, from [39], shows the impedance versus the voltage difference between input and output. Pseudo-resistors were also used in [40] and [41], and they measured a resistance of approximately $12 \text{ T}\Omega$. Assuming that we will obtain a similar value, our LNA will have a high-pass 3 dB cutin at $f_{3dB} = \frac{1}{2 * \pi * 12 \text{T}\Omega * 200 \text{fF}} = 0.065 \text{ Hz}$. The 3 dB cutoff was set significantly lower than 0.5 Hz to compensate for process variation, which, according to simulations, can vary by over a factor of two between slow and fast corners.

EEG signals can be as low as 0.5 Hz , so one of the main concerns in the LNA design is flicker noise. Chopping is one of the most common methods used to reduce it. The intention is to modulate the input signal before it goes through the LNA so

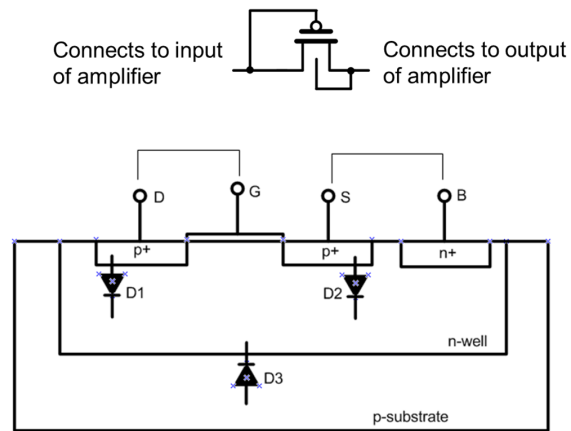


Figure 5-3: PMOS transistor with parasitic diodes that affect pseudo-resistor.




Figure 5-4: Measured resistance of the pseudo-resistor as the difference between input and output voltages are varied [39].

that flicker noise's spectrum does not overlap with the modulated signal. The signal is recovered by low-pass filtering and demodulating, which remove the flicker noise [42]. There are two main problems associated with chopping: first, it is done before the large input capacitors, reducing the amplifier's input resistance beyond the level desired for scalp electrodes [43]; second, finite gain of the amplifier causes first-order transient response leading to output errors [42]. To solve these problems, feedback is applied. However, noise from all the feedback paths starts to couple to the input hurting the noise level. Instead of chopping, we chose to use large transistors to minimize the effect of flicker noise. We are trading area for power to reduce the noise.

The chosen topology for the amplifier is fully-differential one-stage telescopic. This topology, shown in Figure 5-5, was chosen because only four transistors contribute to noise (M1, M2, M7, and M8), and with only one stage, power is minimized. The output-referred flicker noise power spectral density for this amplifier operating in subthreshold is given by Equation 5.1, where K is a process-dependent constant, W is the transistor's width, L is the transistor's length, C_{ox} is the oxide capacitance, and f is frequency [44]. There is a factor of 2 in the equation because M2 and M8 contribute the same amount of noise as M1 and M7, respectively.

$$\frac{v_o^2}{\Delta f} = \frac{2K_1}{W_1 L_1 C_{ox}^2 f} + \frac{2K_7}{W_7 L_7 C_{ox}^2 f} \quad (5.1)$$

To minimize flicker noise, we used PMOS input transistors since their K coefficient is smaller than NMOS's, and we also used large transistors for M1, M2, M7, and M8. The cascode transistors M2, M3, M4, and M5 contribute almost no noise [45]. They were included to increase the gain of the amplifier by increasing the output impedance. Table 5.3 lists the transistors dimensions. Transistors M3 through M6 had a length three times the minimum in order to reduce the effect of process variation. The length of M7 and M8 was set to 20 μm since it was the limit of the process technology. Their width was set to 150 μm to reduce flicker noise. The length of M1 and M2 didn't have to be as large as M7 and M8 because PMOS transistors generate less $1/f$ noise. A length of 2 μm was set to reduce transistor's geometry mismatch. The width was

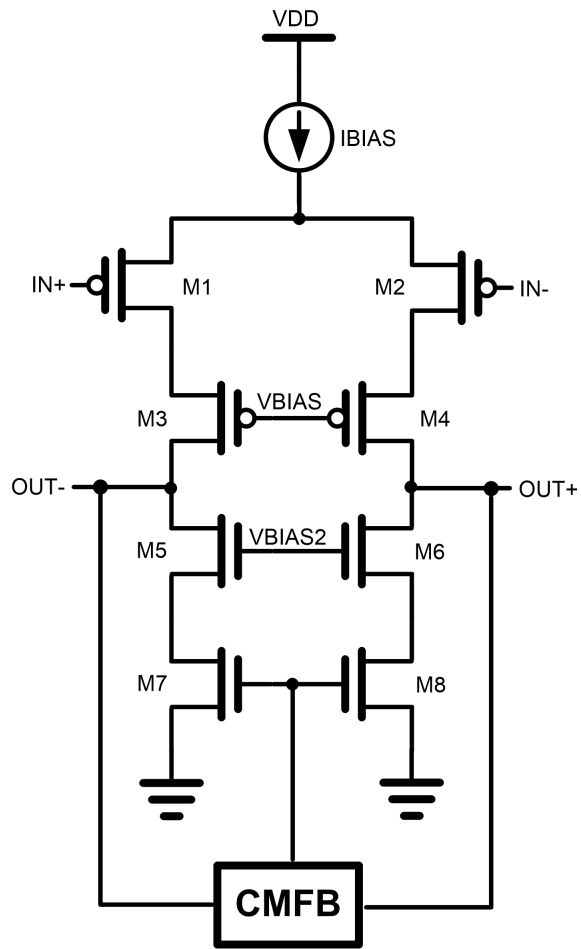


Figure 5-5: LNA schematic.

Table 5.3: Width and length of transistors used in the LNA (0.18 μm process)

Transistor	W (μm)	L (μm)
M1, M2	700	2
M3, M4	75	0.54
M5, M6	75	0.54
M7, M8	150	20

set to 700 μm to increase gain and reduce flicker noise.

We designed the LNA to have less than 1 μV_{rms} of input-referred noise in our bandwidth of interest (0.5 Hz - 500 Hz) to meet the ACNS guideline listed in Table 5.2. According to Spectre simulations, flicker noise contributed with 0.74 μV_{rms} to input-referred noise. Flicker and thermal noise are uncorrelated, so total noise power equals to the sum of their powers. As a result, thermal noise had to be limited to 0.67 μV_{rms} in order to meet the 1 μV_{rms} budget. Thermal noise is inversely proportional to the bias current going through the transistor [45]. As a result, according to Spectre simulations, the bias current had to be at least 1.2 μA to meet the noise budget.

In order to minimize power consumption, the supply voltage was reduced. All transistors in the LNA are in subthreshold, requiring approximately 100 mV between the drain and source to remain in saturation [45]. We have five transistors between VDD and ground, so the minimum supply voltage to have all transistor in saturation is approximately 0.5 V. We also need to take into consideration the voltage swing at the output nodes. The gain of the LNA is set to 120 and the maximum input voltage is 1 mV, thus the output node can swing between +120 mV to -120 mV. In order to keep M3, M4, M5, and M6 always in saturation, regardless of the input voltage, they require at least 220 mV between the drain and source. Therefore, the minimum supply is 740 mV. We decided to set the supply to 0.9 V, so that we are not running the transistors at the edge of saturation.

The amplifier was designed as fully differential to increase common-mode rejection and dynamic range. A common-mode feedback (CMFB) circuit, shown in Figure 5-6, is included to guarantee a stable common-mode voltage at the output by adjusting

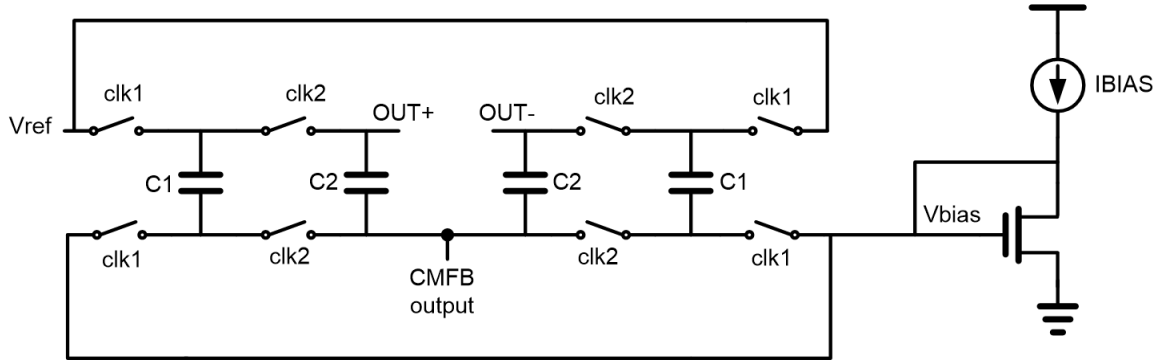


Figure 5-6: CMFB schematic.

the gate voltage of M7 and M8. Clk1 and Clk2 are non-overlapping clocks running at 2.56 kHz. The CMFB circuit tries to force the output common-mode voltage to be very close to V_{ref} , which equals $\frac{V_{DD}}{2}$ in our design. For example, if the output common-mode voltage increases beyond V_{ref} , then the voltage at CMFB output (gate voltage of M7 and M8) goes up. Since the gate voltage of M7 and M8 went up, the current going through them increases, forcing the output common-mode voltage to go back to V_{ref} [44].

Table 5.4 summarizes the simulated results of the amplifier. In Mode 2, the current and the bandwidth are reduced by a factor of five. This results in input noise staying practically constant since thermal noise is proportional to bandwidth and inversely proportional to current. The total power consumption reduced by less than a factor of five because the current used in some current mirrors and bias voltages was not lowered proportionally.

Table 5.5 compares the simulated results to the ACNS’s guideline. All of our simulated results exceed the recommended values.

5.2.2 Programmable Gain Amplifier

The second block in the AFE is the programmable gain amplifier (PGA). Figure 5-7 illustrates the top-level schematic of the PGA. Since the LNA’s bias current is very small ($1.6 \mu\text{A}$) we decided to use capacitors instead of resistor to set the gain in the PGA. If we had used resistors instead, their values would have to be very large (on

Table 5.4: Simulated results of the LNA in Modes 1 and 2

Parameter	Simulated Results Mode 1	Simulated Results Mode 2
Open loop DC gain	80.3 dB	72.6 dB
Unity gain frequency	89.9 kHz	19.9 kHz
Phase Margin	89°	88°
Closed loop midband gain	41.3 dB	41.1 dB
Closed loop bandwidth	0.07 Hz - 660 Hz	0.07 Hz - 159 Hz
CMRR at 60 Hz	77.8 dB	77.6 dB
Input Referred Noise	0.88 μ Vrms (0.5 Hz - 500 Hz)	0.91 μ Vrms (0.5 Hz - 100 Hz)
Power Consumption	1.39 μ W	0.47 μ W

Table 5.5: Simulated results of the LNA in Modes 1 and 2 compared to the American Clinical Neurophysiology Society guideline

Parameter	Recommended Value	Simulated Results Mode 1	Simulated Results Mode 2
Low-Frequency Response	0.5 Hz or lower	0.07 Hz	0.07 Hz
High-Frequency Response	70 Hz or higher	500 Hz	100 Hz
Noise Level	< 1 μ Vrms	0.88 μ Vrms	0.91 μ Vrms
Input Impedance	> 1 M Ω	1.6 G Ω	1.6 G Ω
Common Mode Rejection	> 60 dB	77.8 dB	77.6 dB
Dynamic Range	> 40 dB	60 dB	60 dB




Figure 5-7: PGA top-level schematic.

Table 5.6: Possible values of C_1 and the corresponding PGA gain

C_1 Value	PGA Gain
400 fF	12 dB
500 fF	14 dB
750 fF	17.5 dB
1.5 pF	23.5 dB

the order of tens of $M\Omega$) in order to not load the output of the LNA, increasing noise and also layout area. The pseudo-resistors (PMOS in parallel with C_2), just like in the LNA case, were included to set up the DC bias point.

The gain is set by $\frac{C_1}{C_2}$, where C_2 equals 100 fF and C_1 is a 2-bit programmable capacitor bank. Table 5.6 lists the possible values of C_1 and the corresponding gain setting. With the gain set to 17.5 dB, the PGA will saturate when the LNA input is over 1 mV, which is typically the maximum amplitude during a seizure. The gain setting of 12 dB and 23.5 dB were included so that the input amplitude that causes the PGA to saturate could be changed to 500 μV and 2 mV.

The PGA was designed as a fully-differential two-stage Miller-compensated op-amp, depicted in Figure 5-8. Two stages were used instead of one to maximize the

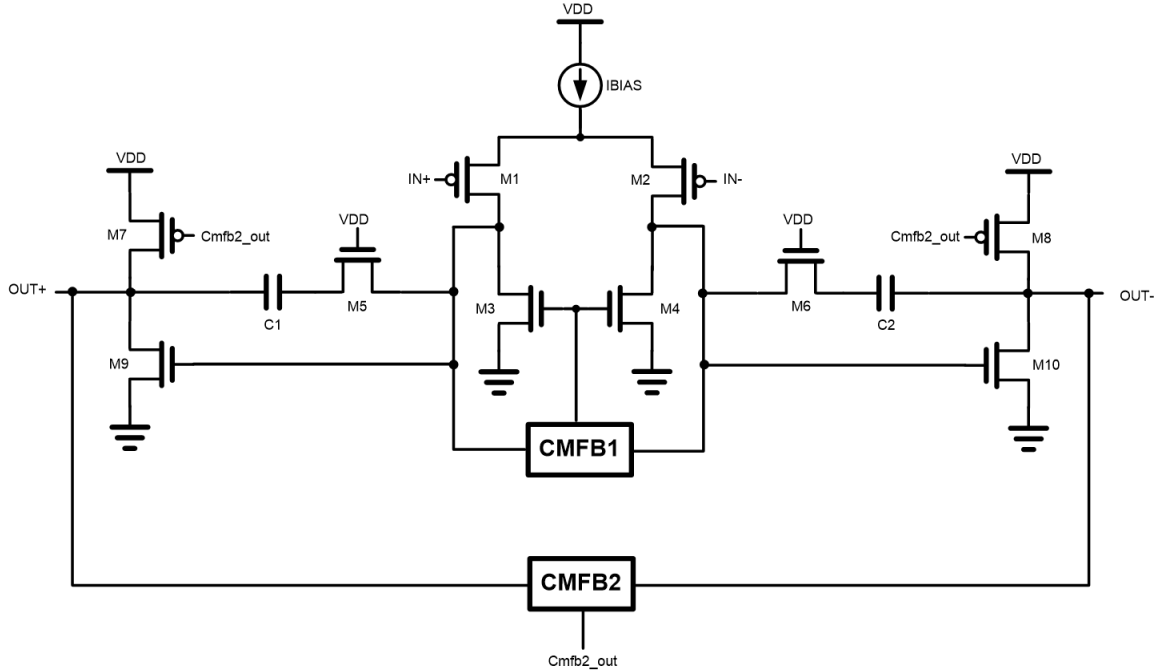


Figure 5-8: PGA schematic.

output swing because the output can be very close to the rails depending on the input. Flicker and thermal noise were not a major concern because the PGA comes after the LNA, which has a gain of 120. Therefore, the input transistors were not as big as in the LNA case and the bias current was reduced. The CMFB1 and CMFB2 were designed using switched capacitors similar to the one used in the LNA. Their operation is the same as explained in Section 5.2.1. Table 5.7 lists the simulated results of the PGA. Again, in Mode 2 the current and bandwidth are reduced by the same factor, causing noise to be similar to Mode 1's value.

5.2.3 Anti-Alias Filter

The anti-alias filter limits the bandwidth of the signal to avoid aliasing when sampling in the ADC. The ADC sample rate is 5.12 kHz and 1.024 kHz in Modes 1 and 2, respectively. Since the resolution of the ADC is nine bits, the required attenuation at half the sampling frequency (2.56 kHz in Mode 1 and 512 Hz in Mode 2) is 56 dB to ensure that the signal is below the ADC quantization noise. In Mode 1, the bandwidth of interest is 500 Hz, so we need a third-order filter with poles located

Table 5.7: Simulated results of the PGA in Modes 1 and 2

Parameter	Simulated Results	Simulated Results
	Mode 1	Mode 2
Open loop DC gain	86.2 dB	83.0 dB
Unity gain frequency	71.3 kHz	14.8 kHz
Phase Margin	70°	63°
Closed loop midband gain	17.3 dB	17.3 dB
Closed loop bandwidth	0.16 Hz - 5.9 kHz	0.16 Hz - 1.2 kHz
Input Referred Noise	9.6 μ Vrms (0.5 Hz - 500 Hz)	10.3 μ Vrms (0.5 Hz - 100 Hz)
Power Consumption	0.70 μ W	0.17 μ W

close to 500 Hz to get the required attenuation before 2.56 kHz (half the sample rate). The LNA has a pole located at 660 Hz, so the anti-alias filter only needs to be a second-order low-pass filter (LPF).

Several topologies were considered for the anti-alias filter, such as Gm-C, sinc, and active-RC. The second-order Gm-C filter described in [46] and seen in Figure 5-9 requires four transconductors to implement a second-order transfer function. This translates into higher power and also distortion at the output since the transconductance value varies with input amplitude. The sinc filter shown in Figure 5-10 and described in [47] also suffers from transconductance variation with input amplitude. To make matters worse, in order to get the required attenuation, two sinc filters would be needed. The input of the anti-alias filter can be very close to the rails during a seizure, resulting in significant distortion at the output of the sinc filters. A second-order LPF can also be implemented with an active-RC with just one amplifier and passive components. The gain is not dependent on the input amplitude, which minimizes distortion. For these reasons an active-RC was chosen. A Rauch topology based on [48] and shown in Figure 5-11 was used because it provides a second-order transfer function with reduced sensitivity to component variation as is explained later in this section. The transfer function is described by Equation 5.2.

$$H(s) = -\frac{1}{\frac{R_1}{R_3} + sC_2(R_1 + R_2 + \frac{R_1R_2}{R_3}) + 2s^2R_1R_2C_1C_2} \quad (5.2)$$

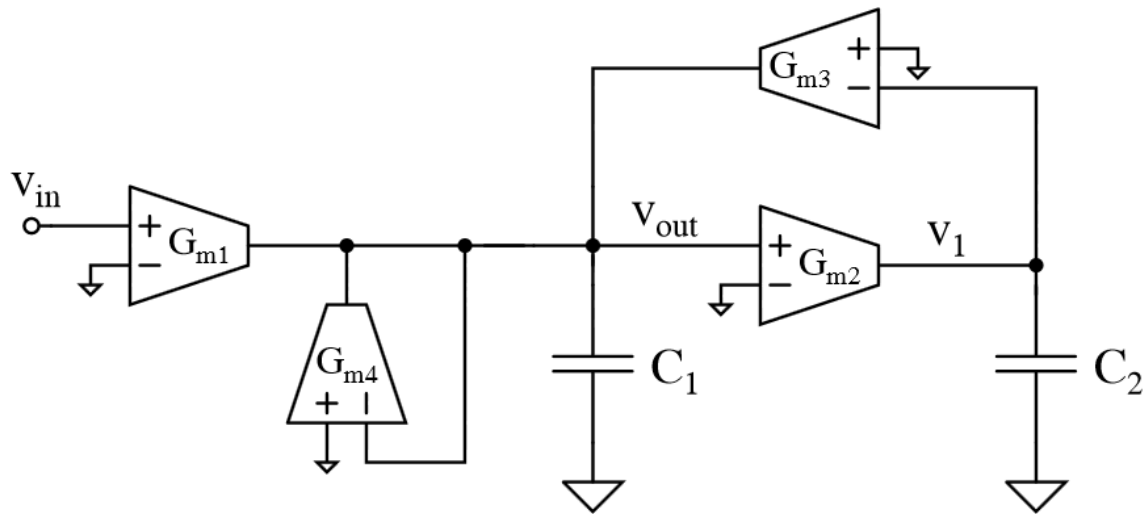


Figure 5-9: Schematic of a second-order Gm-C filter [46].

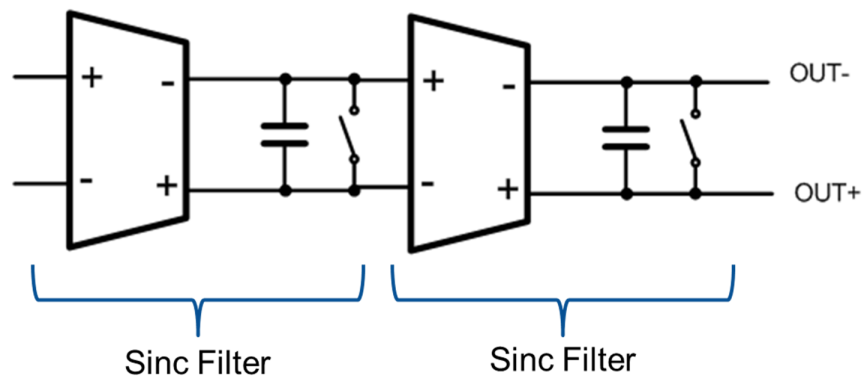


Figure 5-10: Schematic of a second-order sinc filter [47].




Figure 5-11: Schematic of Rauch filter based on [48].

In order to design the Rauch filter as being maximally flat, meaning that the passband will have a flat frequency response like a Butterworth filter, the values of the resistors and capacitors need to be set according to Equation 5.3. The location of the poles are then given by Equation 5.4 [49].

$$R_1 = R_3 = 2R_2 = 2R \quad C_1 = 2C_2 = 2C \quad (5.3)$$

$$f_p = \frac{1}{4\pi RC\sqrt{2}} \quad (5.4)$$

According to Equation 5.4, the poles location depends on the absolute value of the resistor and capacitor, which can fluctuate due to process variation. To minimize the effect of process variation, the poles should be dependent on the ratio of resistors or capacitors and not on their absolute values. This can be accomplished by implementing the resistors using switched capacitors, as shown in Figure 5-12. The switches S1 and S2 are turned on and off by two non-overlapping clocks running at the same frequency. Equation 5.5 describes the resistance value between input and




Figure 5-12: Implementing a resistor using switched capacitor.

output using a switched capacitor.

$$R = \frac{1}{f_s C} \quad (5.5)$$

By replacing all the resistors with switched capacitors, the schematic of the Rauch filter changes to the one depicted by Figure 5-13, where Clk1 and Clk2 are non-overlapping clocks. We used transmission gates with minimum sizes for the switches to reduce the impact of charge injection. The architecture is fully differential, which also helps minimize charge injection.

If we use Equation 5.3 to obtain maximum flatness, then the switched capacitors values should follow the relationship $C_4 = 2C_3 = 2C_5 = C_{req}$. With these values, the location of the poles is described by Equation 5.6. Instead of being dependent on the absolute value of the resistors or capacitors, the poles are now determined by the ratio of capacitors ($\frac{C_{req}}{C}$), which is not significantly impacted by process variation if common-centroid technique is used in layout. Table 5.8 lists the capacitors values used in the design.

$$f_p = \frac{C_{req} f_s}{4\pi C \sqrt{2}} \quad (5.6)$$

In Mode 1, the EEG bandwidth of interest is 500 Hz, so the poles were designed to be at 560 Hz to compensate for any clock frequency mismatch or process variation that could push them lower in frequency. In Mode 2, the bandwidth is reduced by a factor of five, so the poles should be closer to 100 Hz. According to Equation 5.6, this is easily accomplished by reducing the clock frequency by a factor of five.

As presented in Figure 5-13, the output of the PGA is sampled into capacitors




Figure 5-13: Rauch filter schematic using switched capacitors to replace resistors. Clk1 and Clk2 are non-overlapping clocks.

Table 5.8: Capacitor values used in the anti-alias filter

Capacitor	Value
C_1	5.12 pF
C_2	2.56 pF
C_3	320 fF
C_4	640 fF
C_5	320 fF

Table 5.9: Simulated results of the anti-alias amplifier in Modes 1 and 2

Parameter	Simulated Results Mode 1	Simulated Results Mode 2
Open loop DC gain	85.1 dB	80.2 dB
Unity gain frequency	43.4 kHz	10.3 kHz
Phase Margin	76°	68°
Closed loop DC gain	-0.23 dB	-0.27 dB
Closed loop bandwidth	560 Hz	117 Hz
Input Referred Noise	4.2 μ Vrms (0.5 Hz - 500 Hz)	4.9 μ Vrms (0.5 Hz - 100 Hz)
Power Consumption	0.84 μ W	0.20 μ W

C3, so we need to make sure there is no aliasing. The inputs are sampled at 40.96 kHz in Mode 1, which means that the signal’s bandwidth must be limited to 20.48 kHz. This is accomplished by the LNA and PGA. The LNA has a pole located at 660 Hz and the PGA at 5.9 kHz. These two poles provide an attenuation of 44 dB at 20.48 kHz. We don’t expect any interferer at frequencies higher than 20.48 kHz, so the attenuation provided by the poles should be sufficient. In Mode 2, the anti-alias samples the input at 8.192 kHz (five times lower than in Mode 1). Since the closed-loop bandwidth of the LNA and PGA are also reduced by a factor of five as a result of the bias current reduction, the attenuation is still 44 dB at half the sample frequency.

The amplifier used in the anti-alias filter was designed as a fully differential two-stage Miller-compensated similar to the PGA’s amplifier. Table 5.9 lists the simulated results of the anti-alias filter. In Mode 2, the bias current and clock frequency used by the switched capacitors are reduced by a factor of five.

5.2.4 ADC

The ADC employed in this ASIC was reported in [50]. It is a 9-bit successive approximation register (SAR) ADC with expected 53.6 dB of SNDR and 8.6 effective number of bits (ENOB). In Modes 1 and 2 the ADC converts 5,120 and 1,024 samples a second, respectively. The ADC oversamples the data by a factor of four in order

to increase the resolution from 9 to 10 bits in the digital core. It runs at 0.9 V and consumes 158 nW and 32 nW when in Modes 1 and 2, respectively.

5.3 Digital Core

The digital core has three main functions: decimate the EEG data by a factor of four, store all registers that control the AFE bias currents, and communicate with external device through an inductive link. The first two functions are described in more detail in the next subsections. The last one is explained in Chapter 6. One important consideration is that the seizure detection implementation is not included in the digital core — it is a block by itself.

5.3.1 Data Decimation

The digital core decimates the EEG data by a factor of four to increase the resolution from 9 to 10 bits. The decimation uses polyphase filters to reduce computation requirements [29]. The filter was designed using MATLAB’s filter design tool, which also generated the FIR Verilog code [51] and [52]. In Mode 1, we set the passband and stopband to 500 Hz and 640 Hz while in Mode 2, the passband and stopband were set to 100 Hz and 128 Hz, respectively. Since the passband and stopband of Modes 1 and 2 differ by a factor of five, we were able to use the same filter block in both modes. The only difference is the clock frequency running the decimation block.

The filter was designed using the MATLAB filter design toolbox and has 106 taps. The magnitude response in Mode 1 is shown in Figure 5-14. The filter was designed using a coefficient word length of 16 bits to minimize the impact of quantization in the transfer function. This resulted in a difference of less than 1 dB between the ideal transfer function and the quantized one.

The expected power consumption of the decimation filter is 1.01 μ W in Mode 1 and 0.21 μ W in Mode 2. The difference in power is due to the slower clock in Mode 2.




Figure 5-14: Magnitude response of the decimation filter in Mode 1. The difference between the quantized and reference is less than 1 dB.

5.3.2 Registers for AFE

The bias current of the LNA, PGA and anti-alias filter can be changed through a two-bit register stored in the digital core. As illustrated in Figure 5-15, BIT1 and BIT0 adjust the total bias current going to the analog block by turning on or off switches M1 and M2. Each block has its own register, so their currents can be changed independently of each other. We used two bits to change the bias, so there are a total of four possible current values: I_{max} , $I_{max}/2$, $I_{max}/4$, and $I_{max}/5$. Typically, the bias is set to the maximum (I_{max}) in Mode 1 (BIT = 11) and to the minimum ($I_{max}/5$) in Mode 2 (BIT = 00). However, the registers can be changed at any time in any mode to give maximum flexibility.

5.4 Seizure Detection

In Mode 2, we are counting the number of seizures a patient has, so the EEG data goes through the seizure detection block, as shown in Figure 5-16. The first step in this process, as explained in Chapter 4, is to measure the energy in four different frequency bands. Since we are not concerned about the phase of the signal, we can use IIR instead of FIR filters to significantly save computational requirements. The IIR filters were implemented as Chebyshev Type II in order to minimize the filter's order while still keeping the passband maximally flat [29]. The four filters were designed

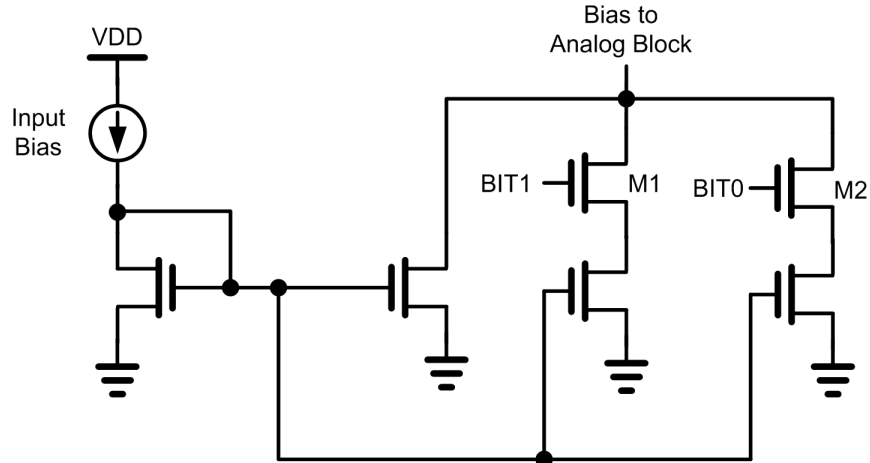


Figure 5-15: Schematic of current gain block.

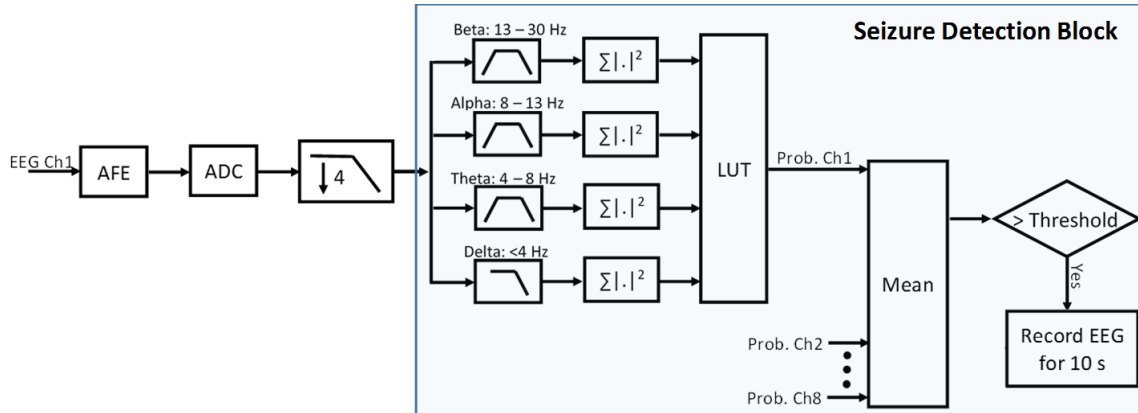


Figure 5-16: Block diagram of the seizure detection implementation.

using MATLAB’s filter design tool, which also generated the Verilog code [51] and [52]. Figures 5-17, 5-18, 5-19, and 5-20 show the magnitude response of the Beta, Alpha, Theta, and Delta filters, respectively. The coefficient word length of each filter was chosen so that the impact of quantization is minimized.

The next step in the algorithm is to measure the energy at the output of each filter for two seconds. As explained in Chapter 4, this is accomplished by summing the square of the voltage values, so we square the output of the IIR filters and add them in an accumulator. The output from the IIR is represented using 10 bits, so when the value is squared, we need 20 bits to represent the new value. The accumulator adds the squared values for two seconds, and the sample rate is 256 Hz, so the accumulator




Figure 5-17: Magnitude response of the reference (red) and quantized (blue) Beta filter (13 Hz - 30 Hz).




Figure 5-18: Magnitude response of the reference (red) and quantized (blue) Alpha filter (8 Hz - 13 Hz).




Figure 5-19: Magnitude response of the reference (red) and quantized (blue) Theta filter (4 Hz - 8 Hz).




Figure 5-20: Magnitude response of the reference (red) and quantized (blue) Delta filter (0.5 Hz - 4 Hz).

requires at least 29 bits to not overflow. The output of the accumulator, which is the measured energy, is binned into eight possible values, requiring three bits. The four binned energy values (one from each filter output) are used as an index in the lookup table (LUT) to find the probability value of having a seizure. The LUT is limited to 50 rows, as it was discussed in Chapter 4, and each row contains 12 bits for the energy values (four filters and each one requires three bits to represent the energy value) and eight bits for the probability value. As a result the LUT of each channel requires a total of 1,000 bits.

Finally, we take the mean of the probability output from each channel's LUT and compare it with a stored threshold value. If the output of the mean is higher than the threshold, a seizure is declared and ten seconds of data are stored into an external flash memory for later analysis. The system only stores data after it detects a seizure, which means that, on average, the first nine seconds of the seizure are not recorded due to detection delay. We implemented the system this way to minimize power, otherwise, the system would need to continuously record the data.

Table 5.10 lists the estimated power consumption of each block according to our ASIC simulations. The "other" block shown in the table includes the multiplier to measure energy and also some overhead, such as writing and reading from registers, etc. As observed, no block dominates the power consumption — all consume roughly the same percentage.

Table 5.10: Estimated power consumption of individual blocks in seizure detection

Block	Expected Power Consumption (nW)	Area mm ²
Beta Filter	107	0.083
Alpha Filter	77	0.057
Theta Filter	89	0.064
Delta Filter	98	0.072
LUT	108	0.106
Other	75	0.084

5.5 Peripheral Circuits

5.5.1 Clock Generation

The clock used in the ASIC is based on [53] and illustrated in Figure 5-21. When ϕ is low, S1 and S4 are on. The current I going through S4 induces a voltage at $V_2 = IR$. The current going through S1 starts charging the capacitor C_1 . When V_1 becomes greater than V_2 , ϕ goes high turning switches S1 and S4 off and turning S2 and S3 on. Now, $V_1 = IR$ and capacitor C_2 is charged with current I. This results in an output frequency equal to $\frac{1}{2(RC+t_{delay})}$, where t_{delay} is the comparator delay and $C_1 = C_2 = C$. By switching the input node of the comparator that is connected to the resistor R, the comparator's offset is eliminated, improving frequency stability [53]. In our design we set the current to 100 nA, the capacitors to 900 fF, and the resistor to 930 k Ω in order to get a clock frequency of 491.52 kHz. The clock frequency is set to this value because it is a common multiple between the required ADC clock (81.92 kHz) and the block that transmits EEG data through the inductive link (450.56 kHz). In order to get a clock frequency very close to the desirable value, we implemented the resistor R using an off-chip component. The expected power consumption of this block is 890 nW.

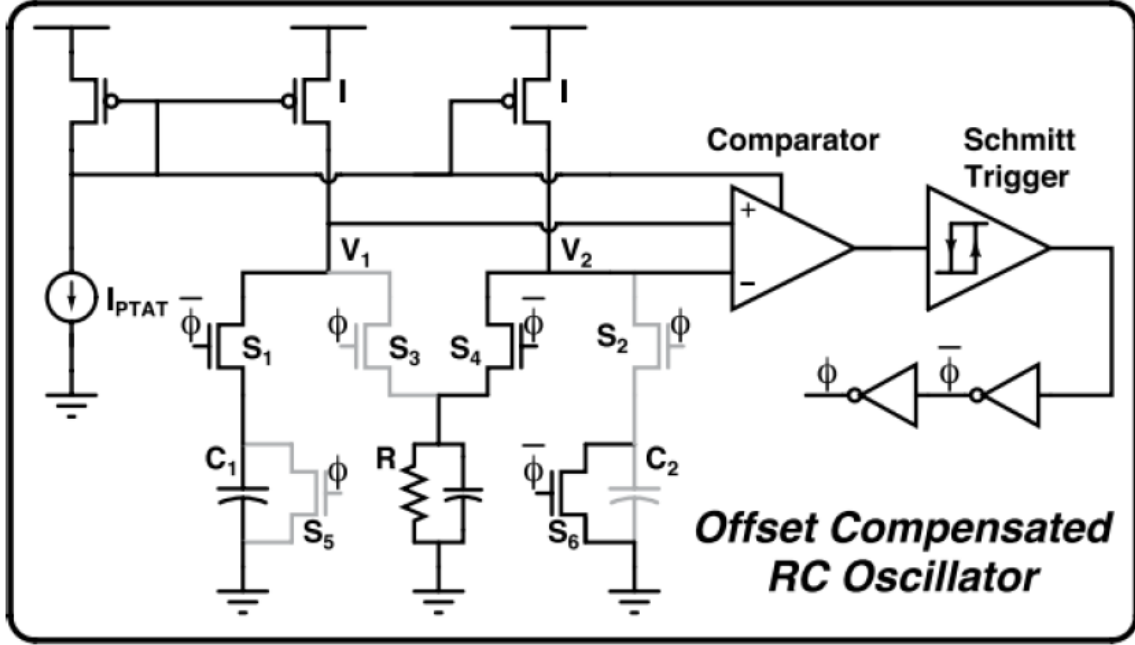


Figure 5-21: Schematic of clock circuit [53].

5.5.2 Current Reference

The current reference used to bias the LNA, PGA and anti-alias filter is based on the work done in [54] and is shown in Figure 5-22. The actual current reference is formed by transistors M1 through M5, resistor R, and capacitors C1 and C2. The remaining components (M6, M7, and C3) are included to guarantee that the reference starts up. If the startup circuit is not included, then a stable condition exists where the gates of M1 and M2 are at VDD while the gates of M3 and M4 are at ground, resulting in no output current. To prevent this condition, the startup circuit is required. When VDD turns on, capacitors C1 through C3 force the gates of M1 and M6 to be at VDD while M3 is forced to be at ground. M6 starts charging the gate of M3 until M7 turns on and shuts M6 off. At this point, the current reference is on and the startup circuitry is off.

The current is determined according to Equation 5.7. In our case, we wanted I_{out} to equal 25 nA, so we set the ratio W_3/W_4 to 2 and R to 835 k Ω .

$$I_{out} = \frac{V_{thermal}}{R} \ln(W_3/W_4) \quad (5.7)$$




Figure 5-22: Current reference schematic.

5.5.3 Voltage Reference

A switched-capacitor voltage reference, illustrated in Figure 5-23, generates a $\frac{V_{DD}}{2}$ reference that is used by the amplifiers' common mode feedback circuits (CMFB). Depending on the clock phase, the capacitor on the left is either connected between VDD and right capacitor, or between right capacitor and ground. Since both capacitors are equal, the output voltage equals $\frac{V_{DD}}{2}$. The CMFB was also designed with switched capacitors, so it samples the voltage reference in every clock cycle. In order to minimize the transient noise caused by the CMFB sampling and the reference switching, an output buffer was included and both clocks (reference and CMFB) run at the same frequency. The expected power consumption of the reference is 87 nW.

5.6 Physical Layout Considerations

Careful layout is crucial when amplifiers require high CMRR and have transistors that operate in subthreshold. For example, our first layout version of the LNA did not have a perfectly symmetric routing between the input capacitors and input terminals of the amplifier. This resulted in a 2 fF parasitic mismatch between the input terminals of the amplifier and ground, as illustrated by Figure 5-24, reducing CMRR by 21 dB.




Figure 5-23: Voltage reference schematic.




Figure 5-24: Top-level LNA schematic with mismatched parasitics at input terminals of the amplifier.

All transistors and capacitors that required matching were laid out using common-centroid with edge dummy technique. When capacitors or transistors had different values, they were all multiples of a specific unit cell such that common-centroid could still be applied.

5.7 Measurement Results

An ASIC, which measures 5 mm by 5 mm, was fabricated using the TSMC 0.18 μm process. A photomicrograph can be seen in Figure 5-25. It contains eight channels to measure EEG and a test channel that was included in order to characterize the outputs from each block in the analog front end (LNA, PGA, anti-alias filter, and ADC).




Figure 5-25: Die micrograph of the ASIC with circuit blocks labeled.




Figure 5-26: LNA frequency response in Mode 1 (blue) and Mode 2 (red).




Figure 5-27: PGA frequency response in Mode 1 (blue) and Mode 2 (red).

5.7.1 Frequency Response

This section shows the frequency response plots for each test channel block in Modes 1 and 2. The frequency response was measured by applying a 1 mV sine wave at different frequencies to the input. Figures 5-26, 5-27, and 5-28 show the Mode 1 and 2 frequency response of the LNA, PGA, and anti-alias filter, respectively. Figure 5-29 plots the frequency response from the LNA, PGA, and anti-alias filter combined. This is the frequency response that the EEG data goes through from the input of the LNA until the input of the ADC.

In Mode 2 the bandwidth of the LNA and PGA are reduced due to a lower bias current. The anti-alias filter's bandwidth decreases because the clock frequency is reduced by a factor of five, pushing the poles from 650 Hz to 125 Hz approximately.




Figure 5-28: Anti-alias filter frequency response in Mode 1 (blue) and Mode 2 (red). Less than 1 dB difference from simulation.




Figure 5-29: AFE (from LNA to Filter) frequency response in Mode 1 (blue) and Mode 2 (red).




Figure 5-30: Output noise power spectral density in Modes 1 (blue) and 2 (red) measured with Agilent 35670A Dynamic Signal Analyzer.

5.7.2 Noise Measurement

The noise was measured at the output of the PGA with its gain set to the maximum value of 23.5 dB in order to guarantee that the measured noise was well above the equipment's noise floor. The measurement was taken with the Agilent 35670A Dynamic Signal Analyzer for both modes of operation.

Figure 5-30 plots the output noise power spectral density (PSD) in Modes 1 and 2. To calculate the input-referred noise PSD, we must divide the output noise by the square of the LNA and PGA gain. The measured frequency response of the LNA and PGA combined is shown in Figure 5-31 and the input-referred noise PSD in both modes is shown in Figure 5-32.

To calculate the total input-referred noise in Mode 1, we need to integrate Figure 5-32 in the bandwidth of interest, which is 0.5 Hz to 500 Hz, and take the square-root of the value. Following this procedure, we get a total input noise of 1.13 μVrms . This value is 17% higher than the simulated value of 0.96 μVrms . In Mode 2, the bias current is reduced by a factor of five, so flicker noise stays constant while thermal noise increases roughly by a factor of five. This happens because flicker noise does not depend on bias current while thermal noise is inversely proportional to it. The bandwidth of interest now is 0.5 Hz to 100 Hz, so the total input referred noise equals 1.06 μVrms . This value is very close to one measured in Mode 1 because bias current




Figure 5-31: Frequency response measured at the output of the PGA with its gain set to maximum in Modes 1 (blue) and 2 (red).




Figure 5-32: Input-referred noise power spectral density in Modes 1 (blue) and 2 (red).

Table 5.11: Measured versus simulated power of main blocks in the ASIC

Block	Simulated (μW)	Measured (μW)
AFE (Mode 1)	23.5	22.0
AFE (Mode 2)	6.9	6.7
Digital Core (Mode 1)	7.7	8.0
Digital Core (Mode 2)	3.3	3.5
Seizure Detection	3.9	3.6

and bandwidth were reduced by the same factor.

5.7.3 Power Measurement

Table 5.11 compares the measured versus simulated power of different blocks in the ASIC. The measured power of all blocks matched simulated values very closely.

Figure 5-33 shows a pie chart indicating the power breakdown in Mode 2. The total power consumption per channel is $1.86 \mu\text{W}$.

5.7.4 AFE Comparison

A widely used figure of merit for low-noise amplifiers is the Noise Efficiency Factor (NEF), which was proposed in [55] and described by Equation 5.8, where $V_{ni,rms}$ is the input-referred RMS noise voltage, I_{total} is the total amplifier supply current, $V_{thermal}$ is the thermal voltage, and BW is the amplifier bandwidth. The measured NEF for our LNA is 2.2. Figure 5-34 illustrates how this value compares with other ones published in literature. Ideally, the amplifier should be as close to the $NEF = 1$ line as possible.

$$NEF = V_{ni,rms} \sqrt{\frac{2I_{total}}{\pi V_{thermal} 4KTBW}} \quad (5.8)$$

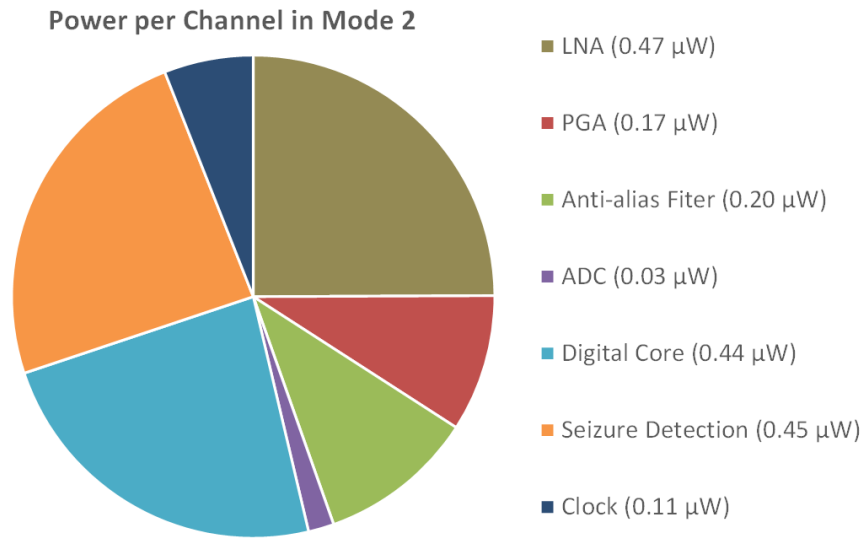


Figure 5-33: Pie chart indicating the power breakdown of each block in Mode 2.




Figure 5-34: Comparison of the LNA with published results in literature. The data is adapted from [39] and [56].

Table 5.12: Our AFE compared with published literature

Parameter	[56]	[34]	[57]	This System
AFE Power (μ W/channel)	0.69	8.25	1.26	2.75
EEG Bandwidth (Hz)	150	100	360	500
Input Referred Noise	1.15 μ Vrms (0.5 Hz - 150 Hz)	0.91 μ Vrms (0.5 Hz - 100 Hz)	2.2 μ Vrms (0.1 Hz - 360 Hz)	1.13 μ Vrms (0.5 Hz - 500 Hz)
NEF	2.2	5.1	3.5	2.2
Supply Voltage	1 V	1.8 V	1.8 V	0.9 V
Number of Channels	32	8	8	8

Table 5.12 compares our AFE with other ones published in literature. The best power per channel was achieved by [56]; however, the bias voltages used by his LNA are produced off-chip, which helped reduce his power number. Our NEF is comparable to state of the art while also operating at a higher bandwidth, enabling the recording of high-frequency oscillation in the EEG.

5.7.5 Seizure Detection

The seizure detection algorithm in the ASIC was tested using the same PhysioNet dataset described in Chapter 4. The first step was to train on all seizure records except for one. Following the procedure explained in Chapter 4, we generated the LUT for each channel in MATLAB. Then, we programmed the IC’s LUT using an SPI interface. We set the threshold to be 0.25 since, according to Figures 4-20 and 4-21, it provides maximal seizure detection without an excessive false alarm rate. Finally, the test record was sent to the IC through another SPI interface and the IC declared every second of data as non-seizure or seizure. The results were compared with the ones obtained in MATLAB and they agreed. So, for a threshold of 0.25, the IC detected 98.5% of the seizures while having a false alarm rate of 4.4 per hour.

Table 5.13 compares our results with previous published ones. Our algorithm consumes over an order of magnitude less power than comparable ones in the literature while still detecting 98.5% of seizures. It does have a higher number of false alarms, which means we need to store more data to be later analyzed than others. As mentioned in Chapter 4, the recorded data goes through a more complex algorithm in a

Table 5.13: Seizure detection results compared with published literature

Parameter	[34]	[58]	[59]	This System
Power ($\mu\text{W}/\text{channel}$)	24.23	21.8	5.9	0.45
Sensitivity	84.4%	95%	96%	98.5%
False Alarm per Hour	1.3	0.27	0.2	4.4
Number of Channels	8	8	18	8

computer to reduce the number of false alarms to a value small enough that a doctor can review. In order to estimate the power penalty of recording extra data, we used a 128 Mbit flash memory (MX25U12835FZ2I) and measured its power when writing to it at the expected rate. The memory consumes $1.3 \mu\text{W}$ per channel. Therefore, if we include this power penalty, our algorithm still performs better than the other systems by a factor of three.

The system only starts to record the EEG after the algorithm detects a seizure, which on average takes 9.1 seconds. Another option would have been to keep buffering ten seconds of EEG data into a memory so that when a seizure is detected, its beginning wouldn't be missed. In order to buffer ten seconds of EEG, the memory must be at least 200 kbits (8 channels, 10 bits per sample, 256 Hz sample rate, and ten seconds of data). The power required to buffer the data is approximately $90 \mu\text{W}$; it was estimated using a 256 kbit SRAM from Cypress (CY62256EV18) because it was one of the most energy efficient SRAM available. In Mode 2, the IC's power consumption (including AFE and seizure detection) is only $13.8 \mu\text{W}$, so including the buffer would significantly degrade the overall system's power consumption.

5.8 Summary

This paper presented the design of an ASIC to be used in a minimally invasive sub-dermal implantable EEG recorder and seizure detector. The ASIC can record up to eight channels of EEG and uses a power-efficient algorithm to detect seizures. In the diagnosis mode, the ASIC records EEG up to five times the typical bandwidth.

The higher bandwidth enables recording of HFO, which can potentially improve diagnosis. The ASIC's low-power consumption enables the design of smaller and better ambulatory EEG systems. It offers patients with epilepsy a more convenient way of monitoring their seizures and gives doctors more accurate information for diagnosis and treatment options.

Chapter 6

Subdermal Implantable System

This chapter presents the design of a subdermal implantable system that was tested in a swine at MGH. The system consists of two parts: an implant located between the skin and skull behind the right or left ear, and an external device, designed by Jason Yang, located on top of the skin.

6.1 System Overview

Figure 6-1 presents the system's block diagram. The system has eight EEG channels but requires nine electrodes since all EEG channels are measured with respect to a reference. The two coil blocks form an inductive link that is used by the implant and the external device.

6.2 Inductive Link

The inductive link is responsible for: transferring power from the external to the implant and transmitting data between them. We decided to use the same coil to send power and data in order to reduce implant's size.

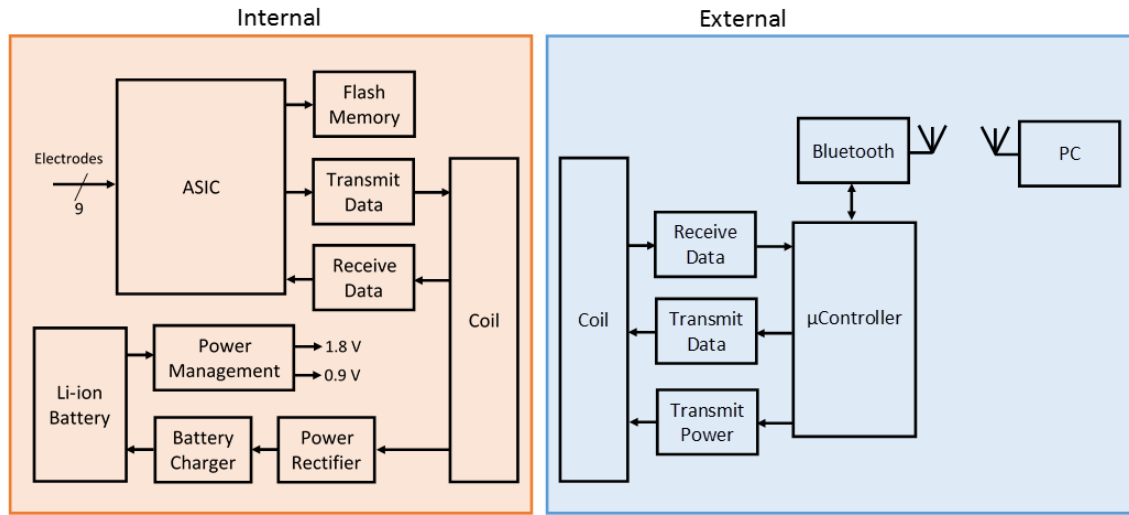


Figure 6-1: Block diagram of the system.

6.2.1 Power Transfer

The implant uses the power received from the external device to recharge a battery. The power transferred through the inductive link is generated at the external device by a class-E power amplifier, which is shown in Figure 6-2. This type of power amplifier was chosen because its efficiency can approach 100% [60].

Class-E Power Amplifier

According to Figure 6-2, the microcontroller sets the switching frequency by controlling the gate voltage of M1, turning it on/off with a 50% duty cycle clock. L1 is chosen to be a large inductor so that the current that goes through it is almost constant. When M1 is off, the current through L1 charges the capacitor C1 and increases M1's drain voltage. As the voltage goes up, current starts flowing through the resonant circuit consisting of C2 and external coil, causing the drain voltage to discharge. When the drain voltage is close to zero, the microcontroller turns M1 on for half of the clock period. Then, M1 is turned off and the process continues. The resonant circuit (C2 and external coil) filters any harmonics at the coil, producing a sinusoidal waveform at the resonant frequency. A more in depth explanation of a class-E amplifier and the equations to calculate the inductors and capacitors values




Figure 6-2: Schematic of a class-E power amplifier.




Figure 6-3: Schematic of circuit used to rectify the power from the coil.

can be found in [60] and [61].

Receiving Power at the Implant

The implant needs to rectify the power coming from the inductive link to generate a stable voltage to recharge a lithium-ion battery. The voltage needs to be at least 4.25 V, the minimum acceptable input voltage for most battery charger ICs.

To rectify the sinusoidal voltage at the internal coil, we used a diode and a large capacitor ($22 \mu\text{F}$), illustrated in Figure 6-3 as D1 and C2, respectively. The capacitor C1 was set so that it resonates with the internal coil at 5 MHz, which is the frequency of power transfer. Section 6.2.4 explains why we used 5 MHz.

To estimate the ripple at capacitor C2, we took into account the 2 mA needed to recharge the battery in 15 minutes (more details in Section 6.3.2). The power transfer frequency is 5 MHz, so capacitor C2 gets recharged every 200 ns. To calculate the ripple, we use Equation 6.1, where we set Δt to 200 ns and C to $22 \mu\text{F}$. With these values, the ripple voltage is only $18 \mu\text{V}$.

$$\Delta V = \frac{I\Delta t}{C} \quad (6.1)$$

The zener diode Z1 is included as a protection device. If the voltage goes above 5.1 V, the zener starts conducting current and limits the voltage at that node. This is necessary because the external device doesn't get any feedback on what the voltage at the internal coil is, so it might send more power than needed.

6.2.2 Programming the ASIC

The ASIC has registers that can be programmed by the external device: the probability LUT for each channel, the threshold value, and the bias currents of the AFE. Ideally these registers would be programmed only once, so we traded spectrum efficiency with circuit simplicity to transmit and decode. We decided to use on-off keying (OOK) since this could easily be implemented in the external device by simply turning on/off the class-E power amplifier. To decode a typical OOK, the decoder needs to count how long the input is on or off. An example is presented in Figure 6-4, where each unit of time corresponds to one bit. A problem arises if the receiver and transmitter timers don't match, then the receiver might decode the data incorrectly. For example, instead of decoding bits 11 between time 3 and 5 in Figure 6-4, it could decode just 1 or even 111.

To make the data transfer more insensitive to timer mismatch, we used a scheme similar to Manchester encoding, but requiring two transitions per bit instead of one. For each bit, the external power needs to be on for some time and off for a different amount of time. If the power is on for a longer period of time than it is off, then the bit is 1, otherwise it is 0. Figure 6-5 illustrates this scheme. To minimize error, the longer period always equals 75% of the time taken to transfer each bit while the short period equals the remaining 25%. For example, if each bit takes four seconds to transmit, then a bit 1 would stay on for 3 s and off for 1 s while a bit 0 would stay on for 1 s and off for 3 s. This scheme is more robust to timer mismatch because the bit is decoded based on the ratio of time spent on versus off. It is, however, more




Figure 6-4: Example of a typical OOK.

inefficient in terms of how many bits it can transfer for a given amount of time, but since we only program the ASIC once, we are not concerned about maximizing the spectral efficiency.

The decoder clock runs at the maximum frequency of 491 kHz, and each bit requires at least 4 clock cycles to be received correctly (1 clock during on and 3 during off or vice versa), so the fastest data rate acceptable is 122.75 kbits per second. The register used to count how long the input is high or low is eight bits long. So, the slowest rate that you can program the ASIC before the counter overflows is $\frac{491kHz}{\frac{256}{0.75}} = 2.56$ kbits per second. It is important to send data faster than the minimum because if the counter overflows, the ASIC might decode the bit incorrectly.

In order to decode the data coming in, we used a simple envelope detector, depicted by Figure 6-6, with resistor R1 and capacitor C3 set to 40 k Ω and 100 pF, respectively. These values yield a corner frequency of 40 kHz, which is well below the 5 MHz carrier frequency. The Schmitt trigger was included to avoid glitches during the transitions. The output of the Schmitt trigger goes to the ASIC, which decodes the data. Figure 6-7 shows an example of bits 1 (between 0 and approximately 105 μ s) and 0 (between

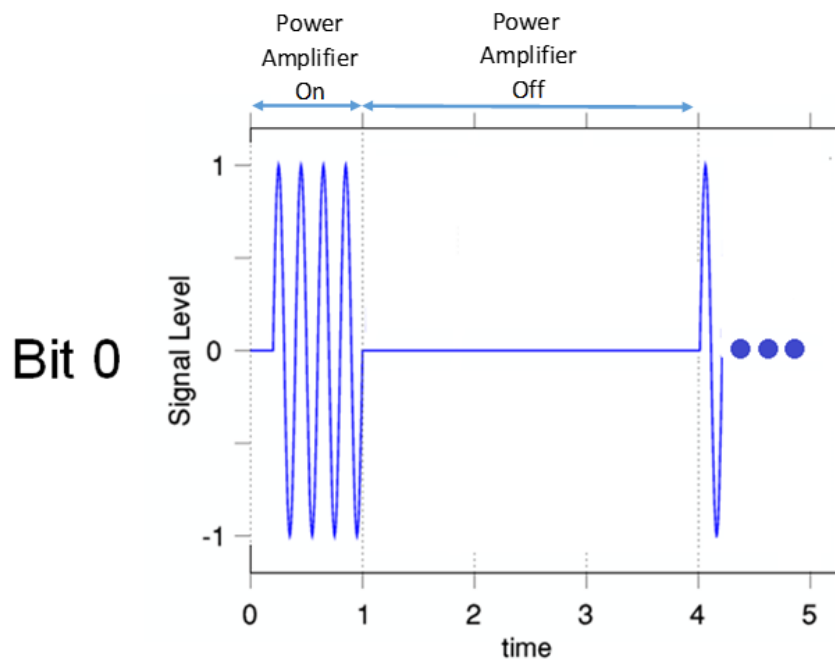



Figure 6-5: Proposed data encoding using OOK.




Figure 6-6: Schematic of circuit used to receive data from external device.




Figure 6-7: Example of bits 1 and 0 being transmitted to program the ASIC. Top and bottom graphs are the clock that controls the class-E power amplifier and the output from the Schmitt trigger, respectively.

105 μs and 210 μs) being transmitted. The top and bottom graphs are the clock that controls the class-E power amplifier and the output from the Schmitt trigger, respectively. According to the figure, the envelope detector removes the 5 MHz carrier frequency and the Schmitt trigger removes any glitches during the transitions.

6.2.3 Transmitting Data to External Device

In Mode 1, the ASIC samples the EEG channels at 1.28 kHz and each sample contains 10 bits. Therefore, the ASIC needs to transmit at least 102.4 kbit/s, assuming no overhead bits are included. In our case, we included one header byte (hex value AA) with every packet of data, which consists of one sample from each channel. As a result, our data rate is actually 112.64 kbits/s. For implants that use an inductive link, one




Figure 6-8: Schematic of circuit used to transmit data to external device.

of the most power efficient ways of transmitting data is through load modulation of the internal coil [49]. The idea, illustrated in Figure 6-8, is to change the coupling between external and internal coils by turning transistor M1 on and off. When M1 is on, the internal coil couples to the external coil. When M1 is off, it does not.

The coupling between the two coils creates a load on the external coil. This load was derived in [62] and is given by Equation 6.2, where $Z_{reflected}$ is the reflected impedance from the internal coil into the external one, k is the coupling between the coils, Q_1 and Q_2 are the quality factors of the internal and external coils, and R_2 is the series resistance of the external coil. According to Equation 6.2, if the transistor in Figure 6-8 is off, then Q_1 equals zero and the internal coil doesn't load the external coil. Load modulation affects the voltage amplitude at the external coil [62], so in order to decode the data, the external device needs an envelope detector similar to the one described in Section 6.2.2. Figure 6-9 illustrates the amplitude of the external coil changing due to load modulation and the decoding of the envelope.

$$Z_{reflected} = k^2 Q_1 Q_2 R_2 \quad (6.2)$$

6.2.4 Frequency of Operation

The inductive link carrier frequency was set to 5 MHz for two reasons. First, an FDA-approved cochlear implant, Nucleus Freedom, used this RF frequency to transmit 100




Figure 6-9: Top and bottom graphs are the voltage at the external coil and the output of the envelope detector, respectively

mW into the implant [63]. We are transmitting approximately 10 mW of power at the same frequency, so any required approval process to implant the device would be simplified since the FDA already approved a similar one. Second, to ease the data decoding by the envelope detector in the external device. The envelope detector has to filter the carrier frequency while not affecting the signal's envelope. In our case, we are transmitting 112.64 kbit/second, which means that the envelope changes around this rate. To ease the detection, the carrier frequency needs to be much higher than this value.

The carrier frequency sets the values of the resonant capacitor and inductor according to Equation 6.3. The higher the carrier frequency, the lower the inductor and capacitor values. Smaller inductors result in lower coupling coefficient between the external and internal coil, which degrades power transfer efficiency [49]. Smaller capacitors are more affected by board parasitics. A good trade-off between ease of envelope detection and small component values was around 5 MHz. With this carrier frequency, the capacitor and inductor used were 280 pF and 3.6 μ H, respectively.

$$\omega_{carrier} = \frac{1}{\sqrt{LC}} \quad (6.3)$$

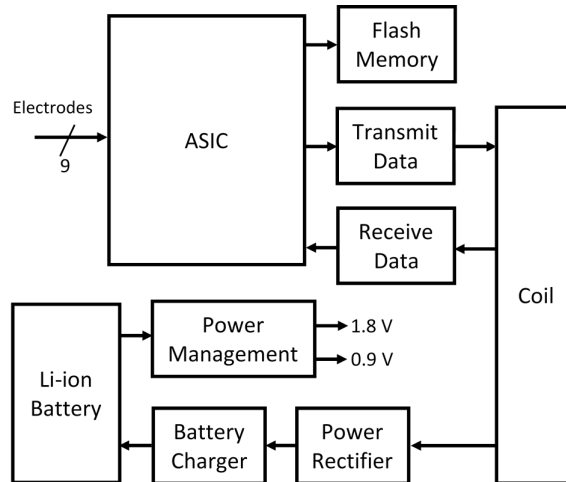


Figure 6-10: Block diagram of the implant.

6.3 Implant

The implant's block diagram is presented in Figure 6-10. The next subsections will describe its power management and battery charger in detail.

6.3.1 Power Management

The power management of the implant must generate three voltage domains from a 3.6 V lithium-ion battery: 1.8 V, 0.9 V for analog, and 0.9 V for digital. The analog and digital blocks inside the ASIC run at 0.9 V, so in order to minimize supply noise in the analog blocks, we decided to separate them.

The ASIC consumes approximately $31 \mu\text{W}$ in Mode 1 and $17 \mu\text{W}$ in Mode 2, so we must consider the quiescent current of the linear and switching regulators to make sure they don't consume more power than the ASIC. The chosen switching regulator was the TI TPS62736 since its quiescent current is only 350 nA and its expected efficiency is approximately 87%. As for the linear regulator, we couldn't find one that had very low quiescent current and could generate 0.9 V from an input around 1 V input. Most linear regulators that would work with approximately 1 V input had a quiescent current on the order of $20 \mu\text{A}$ - $100 \mu\text{A}$, which would cause the system to be power inefficient. Our best option was the Linear LT3009, which consumes $3 \mu\text{A}$, but has a minimum input voltage of 1.6 V.

The 1.8 V is used by the flash memory and some level shifters in the ASIC, allowing for a few mV of tolerable switching noise. As a consequence, we considered the two power management options presented in Figure 6-11. In option 1, the 1.8 V needs to supply approximately 5.8 μA . Both 0.9 V supplies consume about 17 μW in Mode 2 and each linear regulator has a quiescent current of 3 μA . According to the TPS62736 datasheet, the efficiency to supply 5.8 μA when the output is 1.8 V is approximately 77% and the efficiency to supply 25 μA when the output is 1.6 V is approximately 83%. Therefore, the total power consumption from the battery is 61.7 μW . In option 2, the 1.8 V supply all the current, which is 30.8 μA . According to the datasheet, the efficiency in this case is approximately 87%. Therefore, the total power consumption from the battery is 63.7 μW , resulting in option 2 being 3% more inefficient than option 1. However, since this system will be implanted, its area is also a major concern. Option 2 has four fewer parts than option 1: one switching regulator, two capacitors and one inductor. Since the power efficiency difference is only 3%, we decided to go with option 2.

6.3.2 Battery Charger

We included a battery in the implant to eliminate the requirement of wearing the external device continuously in Mode 2. Figure 6-12 shows a picture of the medical grade battery from Quallion (QL0003I). It has a 3 mAh capacity, measures 2.9 mm in diameter by 11.8 mm in height, and weights 0.2 grams. The system is expected to consume 63.7 μW , so it can last up to seven days on a single charge.

In Mode 2, the goal is to have the user wear the external device once a day for about 15 minutes to get the seizure data from that day and also to recharge the battery. In 24 hours, the system consumes about 14% of the battery. So, in order to recharge that amount in 15 minutes, we need to charge the battery with about 1.7 mA. The battery charger used in the system is Maxim's MAX8808 because it can charge with a current as low as 1 mA and only consumes 1 μA from the battery when not charging. The MAX8808 part works with an input voltage as low as 4.25 V, so with a recharge current of 1.7 mA, the external device must transfer at least 7.2 mW

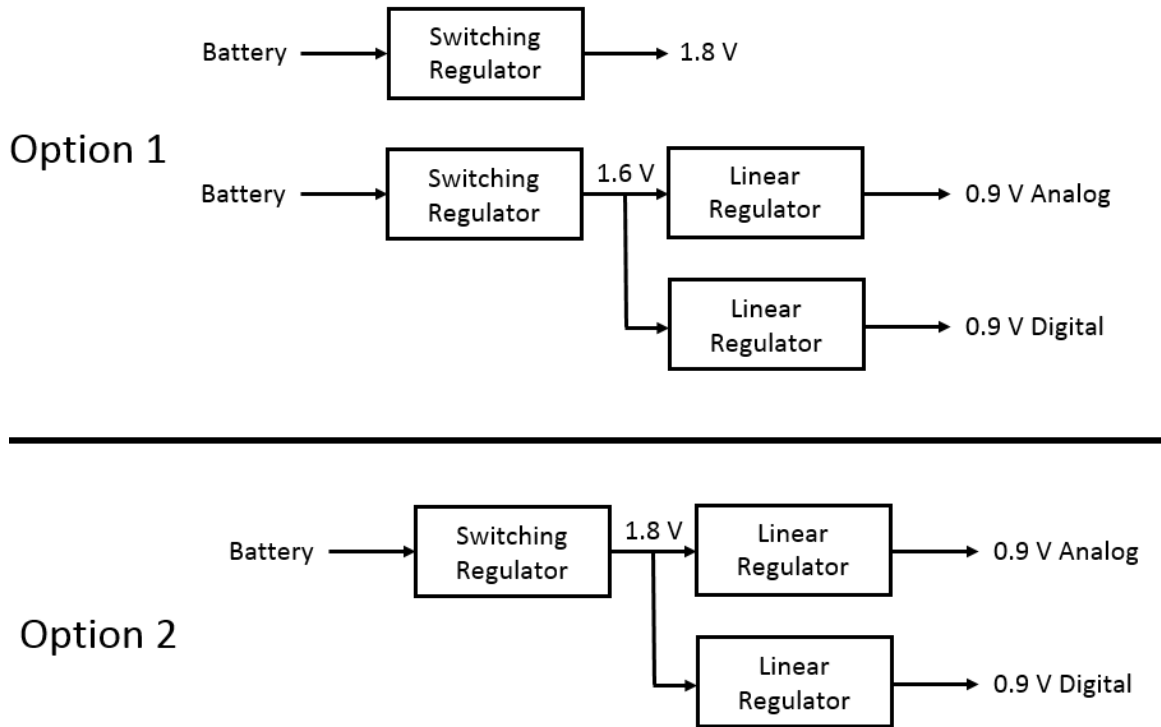


Figure 6-11: Two options considered for the power management.



Figure 6-12: 3 mAh implantable medical battery from Quallion used in the implant.




Figure 6-13: Front and back of the implant PCB.

to meet our goal of recharging in 15 minutes.

6.3.3 PCB

Figure 6-13 shows the front and back of the implant PCB. The battery is located inside the battery holder, which was created using a 3D printer. The ASIC was packaged in a 144-pin LQFP and it takes most of the front of the PCB while the back is devoted to the power management and passives to rectify the power coming through the inductive link. The PCB measures 31 mm by 31 mm.

The coil was made using litz-wire to minimize skin effect [64]. The litz wire was wrapped around a circular disc, created using a 3D printer. The coil has a magnet in the center to guarantee that it is aligned with the external device's coil.




Figure 6-14: Block diagram of the external device.

6.4 External Device

The block diagram of the external device, designed by Jason Yang, is shown in Figure 6-14. The microcontroller, a MSP430 from TI, is responsible for sending the clock that controls the class-E power amplifier, decoding the data received from the envelope detector, and using a Bluetooth module to transmit the EEG data to a computer. The Bluetooth module (LM780) contains all the required components for data transmission (antenna and IC to handle the protocol). The module communicates with the microcontroller through a simple UART interface. In our case, we traded simplicity of the module with power. The module consumes 29 mA when transmitting data and requires a minimum supply of 3 V.

6.4.1 PCB

Figure 6-15 shows the front and back of the external PCB, respectively. The PCB measures 42 mm by 41 mm. Most of the area is taken by the Bluetooth module.

6.4.2 Power Consumption

Table 6.1 lists the power consumption of the main blocks in the external device. The Bluetooth module consumes approximately 65% of the overall power.

To power the external device, we used a Tenergy ER14505 primary lithium-ion battery. It provides an output voltage of 3.6 V and has a capacity of 2.4 Ah. The size is the same of an AA battery: 14.5 mm diameter by 50.5 mm height. The external device




Figure 6-15: Front and back of the external PCB.

Table 6.1: Power consumption of main blocks in the external device

Block	Power Consumption (mW)
Power Amplifier	42
Microcontroller	7.5
Bluetooth	87

consumes approximately 205 mW (including inefficiencies due to power management), so the battery lasts for about forty hours. When choosing the battery, there was a trade off between how long the external device would last and the size of the battery. We chose the Tenery ER14505 because the system would last for more than two days and its size wouldn't cause the external device's form factor to increase significantly.




Figure 6-16: External device housing with the lid opened.

6.5 External Device Housing

Figure 6-16 shows a picture of the external device housing with the top lid opened. The housing will be attached to the pig's ear using an ear tag. The size of the housing is 61 mm (wide) x 65 mm (length) x 23 mm (height) and it weighs 69 grams. Two wires from the coil come off from the package so that the external coil can be placed on top of the implant coil as illustrated by Figure 6-18. The coil has a magnet in the center to guarantee alignment with the internal coil.

Figure 6-17 shows a picture of the external device housing with the top lid closed. The lid is secured by four magnets in order to ease its opening to replace the battery.




Figure 6-17: External device housing with the lid closed.

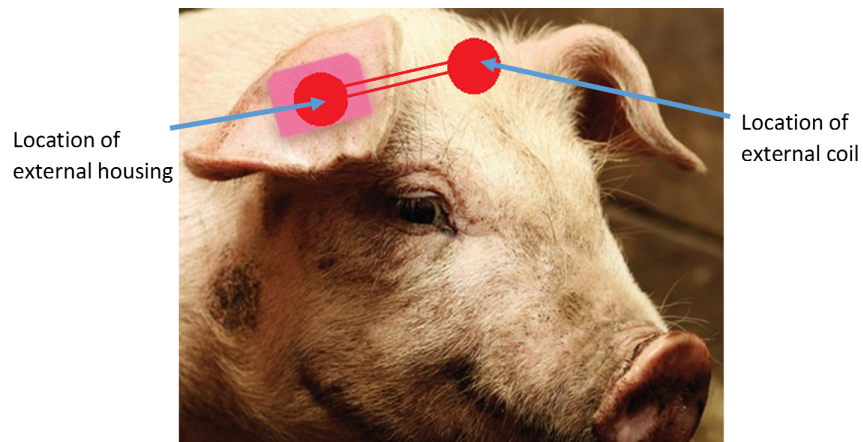


Figure 6-18: Location of the external device and coil on the swine. The distance between the external housing and coil is approximately 10 cm [65].

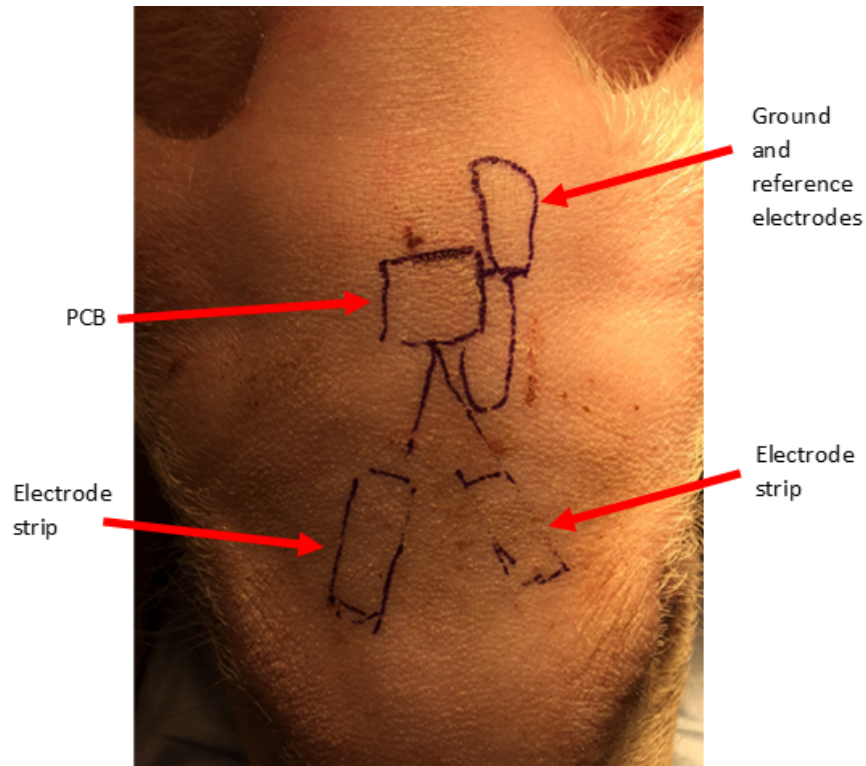


Figure 6-19: Pre-surgical planning.

6.6 Animal Testing

The animal tests were performed at Massachusetts General Hospital (MGH) using a swine model. The protocol (2014N000267) was approved by MGH's Institutional Animal Care and Use Committee (IACUC). All animal experiments were conducted according to IACUC and the USDA guidelines.

Figure 6-19 shows the shaved head of the swine with the pre-surgical planning of where the PCB and electrode strips were implanted. The PCB was placed roughly at the center of the head. The reference and ground electrodes were placed behind the PCB (touching the skull) while the recording electrodes were positioned on top of the temporal regions.

Figures 6-20 and 6-21 are pictures of the first surgery on a swine. An incision of approximately 6 cm was made on the scalp to implant the device. The hardest part of the surgical procedure was to place the electrodes since they did not have a guide wire. Nevertheless, the surgery went well and the device was successfully implanted.





Figure 6-20: Device being implanted in swine.



Figure 6-21: Device is implanted in swine.

Front



Back




Figure 6-22: Implant coated with epoxy.

Even though the surgery was successful, the first animal test did not work as planned. The PCB and coil were coated with silicone (MG Chemicals Silicone Conformal Coating Part Number 422) to prevent blood from shorting the system. The coating wasn't sufficient and some blood got into the coil affecting its inductance, which caused the inductive link to fail. Based on this trial, we decided to coat the system with epoxy (MG Chemicals 832B Epoxy Encapsulating and Potting Compound) since it gives a better protection than the silicone coating. Figure 6-22 shows the implant coated with epoxy.



Figure 6-23: Device implanted in swine and external device glued to ear tag.

The implant coated with epoxy was tested on another swine. Figure 6-23 shows a picture of the implanted device. This time a U-shaped incision was made to simplify the electrodes placement on top of the skull. The external coil, which was also coated with epoxy, was held in place by a magnet and the external device was glued to the swine's ear tag.

During this experiment we were able to record about thirty minutes of EEG. Figure 6-24 depicts a ten-second EEG segment recorded from the swine using eight channels. The data contains some slow activity – oscillations at approximately 3 Hz (Delta band) of approximately $100 \mu\text{V}$ amplitude in the temporal region. The slow activity was probably caused by the anesthesia since the swine hadn't fully recovered from it when the data was recorded. According to the FFT, shown in Figure 6-25, the ten-second segment actually has two peaks: one at 3 Hz and another one at 8 Hz.

We were not able to record for a longer period of time because as the swine started to wake up from anesthesia, it would move its head frequently, causing the external




Figure 6-24: Ten-second EEG segment recorded from the swine.




Figure 6-25: FFT of the ten-second EEG segment recorded from the swine. First peak at 3 Hz and second one at 8 Hz




Figure 6-26: External coil with holes for suture.

coil to fall off. This resulted in the external device not receiving the data from the implant.

To solve this problem, we decided to do another animal test where the external coil would be sutured to the swine's head. The coil was modified, as shown in Figure 6-26, to include holes for the suture. Figures 6-27 and 6-28 show the implant being placed on the swine and the coil sutured to the swine's head, respectively. In this experiment, we also moved the external device from the swine's ear to its back to make it more secure. The swine wore a jacket, manufactured by Lomir Biomedical (item number SS J2YA), that contained a velcro pouch located on the midline of shoulder blades. The pouch was 10 cm long by 8.5 cm wide and 3.5 cm deep. Figure 6-29 shows the swine wearing the jacket.

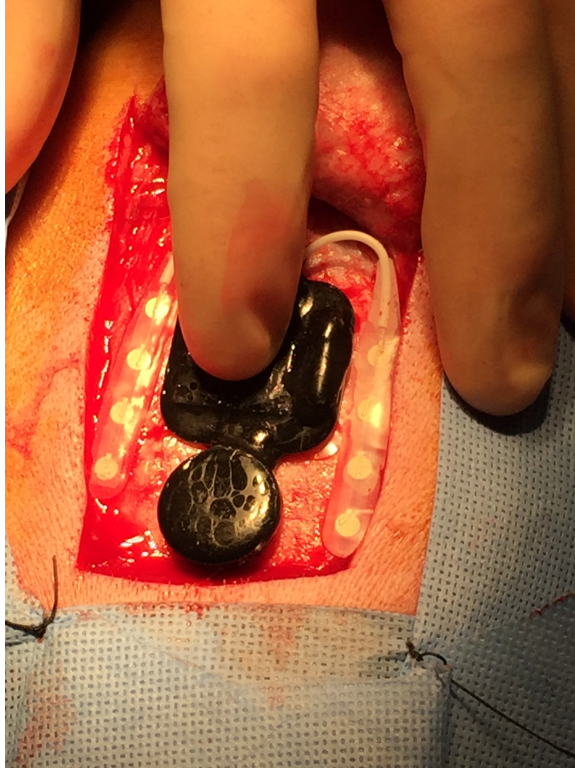


Figure 6-27: Implant being placed on the swine.

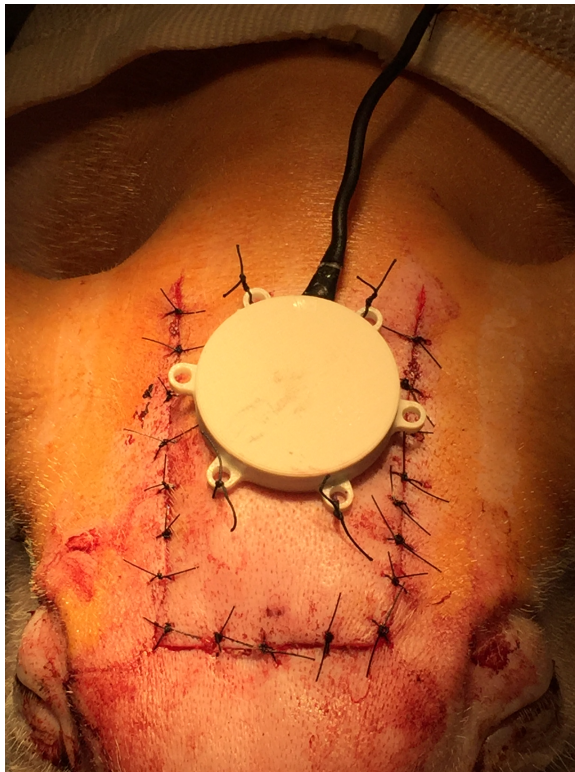


Figure 6-28: Coil sutured to the swine's head.




Figure 6-29: Swine wearing the jacket that holds external device.

This time we were able to record data during two days before the wire that connects the external device to the coil broke due to the swine's movement. The recorded data contained a lot of artifacts because the electrodes could not make a perfect contact to the skull due to the blood caused by the incision. Figure 6-30 shows a ten-second EEG segment recorded from the temporal lobe of the swine as he was recovering from anesthesia. The data is plotted in Sequential montage, where each channel represents the difference between neighboring channels, to reduce the amount of noise. Even though the data is plotted in Sequential montage, only the bottom channel (ch7-ch8) is not dominated by noise. According to the FFT of the channel 'ch7-ch8', shown in Figure 6-31, the data contains oscillations at approximately 1.6 Hz (Delta band), 3.3 Hz (Delta band), and 8.1 Hz (Alpha band). The Delta and Alpha amplitudes are approximately $100 \mu\text{V}$ and $50 \mu\text{V}$, respectively. They are present since the animal was waking up from anesthesia.




Figure 6-30: Ten-second EEG segment recorded from swine. Data is plotted in Sequential montage.




Figure 6-31: FFT of the channel 'ch7-ch8'. Main peak is located at 1.6 Hz.

To get reliable recordings from the swine, its body needs at least a few days before it can absorb all the blood from the incision and any swelling is gone. Unfortunately, due to a wire problem, this system only lasted two days in this animal study and the swine's body did not have enough time to absorb all the blood. To make the external device more reliable, some sort of strain relief must be placed on the wire that connects the external device to the coil, preventing the wire from breaking when the swine moves.

Despite the issues with the external device, the system was successful at recording EEG data from the swine. We saw great improvement in data collection between the second and third animal studies – recording time increased from 30 minutes to two days. Most of the problems we had were due to the fact that we were not able to restrain the animal's movement. Future versions of this system can be made more reliable by including a strain relief for wires and reducing the size of the device.

6.7 Summary

This chapter presented the design of a subdermal implantable EEG recorder and seizure detector. The system is comprised of two parts: an implant that is placed between the skin and skull behind the right or left ear, and an external device located on top of the skin. An inductive link consisting of a coil in the implant and another one at the external device is responsible for data and power transfer. The system was tested using a swine model at MGH and we were able to record EEG during two days before the wire that connected the coil to the external device broke due to swine movement. The animal test proved that the system can record EEG from electrodes placed on the skull and transmit the data through the inductive link to the external device, which transmits the data to a laptop using Bluetooth. This was a major milestone for creating a long-term, subdermal implantable EEG recorder and seizure detector that can be used on patients with epilepsy.

Chapter 7

Conclusion

This thesis presented the development of two ambulatory EEG systems. The first is a one-channel, wearable behind-the-ear EEG recording device, and the second is an eight-channel, subdermal implantable EEG recorder and seizure detector. While the wearable system can only be used to track the number of seizures a patient has, the implantable device can also be used for diagnosis. These new systems are one step forward in the design of smaller and better ambulatory EEG systems. They offer patients with epilepsy a more convenient way of monitoring their seizures and give doctors more accurate information for diagnosis and treatment options.

7.1 Summary of Contributions

In summary, this work presents several contributions:

- A proof-of-concept, one-channel, wearable EEG recorder for long-term monitoring of patients with epilepsy was designed. This device uses an iPod Touch to continuously transmit data to a secure server at the hospital, giving doctors access to the data in real time. Clinical testing was performed on four patients at MGH in collaboration with Dr. Syd Cash. The data recorded with our device was comparable to hospital's recording.
- A novel low-power on-chip seizure detection algorithm was developed. Seizure

detection power is reduced by at least a factor of ten when compared to recently published algorithms. This algorithm is ideal for an ambulatory system that is used to track the number of seizures from a patient to determine medication efficacy.

- A low-power, eight-channel ASIC to record EEG up to five times the typical bandwidth was designed. The higher bandwidth enables recording of HFO, which can potentially improve diagnosis. The analog front end is comparable to state-of-the-art designs. Seizure detection was also included in the ASIC and tested using a PhysioNet dataset.
- A subdermal implantable system using the ASIC was created. Data and power are transmitted to/from the ASIC through an inductive link to/from an external device. Animal testing using a swine model was performed at MGH.

7.2 Future Work

Several areas are worth exploring to improve the current system:

- The goal for the wearable EEG system is to send patients home with it to track the number of seizures they have. However, first the iPod app needs to be upgraded. The iPod communicates with our device through Low Energy Bluetooth and currently, if that connection is lost, the only way to reconnect is by pressing buttons on the app. Ideally the app should reconnect automatically to the device without any interference from the patient. The connection loss can be caused by several factors, such as separating the iPod from our device by more than 20 meters, which might happen frequently. Therefore, updating the app is crucial before sending patients home with the device.
- The current IC is packaged in a 144-pin LQFP, which measures 20 mm by 20 mm. The IC only requires about 50 pins to be bonded to the pads, however, due to the pad placement, we had to use a 144-pin package since the bonding

diagram for a 64-pin package was not feasible. By using a smaller package or bonding the die directly to the PCB, the size of the internal board can be reduced significantly, which simplifies the surgical procedure to implant the system underneath the scalp.

- The ASIC detects seizures by measuring the energy on four frequency bands, but they are only available internal to the chip, not as an output. There are other applications, such as sleep tracking, that monitor the energy of those frequency bands. Therefore, if the energy values are also made available as an output, the number of applications that can use this ASIC increases.
- The only linear regulator (LT3009) that has a low quiescent current comparable to the implant's PCB requires a minimum input voltage of 1.6 V. This degraded our power consumption significantly since the linear regulator produces an output voltage of 0.9 V. Therefore, if the linear regulators were designed on chip, then the overall power efficiency would increase significantly.
- The external device used a Bluetooth module that consumes 65% of overall power. If the module is changed to a Bluetooth Low Energy, power savings on the order of 50% can be achieved, resulting in a smaller battery and package.
- The chosen charger IC (MAX8808) doesn't work when the battery voltage drops below 2 V. Therefore, if the implant's battery falls below 2 V, it can no longer be charged, resulting in a permanent loss of the implant. Since the system is coated with epoxy, it cannot be fixed even if removed from the swine. The MAX8808 was chosen as the charger because it was the only IC found that could charge a battery with a current as small as 1 mA. One way to fix this issue is by finding another charger IC that can charge with currents as low as 1 mA and also work when the battery voltage is close to 0 V. Another way is to bypass the charger IC when the battery voltage is too small by adding two resistors and one diode, as illustrated by Figure 7-1. Resistors R1 and R2 can be chosen to have a value around 100 k Ω to minimize the amount of current




Figure 7-1: Circuit to bypass battery charger.

drawn from the system. The input voltage to the charger is limited to 5 V, so if R1 and R2 have the same value, then when the battery reaches 2.5 V, the bypass circuit turns off. The diode D2 was included to prevent any discharge from the battery through R2.

- The external device used to record data from the implant was the limiting factor on the animal study. On the second day of recording, the wire that connected the coil to the external device broke and interrupted the data collection. Future versions of the external device could be made more reliable by using a thicker wire to connect the coil to the external device, including a strain relief for the wire, and reducing coil and external device's sizes.

Bibliography

- [1] L. D. Iasemidis, “Epileptic seizure prediction and control,” *Biomedical Engineering, IEEE Transactions on*, pp. 549–558, May 2003.
- [2] G. L. Krauss, R. S. Fisher, and P. W. Kaplan, *The John Hopkins Atlas of Digital EEG*. The John Hopkins University Press, 2011.
- [3] R. Senelick, “Epilepsy seizure types and symptoms @ONLINE,” May 2014. URL: <http://www.webmd.com/epilepsy/guide/types-of-seizures-their-symptoms>.
- [4] S. R. Haut, C. B. Hall, J. Masur, and R. B. Lipton, “Seizure occurrence precipitants and prediction,” *Neurology*, vol. 69, pp. 1905–1910, November 2007.
- [5] C. Carmen, “Study of EEGs in the ER improves seizure diagnosis, care @ONLINE,” May 2014. URL: <http://epilepsy.com/blog/study-eegs-er-improves-seizure-diagnosis-care/>.
- [6] S. R. Fisher, D. E. Blum, B. DiVentura, J. Vannest, J. D. Hixson, and J. A. French, “Seizure diaries for clinical research and practice: Limitations and future prospects,” *Epilepsy and Behavior*, pp. 304–310, May 2012.
- [7] NYU Langone Medical Center, “Ambulatory EEG @ONLINE,” March 2014. URL: <http://epilepsy.med.nyu.edu/diagnosis-treatment/eeg/ambulatory-eeg>.
- [8] Sacred Heart HealthCare System, “Ambulatory electroencephalography (EEG) @ONLINE,” March 2014. URL: <http://www.shh.org/hospital-services/neurodiagnostic-testing/ambulatory.asp>.
- [9] B. McLaughlin, L. Mariano, S. Prakash, A. Kindle, A. Czarnecki, M. Modarres, A. Rotenberg, T. Loddenkemper, A. Shoeb, and S. Schachter, “An electroencephalographic recording platform for real-time seizure detection,” in *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pp. 875–878, September 2012.
- [10] Boston Children’s Hospital, “Family education sheet @ONLINE,” March 2014. URL: <http://www.childrenshospital.org/media/Centers%20and%20Services/Departments%20and%20Divisions/Department%20of%20Neurology/AmbulatoryEEG.ashx>.

- [11] B. Do Valle, S. Cash, and C. Sodini, “Wireless behind-the-ear EEG recording device with wireless interface to a mobile device (iPhone/iPod touch),” in *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE*, pp. 5952–5955, August 2014.
- [12] A. Casson, D. Yates, S. Smith, J. Duncan, and E. Rodriguez-Villegas, “Wearable electroencephalography,” *Engineering in Medicine and Biology Magazine, IEEE*, vol. 29, pp. 44–56, May 2010.
- [13] Neuropace, “Brief statement @ONLINE,” May 2015. URL: <http://www.neuropace.com/product/pdfs/Brief-Statement.pdf>.
- [14] Neuropace, “Patient education @ONLINE,” May 2015. URL: <http://neuropace.articulate-online.com>.
- [15] Ontario Spinal Institute, “@ONLINE,” June 2015. URL: <http://www.ontsi.com/services.htm>.
- [16] B. Lee, M. Zubair, Y. Marquez, D. Lee, L. Kalayjian, C. Heck, and C. Liu, “A single center experience with the neuropace rns system: A review of techniques and potential problems,” *World Neurosurgery*, November 2014.
- [17] A. H. Shoeb, *Application of Machine Learning to Epileptic Seizure Onset Detection and Treatment*. PhD thesis, Massachusetts Institute of Technology, 2009.
- [18] M. F. Bear, B. W. Connors, and M. A. Paradiso, *Neuroscience Exploring the Brain*. Lippincott Williams and Wilkins, 2007.
- [19] J. Jacobs, M. Zijlmans, R. Zelman, C.-É. Chatillon, J. Hall, A. Olivier, F. Dubeau, and J. Gotman, “High-frequency electroencephalographic oscillations correlate with outcome of epilepsy surgery,” *Annals of Neurology*, vol. 67, pp. 209–220, February 2010.
- [20] J. R. Hughes, *EEG in Clinical Practice*. Butterworth-Heinemann, 2 ed., 1994.
- [21] G. H. Klem, H. O. Lüders, H. Jasper, and C. Elger, “The ten-twenty electrode system of the international federation,” *Electroencephalography and Clinical Neurophysiology*, vol. 52, no. suppl., p. 3, 1999.
- [22] M. Teplan, “Fundamentals of EEG measurement,” *Measurement Science Review*, vol. 2, no. 2, pp. 1–11, 2002.
- [23] S. D. Muthukumaraswamy, “High-frequency brain activity and muscle artifacts in MEG/EEG: a review and recommendations,” *Frontiers in Human Neuroscience*, vol. 7, April 2013.
- [24] A. Casson, D. Yates, S. Smith, J. Duncan, and E. Rodriguez-Villegas, “Wearable electroencephalography,” *Engineering in Medicine and Biology Magazine, IEEE*, vol. 29, pp. 44–56, May 2010.

- [25] American Clinical Neurophysiology Society, “Guideline twelve: Guidelines for long-term monitoring for epilepsy @ONLINE,” March 2008. URL: <https://www.acns.org/pdf/guidelines/Guideline-12.pdf>.
- [26] E. Niedermayer and F. L. Silva, *Electroencephalography: Basic Principles, Clinical Applications and Related Fields*. Williams and Wilkins, 2005.
- [27] R. Yazicioglu, P. Merken, and C. Van Hoof, “Effect of electrode offset on the CMRR of the current balancing instrumentation amplifiers,” in *Research in Microelectronics and Electronics, 2005 PhD*, vol. 1, pp. 35–38 vol.1, July 2005.
- [28] K. Blinowska and P. Durka, *Electroencephalography (EEG) in Wiley Encyclopedia of Biomedical Engineering*. Wiley-Interscience, 2006.
- [29] A. Oppenheim and R. Schaffer, *Discrete-Time Signal Processing*. Prentice Hall, third ed., August 2009.
- [30] MVAP, “Huggable electrode @ONLINE,” April 2015. URL: <https://secure.mvapped.com/prod-HuggableElectrode-534.aspx>.
- [31] R. Powers, “Batteries for low power electronics,” *Proceedings of the IEEE*, vol. 83, pp. 687–693, April 1995.
- [32] M. Holmes, C. Dodrill, R. Wilkus, L. Ojemann, and G. Ojemann, “Is partial epilepsy progressive? Ten-year follow-up of EEG and neuropsychological changes in adults with partial seizures,” *Epilepsia*, vol. 39, pp. 1189–1193, November 1998.
- [33] A. Shoeb, D. Carlson, E. Panken, and T. Denison, “A micropower support vector machine based seizure detection architecture for embedded medical devices,” in *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, pp. 4202–4205, September 2009.
- [34] J. Yoo, L. Yan, D. El-Damak, M. Bin Altaf, A. Shoeb, H.-J. Yoo, and A. Chandrakasan, “An 8-channel scalable EEG acquisition SoC with fully integrated patient-specific seizure classification and recording processor,” in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, pp. 292–294, February 2012.
- [35] S. Asha, C. Sudalaimani, P. Devanand, T. Thomas, and S. Sudhamony, “Automated seizure detection from multichannel EEG signals using support vector machine and artificial neural networks,” in *Automation, Computing, Communication, Control and Compressed Sensing (iMac4s), 2013 International Multi-Conference on*, pp. 558–563, March 2013.
- [36] J. Gotman, “High frequency oscillations: The new EEG frontier?,” *Epilepsia*, vol. 51, pp. 63–65, February 2010.

- [37] J. Jirsch, E. Urrestarazu, P. LeVan, A. Olivier, F. Dubeau, and J. Gotman, “High-frequency oscillations during human focal seizures,” *Brain*, vol. 129, pp. 1593–1608, April 2006.
- [38] L. Andrade-Valenca, F. Dubeau, F. Mari, R. Zelmann, and J. Gotman, “Interictal scalp fast oscillations as a marker of the seizure onset zone,” *Neurology*, vol. 77, no. 6, pp. 524–531, 2011.
- [39] R. Harrison and C. Charles, “A low-power low-noise CMOS amplifier for neural recording applications,” *Solid-State Circuits, IEEE Journal of*, vol. 38, pp. 958–965, June 2003.
- [40] M. Yip, *Ultra-low-power circuits and systems for wearable and implantable medical devices*. PhD thesis, Massachusetts Institute of Technology, 2013.
- [41] D. He, *A wearable heart monitor at the ear using ballistocardiogram (BCG) and electrocardiogram (ECG) with a nanowatt ECG heartbeat detection circuit*. PhD thesis, Massachusetts Institute of Technology, 2013.
- [42] T. Denison, K. Consoer, W. Santa, A.-T. Avestruz, J. Cooley, and A. Kelly, “A 2 uW 100 nV/rtHz chopper-stabilized instrumentation amplifier for chronic measurement of neural field potentials,” *Solid-State Circuits, IEEE Journal of*, vol. 42, pp. 2934–2945, December 2007.
- [43] N. Verma, A. Shoeb, J. Bohorquez, J. Dawson, J. Guttag, and A. Chandrakasan, “A micro-power EEG acquisition SoC with integrated feature extraction processor for a chronic seizure detection system,” *Solid-State Circuits, IEEE Journal of*, vol. 45, pp. 804–816, April 2010.
- [44] D. Johns and K. Martin, *Analog Integrated Circuit Design*. Wiley, New York, NY, 1996.
- [45] R. Sarpeshkar, *Ultra Low Power Bioelectronics: Fundamentals, Biomedical Applications, and Bio-Inspired Systems*. Cambridge University Press, 2010.
- [46] S. Arfin, S. Mandal, and R. Sarpeshkar, “Dynamic-range analysis and maximization of micropower Gm-C bandpass filters by adaptive biasing,” in *Circuits and Systems, 2009. ISCAS 2009. IEEE International Symposium on*, pp. 2954–2957, May 2009.
- [47] J. Bohorquez, M. Yip, A. Chandrakasan, and J. Dawson, “A biomedical sensor interface with a sinc filter and interference cancellation,” *Solid-State Circuits, IEEE Journal of*, vol. 46, pp. 746–756, April 2011.
- [48] B. Ahuja, “Implementation of active distributed RC anti-aliasing/smoothing filters,” *Solid-State Circuits, IEEE Journal of*, vol. 17, pp. 1076–1080, December 1982.

- [49] F. Maloberti, *Analog Design for CMOS VLSI Systems*. Springer Science and Business Media, October 2001.
- [50] M. Yip, J. Bohorquez, and A. Chandrakasan, “A 0.6V 2.9 μ W mixed-signal front-end for ECG monitoring,” in *VLSI Circuits (VLSIC), 2012 Symposium on*, pp. 66–67, June 2012.
- [51] Mathworks, “Filter design HDL coder user’s guide,” *Mathworks*, September 2012.
- [52] R. Losada, “Digital filters with MATLAB,” *Mathworks*, May 2008.
- [53] A. Paidimarri, D. Griffith, A. Wang, A. Chandrakasan, and G. Burra, “A 120nW 18.5kHz RC oscillator with comparator offset cancellation for 0.25% temperature stability,” in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2013 IEEE International*, pp. 184–185, February 2013.
- [54] S. Mandal, S. Arfin, and R. Sarpeshkar, “Fast startup CMOS current references,” in *Circuits and Systems, 2006. ISCAS 2006. Proceedings. 2006 IEEE International Symposium on*, pp. 1–4, May 2006.
- [55] M. Steyaert and W. Sansen, “A micropower low-noise monolithic instrumentation amplifier for medical purposes,” *Solid-State Circuits, IEEE Journal of*, vol. 22, pp. 1163–1168, December 1987.
- [56] X. Zou, W.-S. Liew, L. Yao, and Y. Lian, “A 1V 22 μ W 32-channel implantable EEG recording IC,” in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International*, pp. 126–127, February 2010.
- [57] H. Liu, K.-T. Tang, J.-Y. Wu, and G. Wang, “A digitally trimmable low-noise low-power analog front-end for EEG signal acquisition,” in *Biomedical and Health Informatics (BHI), 2012 IEEE-EMBS International Conference on*, pp. 208–211, January 2012.
- [58] M. Altaf, J. Tillak, Y. Kifle, and J. Yoo, “A 1.83 μ J/classification nonlinear support-vector-machine-based patient-specific seizure classification SoC,” in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2013 IEEE International*, pp. 100–101, February 2013.
- [59] M. Shoaib, N. Jha, and N. Verma, “A compressed-domain processor for seizure detection to simultaneously reduce computation and communication energy,” in *Custom Integrated Circuits Conference (CICC), 2012 IEEE*, pp. 1–4, September 2012.
- [60] N. Sokal and A. Sokal, “Class E-A new class of high-efficiency tuned single-ended switching power amplifiers,” *Solid-State Circuits, IEEE Journal of*, vol. 10, pp. 168–176, June 1975.

- [61] T. Suetsugu and M. Kazimierczuk, “Design procedure of class-E amplifier for off-nominal operation at 50 percent duty ratio,” *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 53, pp. 1468–1476, July 2006.
- [62] S. Mandal and R. Sarpeshkar, “Power-efficient impedance-modulation wireless data links for biomedical implants,” *Biomedical Circuits and Systems, IEEE Transactions on*, vol. 2, pp. 301–315, December 2008.
- [63] F.-G. Zeng, S. Rebscher, W. Harrison, X. Sun, and H. Feng, “Cochlear implants: System design, integration, and evaluation,” *Biomedical Engineering, IEEE Reviews in*, vol. 1, pp. 115–142, January 2008.
- [64] A. Roskopf, E. Bar, and C. Joffe, “Influence of inner skin- and proximity effects on conduction in litz wires,” *Power Electronics, IEEE Transactions on*, vol. 29, pp. 5454–5461, October 2014.
- [65] C. Outfitters, “Ear tag @ONLINE,” April 2015. URL: <http://www.cckoutfitters.com/EARTAGS.html>.