

A System for Efficient Neural Stimulation with Energy Recovery

by

Shawn Kevin Kelly

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

[February 2004]
October 2003


© Shawn Kevin Kelly, MMIII. All rights reserved.

The author hereby grants to MIT permission to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part.

Author
Department of Electrical Engineering and Computer Science
October 6, 2003

Certified by 2003
John L. Wyatt, Jr.
Professor
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Department Committee on Graduate Theses



BARKER

A System for Efficient Neural Stimulation with Energy Recovery

by

Shawn Kevin Kelly

Submitted to the Department of Electrical Engineering and Computer Science
on October 6, 2003, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering

Abstract

An analog VLSI-based low-power neural tissue stimulator is presented as a part of the MIT and Massachusetts Eye and Ear Infirmary Retinal Implant Project to develop a prosthesis for restoring some useful vision to patients blinded by retinal degeneration. Such a prosthesis would receive image data from an external camera and electrically stimulate surviving ganglion nerve cells. However, power consumption for this type of implant is dominated by the tissue and electrode-tissue interface, and the current source stimulators generally used are inefficient, limiting battery life and generating potentially damaging temperature increases at the retinal surface. A stimulation system has been developed which delivers the required stimulus charge to the electrodes, but uses far less power than typical stimulators. A traditional current source uses output transistors to limit current, but those transistors can drop substantial voltage, and therefore cost power. The aim of this system is to generate a step-ramp voltage waveform which mimics the electrode voltage (modeled as a series resistance and capacitance) during constant current stimulation. This is implemented with a series of voltage steps, each step a separate power supply. Electrodes are switched through a series of steps, and each step is maintained at its prescribed voltage by a controlled synchronous rectifier, which charges the supply capacitor from a single AC secondary power coil.

This novel architecture uses less than half of the power used by an aggressively designed current source stimulator with the same voltage rails, and about one-third of the power consumed by typical stimulators used for this function.

Thesis Supervisor: John L. Wyatt, Jr.

Title: Professor

Acknowledgments

First, thanks to Professor John Wyatt for his hours of guidance, insight, and mentorship, both on this thesis project and in general. John has the uncanny ability to see through large complex problems and develop and convey deep insight and understanding. He is then willing and very able to seek the best possible advice if the solution to the problem lies outside the scope of his expertise. I have learned tremendously from him.

Special thanks to Professors Rahul Sarpeshkar and David Perreault for serving on my thesis committee. Rahul allowed me to work in his lab, where I absorbed a great deal of circuit design knowledge from him and his students. He is always energetic, insightful, and eager to help. Dave brought to the committee not only his knowledge of power systems, but also his incredible common sense. When the rest of us were thinking about design details too far down the road, Dave returned us to engineering fundamentals: building a system one step at a time.

Thanks to Dr. Joe Rizzo, co-director with John of the Retinal Implant Project, for his guidance of the project, his friendship, and his impeccable taste in restaurants.

Without fellow graduate student Luke Theogarajan, I would still be spinning my wheels in analog VLSI circuit design. He has been extremely helpful in my development as a designer, and has become a good friend.

Thanks to fellow Retinal Implant graduate students Mariana Markova, Erich Caulfield, Ken Roach, and Andrew Grumet, and researchers Ofer Ziv, Doug Shire, Ralph Jensen, and Milan Raj as well as Professor Sarpeshkar's graduate students Michael Baker, Chris Salthouse, Micah O'Halloran, Heemin Yang, Ji-Jon Sit, and all the rest for many hours and years of help, discussion, and laughs.

The Catalyst Foundation generously supplied five years of funding for this project, allowing me the freedom to work on this project.

Thanks to my freshman year roommates Dennis Burianek and Luis Ortiz and my friend Brian Schuler for leading the way straight through from undergraduate to PhD. Thanks to all of my friends for their years of support and occasional moments of relief from work.

Thanks to my Dad and the rest of my family for all the support through the long hours and long years. And a special thanks to my Mother, who always believed in me, but wasn't able to see me finish this degree.

Finally, thanks to my bride-to-be Elizabeth Daake for her love, support, and patience through all this time. It's *our* PhD.

Contents

1	Introduction	14
1.1	The Retinal Implant Project	14
1.1.1	Retinal Function	14
1.1.2	Outer Retinal Disease	15
1.1.3	The Retinal Prosthesis	16
1.2	Stimulation Electrodes	17
1.2.1	Commonly Used Electrode Materials	18
1.2.2	Electrical Model of the Electrode	20
1.2.3	Chronic Neural Overstimulation Damage	22
1.2.4	Stimulation Charge Thresholds	24
1.2.5	Electrode Power Requirements	26
1.3	Power and Stimulation System	26
1.3.1	Power Delivery Methods	27
1.3.2	Receiver Coil Placement	28
1.3.3	High-Level System Description	29
1.4	Related Work	30
1.5	Thesis Outline	31
2	Background Theory and Possible Architectures	33
2.1	Methods of Electrical Neural Stimulation	33
2.1.1	Constant Current Drive - The Optimal Charging Waveform	33
2.1.2	Electrode Power and Energy	35
2.1.3	The Canonical Stimulation Parameters	37

2.1.4	An Inefficient Constant Current Drive Method	37
2.1.5	Inter-Electrode Energy Recycling	39
2.1.6	The Optimal Drive Voltage Waveform	41
2.1.7	Optimal Capacitor Number and Voltage	46
2.1.8	Stimulation Power Comparison	51
2.1.9	Electrode Drive Tests in Saline	51
2.2	Power Recovery and Rectification	56
2.2.1	Possible Rectifier Architectures	56
2.2.2	Synchronous Rectifier	58
2.2.3	Magnetic Field Calculations	59
2.2.4	Magnetic Field-Induced Heating	63
2.2.5	Secondary Coil Parameters	66
2.2.6	Coil Resonance and Loading	70
3	Coupled Coils	74
3.1	Secondary Coil Design	74
3.1.1	Calculated Inductance	76
3.1.2	Calculated Resistance	76
3.2	Secondary Coil Construction	76
3.3	Primary Coil Design and Construction	79
3.4	Primary Coil Class E Driver	81
3.5	Coil Alignment Jig	84
3.6	Coil System Testing	84
4	Power System	90
4.1	Power Supply Rectifier	90
4.2	Capacitor Bank Rectifier	90
4.2.1	Timing-Based Rectifier	92
4.2.2	Voltage-Based Rectifier	92
4.2.3	Synchronous Rectifier Architecture	93
4.3	Capacitor Bank Rectifier Implementation	98

4.3.1	Voltage Reference	98
4.3.2	Clocked Comparators	101
4.3.3	Clock Circuitry	103
4.3.4	Continuous Comparator	106
4.3.5	Rectifier Timing	108
4.3.6	Controls	113
4.3.7	Gate Drive and Switches	114
4.3.8	Power On Reset	116
4.3.9	Layout Issues	117
4.4	Power System Testing	121
4.4.1	Voltage Reference Testing	123
4.4.2	Clock Extractor Testing	123
4.4.3	Synchronous Rectifier Comparator Testing	125
4.4.4	Synchronous Rectifier Timing Testing	126
4.4.5	Power On Reset Testing	127
4.4.6	Rectifier Chip Power	129
5	Electrode Drive	130
5.1	Pre-Programmed Test Patterns	130
5.1.1	Electrode Pattern Groups	131
5.1.2	Electrode Drive Groups	132
5.2	Digital Control System	133
5.2.1	Main State Machine	135
5.2.2	Electrode State Controls	140
5.2.3	Switch Driver Controls	141
5.3	Design Problems	143
5.4	Layout Issues	143
5.5	Electrode Drive System Testing	144
5.5.1	1.25 ms Clock	144
5.5.2	Electrode Patterns	145

5.5.3	Driving a Series RC Load	146
5.6	Conclusion	147
6	System Integration	148
6.1	Final VLSI Chip	148
6.2	Full System	150
6.3	Full System Testing	151
6.3.1	Driving Electrodes	153
6.3.2	System Power	155
7	Conclusions and Future Work	164
7.1	Thesis Conclusions	164
7.2	Future Work	167
7.2.1	Changes to This Implementation	167
7.2.2	Broader Changes	168
A	Mathematical Derivations and MATLAB Scripts	170
A.1	Proof of Optimal RC Charging Profile	170
A.2	Capacitor Bank Number and Voltage Calculations	173
A.2.1	3-Step Charging	173
A.2.2	4-Step Charging	176
A.2.3	5-Step Charging	180
A.3	Exact Magnetic Field Calculations	184
A.4	Geometric View of Coil Loading and Resonance	186

List of Figures

1-1	The Eye	15
1-2	The Retina	16
1-3	A Retinal Implant	17
1-4	Iridium Oxide Cyclic Voltammetry	20
1-5	Electrical Circuit Model of Iridium Oxide Electrode	22
1-6	Stimulation Current and Electrode Voltage Waveforms	23
1-7	Perception Threshold Charge Densities	25
1-8	Possible Retinal Implant Locations	29
1-9	System Block Diagram	30
2-1	Stimulation Current and Electrode Voltage Waveforms	34
2-2	Electrode Power Waveform	35
2-3	Back Voltage Waveforms for 2 Special Cases	36
2-4	Dual Supply Current Source Drive	38
2-5	Inter-Electrode Energy Recycling	39
2-6	Electrode-to-Electrode Recycling - Hydraulic Analogy	40
2-7	Ramping Capacitor Drive	43
2-8	Voltage and Current Waveforms for Ramping Capacitor Drive	43
2-9	Multiple Capacitor Bank	45
2-10	Voltage and Current Waveforms for Multiple Capacitor Bank	45
2-11	Three-Step Voltage Waveforms (Calculated)	48
2-12	Three-Step Current Waveforms (Calculated)	48
2-13	Four-Step Voltage Waveforms (Calculated)	49

2-14	Four-Step Current Waveforms (Calculated)	49
2-15	Five-Step Voltage Waveforms (Calculated)	50
2-16	Five-Step Current Waveforms (Calculated)	50
2-17	Test Circuit for Electrode Drive Tests in Saline	52
2-18	Test System for Electrode Drive Tests in Saline	52
2-19	Array with 40 Electrodes	53
2-20	Four-Step Voltage Waveforms (Measured)	54
2-21	Four-Step Current Waveforms (Measured)	54
2-22	Step-Ramp Voltage Waveforms (Measured)	55
2-23	Step-Ramp Current Waveforms (Measured)	55
2-24	Full Bridge Rectifier	56
2-25	Voltage Doubling Rectifier	57
2-26	Dual Half-Wave Rectifier	57
2-27	Charging a Negative-Voltage Capacitor from the AC Voltage	58
2-28	Synchronous Rectifier With Comparator and One-Shot	59
2-29	Magnetic Field Intensity vs. Axial Displacement From Primary	60
2-30	Coil Alignment (L) and Angle Calculations (R)	61
2-31	Magnetic Field Intensity vs. Radial Displacement From Center of Secondary	63
2-32	Magnetic Fields in the Eye	64
2-33	Calculating the Power Dissipated by the Electric field	65
2-34	Definitions for power volume integral	65
2-35	Secondary Coil Cross-Section	67
2-36	Simplified Coil Model	68
2-37	Impedance to Current Mapping	70
2-38	Minimum Power and Voltage Geometry	71
2-39	Geometric View of Coil Loading	72
3-1	Coil Cross-Section	75
3-2	The Secondary Coil	77

3-3	Secondary Coil Windings	77
3-4	Secondary Coil Impedance Magnitude and Phase	78
3-5	Secondary Coil Resistance and Inductance	79
3-6	The Primary Coil	80
3-7	Class E Driver	81
3-8	Class E Timing Controller	82
3-9	Class E Driver Circuit	83
3-10	Coil Alignment Jig	85
3-11	Photograph of the Coil Alignment Jig	85
3-12	Coil Coupling Test Circuit	85
3-13	Magnetic Field at Center of Primary	87
3-14	Rectified Power Supply Voltage vs. Class E Choke Supply	87
3-15	Equivalent Rectifier Thevenin Resistance Plots	88
3-16	Rectified Power Supply Voltage vs. Axial Coil Displacement	89
4-1	Power Supply Rectifier Implementation	91
4-2	Synchronous Rectifier Block Diagram	95
4-3	Clock Timing Diagrams	96
4-4	Synchronous Rectifier Control Function	98
4-5	V_T -dependent Voltage Reference	99
4-6	Voltage References	100
4-7	PMOS Input Voltage Follower	100
4-8	Clocked Voltage Follower	101
4-9	Clocked Comparator	102
4-10	Clock Extractor	104
4-11	Clock Extractor Comparator	105
4-12	Clock Distribution	106
4-13	Clock Generator for the Clocked Comparator in Figure 4-9	106
4-14	Rectifier Block	107
4-15	Continuous Comparator	108

4-16	Bazes Very Wide Common Mode Range Differential Amplifier	109
4-17	Typical One-Shot Circuit	109
4-18	Clock Subdivider	112
4-19	Capacitor Charging Controller	114
4-20	Gate Drive Circuitry	115
4-21	Power On Reset Circuitry	116
4-22	Power On Reset Comparator Circuit	117
4-23	Poly-Poly2 Resistor Layout	118
4-24	Clocked Comparator and Clock Generator Layout	118
4-25	Closer View of Clocked Comparator and Clock Generator Layout	119
4-26	Die Photograph of Clocked Comparator and Clock Generator Layout	120
4-27	Corrected Clocked Comparator and Clock Generator Layout	121
4-28	Synchronous Rectifier Testing Block Diagram	122
4-29	Reference Voltage Circuit Test Results	123
4-30	Clock Extractor Test Results	124
4-31	Clock Extractor Test Results - Rising Edge	124
4-32	Clock Extractor Test Results - Falling Edge	125
4-33	Synchronous Rectifier Test Results - 400 <i>ns</i>	126
4-34	Synchronous Rectifier Test Results - 800 <i>ns</i>	127
4-35	Synchronous Rectifier Test Results - Full Time Rectification	128
4-36	Power On Reset Test Results	128
5-1	“X” Test Pattern	131
5-2	“MIT” Test Pattern	132
5-3	Digital Control System Block Diagram	134
5-4	Main State Machine	136
5-5	Clock Divider	138
5-6	Divide By 2 Flip-Flop	139
5-7	Switch Control Logic	139
5-8	Electrode Control Logic	140

5-9	MIT Pattern Loop Logic	141
5-10	Switch Driver Logic and Switches	142
5-11	Bias Current Source	143
5-12	Switch and Driver Layout	144
5-13	Divide by 2 Timing Diagram	145
5-14	Stimulation System Driving a Series Resistor and Capacitor	146
6-1	Layout and Chip Micrograph of Full VLSI System	148
6-2	Labeled Chip Micrograph of Full VLSI System	149
6-3	Secondary Side Circuit Board	150
6-4	15-Electrode Array	151
6-5	Full Power and Stimulation System	152
6-6	Instrumentation Amplifier for Measuring Electrode Current	153
6-7	Measured Electrode Current and Voltage Waveforms	154
6-8	Instrumentation Amplifier for Measuring Secondary Coil Current	155
6-9	Instrumentation Amplifier for Measuring Secondary Coil Voltage	156
6-10	Coil Voltage and Current Waveforms, 15 Electrodes	157
6-11	Coil Voltage and Current Waveforms, No Electrodes	158
6-12	Power Consumption Histogram for 15 Electrodes	159
6-13	Power Consumption Histogram for No Electrodes	160
6-14	Power Into an Electrode During a Stimulation Cycle	162
A-1	Variables for Capacitor Charging Proof	170

Chapter 1

Introduction

1.1 The Retinal Implant Project

The Retinal Implant Project is a collaboration of the Massachusetts Institute of Technology and the Massachusetts Eye and Ear Infirmary to develop a prosthesis to restore some useful vision to victims of certain types of blindness [42]. Two common diseases, retinitis pigmentosa and macular degeneration, attack the photoreceptors of the outer retina. The ganglion nerve cells of the inner retina, which carry signals to the brain, remain largely intact. By electrically stimulating these nerves, we hope to elicit visual perception that can help a patient navigate a room, or possibly even read.

1.1.1 Retinal Function

In normal vision, when light enters the eye, it is focused by the cornea and lens onto the retina, the thin layer of tissue lining the rear wall of the eye, which serves as the visual sensory organ (Figure 1-1). Light passes through the inner retina to the outer retina, where the photoreceptor cells, the rods and cones, are located. Found throughout the outer retina, rods specialize in transducing low levels of light, aiding in night vision, while cones, primarily concentrated in the fovea, the center of the visual field, are sensitive to different colors and yield an image of higher resolution. The rods and cones convert light to neural signals, which are carried forward to the

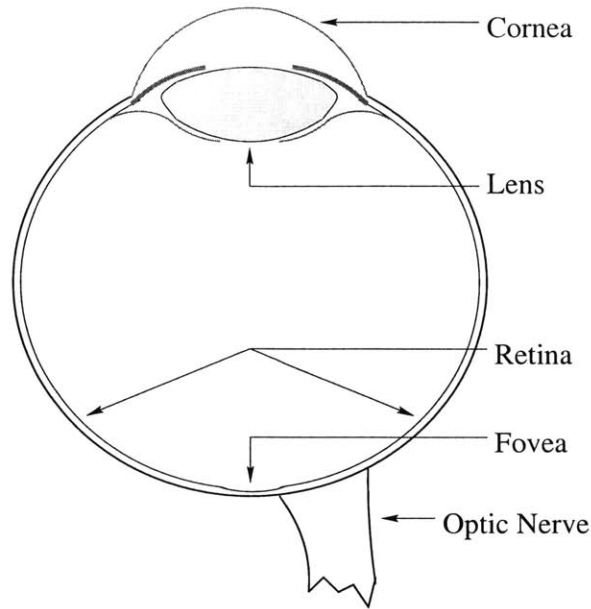


Figure 1-1: The Eye

inner retina via interneurons (horizontal, bipolar, and amacrine cells). There the signals reach the ganglion nerve cells, whose axons form the optic nerve to the vision centers of the brain (Figure 1-2).

One important feature of the retina is its geometric mapping. The photoreceptors, interneurons, and ganglion cells are laid out across the retina, and their locations correspond with the visual stimuli. For example, light coming from the upper left part of the visual field enters the eye and strikes the retina in the lower right part of the eye. The photoreceptors, interneurons, and ganglion cells which process that light are all in the lower right region.

1.1.2 Outer Retinal Disease

Two degenerative diseases of the outer retina eliminate the rods and cones over a period of years or decades, resulting in partial or full blindness. Retinitis pigmentosa is a congenital disease that can develop at any age, and affects 1.6 million people worldwide. Macular degeneration has both congenital and age-related forms, and is the leading cause of blindness in the western world. It affects an estimated 8 million people worldwide. As these diseases deplete the rods and cones, however, they leave




Figure 1-2: The Retina

the interneurons and ganglion cells largely unharmed.

1.1.3 The Retinal Prosthesis

The spatial mapping of the retina allows us to imagine portraying an image to a victim of outer retinal disease by electrically stimulating the ganglion cells in the pattern of the image. A prosthesis with this goal might consist of a camera worn by the patient on a pair of glasses, which would transmit the image to a stimulator circuit in or on the eye, which in turn would stimulate the appropriate ganglion cells. One possible type of implant is shown in Figure 1-3.

There are many significant challenges facing the development of such a retinal implant, not the least of which is learning how to stimulate the nerves in a way which gives a blind person some level of useful vision. Of the many biological and engineering challenges in this project, the one examined by this thesis is that of power consumption within the implant, particularly power consumed by the stimulation electrodes [16].




Figure 1-3: A Retinal Implant

The external camera transmits images to the implant, which are converted to currents stimulating specific patches of retinal ganglion cells.

1.2 Stimulation Electrodes

The electrical signals generated by the implant, typically biphasic constant-current pulses, interface the biological tissue via the stimulation electrodes. These electrodes usually consist of an exposed metal disc for delivering current, with some other area of exposed metal, either local to or distant from the current-delivering electrode, through which the current is returned.

Current is conveyed in the metal electrode by electrons in the conduction band, and in the biological tissue by ions in the fluid. The two current conduction systems interface at the electrode surface by capacitive coupling and/or reduction and oxidation reactions. Some of these reactions are reversible, some are not. The maximum level of charge driven through an electrode of a given size which causes no irreversible redox reactions is considered to be the safe stimulation charge level for that electrode (though not necessarily safe for surrounding biological tissue). This limit is often given as a safe stimulation charge density for the metal [2]. Charge densities above this level may generate irreversible redox reactions, which may degrade the electrodes over time and deposit potentially toxic materials into the tissue.

1.2.1 Commonly Used Electrode Materials

Stimulation electrodes have commonly been made from the noble metals: rhodium, palladium, iridium, platinum, and gold. The three materials considered for stimulation electrodes in this implementation of a retinal implant are discussed briefly.

Gold

Gold is a relatively common material in experimental microfabrication, and is therefore attractive for microfabricated stimulation electrodes. However, when used to deliver the levels of charge necessary for the retinal implant, gold is known to corrode over time [14, 37] and to generate irreversible redox reactions as surface oxides develop [23, 24]. Both of these effects limit the stimulation lifetime of the electrode, and may deposit potentially toxic material in the eye.

Platinum

Platinum is more commonly used for stimulation electrodes than is gold, yet it, too, will corrode with prolonged stimulation [5, 14, 29, 37]. Brummer and Turner found safe stimulation charge density limits for platinum electrodes to be $0.3 - 0.35 \text{ mC/cm}^2$ of real area [6]. This corresponds to as much as 0.45 mC/cm^2 of geometric area due to platinum surface roughness. In Agnew and McCreery's book [2], Robblee and Rose warn that this charge density is a best case number, and that conservative limits should be lower. For chronic stimulation, they suggest limits of $0.05 - 0.1 \text{ mC/cm}^2$ of geometric area for biphasic current pulses with the anodic (positive) current pulse first, and $0.1 - 0.15 \text{ mC/cm}^2$ for cathodic (negative) first pulses [30]. These charge densities are far lower than the $0.8 - 1 \text{ mC/cm}^2$ necessary for retinal stimulation (See Section 1.2.4).

Platinum electrodes were used by the MIT-MEEI group for *in vitro* experiments and rabbit *in vivo* experiments leading up to our acute human stimulation trials, but the Brummer-Turner limits for platinum electrodes would not have been sufficient to induce perceivable responses in a human subject. Platinum electrodes were used by

the Johns Hopkins group in their acute human trials. Perception was achieved, but usually with charge densities far above even the Brummer-Turner limits [12, 13].

Iridium Oxide

The material commonly considered to give the best safe stimulation charge limits is oxidized iridium, or activated iridium oxide. This material can be made either by depositing iridium and then activating the oxide by potential cycling, as described in James Weiland's PhD dissertation [35], or by directly depositing iridium oxide.

Iridium oxide allows for higher safe charge densities than the other materials studied [2]. An early study on iridium oxide showed that anodic-first pulses with charge densities as high as $30 - 40 \text{ mC/cm}^2$ could be applied without generating water electrolysis and gas bubbles [28]. A later study states that the electrolysis limit is too high, allowing irreversible redox reactions, and a limit of $2 - 4 \text{ mC/cm}^2$ is recommended to maintain reversible reactions [4]. Iridium oxide electrodes were used in the first and third through sixth human trials performed by the MIT-MEEI group, and most stimulations stayed below the 4 mC/cm^2 limit [25, 26].

A brief examination of the current *vs.* voltage curve for an iridium oxide electrode, obtained by cyclic voltammetry, shows far more charge capacity at positive voltages than at negative voltages (Figure 1-4). This is taken into account in Beebe and Rose's study [4], which states that a higher safe charge density may be obtained by using anodic-first biphasic current pulses. However, *in vivo* rabbit experiments and *in vitro* retinal experiments have shown that lower charge thresholds are obtained with cathodic-first biphasic waveforms. This conflict may be resolved by anodically biasing the electrodes with respect to the surrounding tissue, making more of the charge capacity available to a cathodic-first waveform. Dr. Stuart Cogan recommends about 0.3 V bias above a platinum electrode.

Iridium oxide, however, does have drawbacks. Continued *in vivo* stimulation through iridium oxide electrodes can alter their electrical properties [36]. Some of the model parameters discussed in Section 1.2.2 are changed temporarily (for hours) while others are changed permanently. In addition, an organic film developed on

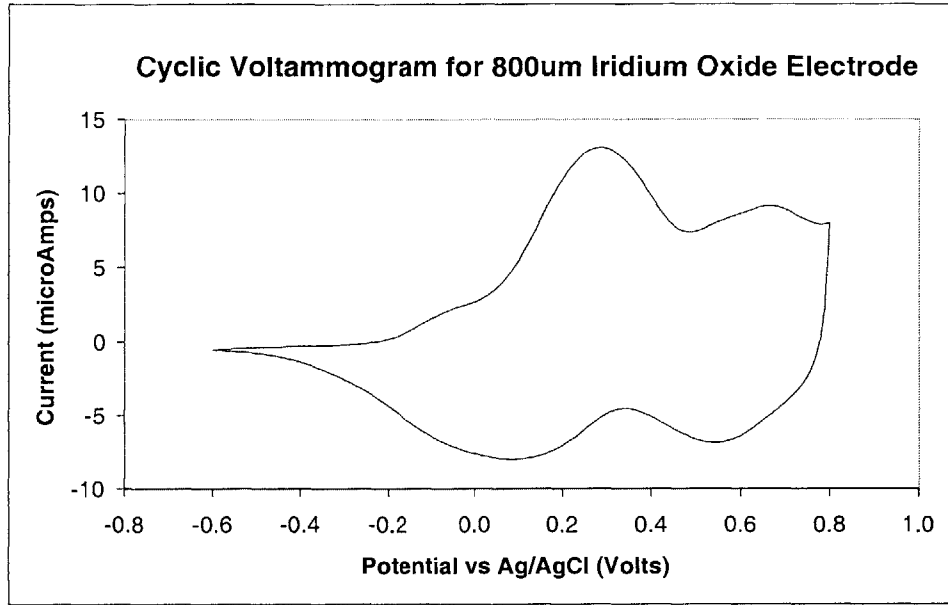


Figure 1-4: Iridium Oxide Cyclic Voltammetry

This figure shows current through electrodes similar to ours as it relates to electrode potential (for a very slowly changing potential). Notice the hysteresis in the plot; positive currents correspond to positive changes in voltage, and *vice versa*. This figure is reproduced from Ken Roach's M.Eng. thesis [27] with permission of the author.

the electrodes used for stimulation. The impedance changes and organic film are a concern, and should be researched more carefully before iridium oxide is used in permanent patient implants, but they do not *de facto* disqualify iridium oxide as an electrode material. In fact, it still seems to be the most promising material for chronic neural stimulators.

1.2.2 Electrical Model of the Electrode


A detailed electrical model of iridium oxide electrodes can be found in Kenneth Roach's M.Eng. thesis [27]. This model, shown in Figure 1-5A, includes a resistance in series with the parallel combination of a capacitive path and charge transfer path. The series resistance represents the resistance of the neural tissue or electrolyte fluid, plus the resistance of cables and wires connecting the electrode to the stimulator. The double-layer capacitance is made up of charges in the iridium electrode

and ions in the solution. The charge-transfer resistance represents current flow via redox reactions at the electrode-fluid interface. Z_W is a Warburg impedance element that is only relevant at low frequencies. For the stimulation pulses we are using, the Warburg impedance merely rounds the corners of the waveforms, producing a minor effect.


The Warburg impedance element arises physically from mass transport limits of ions in the solution. The reaction rate at the electrodes is assumed to be very fast, leaving mass transport as the limiting factor. Since a fluid boundary layer around the electrode shields these ions from convection, and the low electric fields (due to high fluid conductivity) limit migration (drift in electrical engineering parlance), diffusion is the primary mass transport method. Using Fick's laws, and viewing the ion concentrations as small-signal sinusoids about the bulk fluid values (as demanded by a sinusoidal voltage input), the diffusion rate, and therefore current, can be calculated. The resulting impedance has the form $\frac{k}{\sqrt{j\omega D}}$, where D is the diffusion constant from Fick's first law. This happens to have the same frequency dependence as the impedance of an infinite RC transmission line. Far more detailed derivations of the Warburg impedance are provided in [27].

At the frequencies of interest, then, the model simplifies to C_{dl} in parallel with R_{ct} , all in series with R_s . Our tests have further shown that R_{ct} is usually large enough that it may be neglected, leaving Figure 1-5B as a reasonable model.

The back voltage created across the electrode by a biphasic current stimulus was measured to determine the values of R_s and C_{dl} . The waveforms of the current stimulus and back voltage are shown in Figure 1-6. The resistance was found to be approximately $1.1 - 2.5 K\Omega$, and the capacitance, from $300 - 3000 nF$, depending on the level of oxidation. The circuitry used for stimulation in this thesis was designed to interface with electrodes with approximately $1.1 K\Omega$ of resistance and $1 \mu F$ of capacitance. The electrodes used to test the stimulation system were designed to have that target impedance, but ended up with $2.3 K\Omega$ of resistance and $490 nF$ of capacitance, limiting the final system performance somewhat. In final tests, two electrode arrays were used in parallel, giving resistance and capacitance of approximately $1.15 K\Omega$



(A)



(B)

Figure 1-5: Electrical Circuit Model of Iridium Oxide Electrode

(A) A detailed electrode model, including series resistance R_s , double-layer capacitance C_{dl} , charge-transfer resistance R_{ct} , and the Warburg impedance Z_W .

(B) The simpler model primarily used herein, including only the series resistance and double-layer capacitance.

and 980 nF . The target numbers are used in calculations throughout the thesis, since they are the basis of the design.

1.2.3 Chronic Neural Overstimulation Damage

While toxic deposits from electrode corrosion present an easily definable upper bound on safe charge densities for chronic stimulation, it has been found that simply overstimulating neural tissue can cause injury to the tissue [21]. This effect was separated from toxic electrochemical effects by comparing neural tissue stimulated by standard faradaic (exposed metal) electrodes with tissue stimulated by capacitor electrodes (metal electrodes anodized with a thin layer of dielectric, such as tantalum pentoxide). Both of these types of electrodes were placed flush against the parietal cortex of a cat. Tissue stimulated in a similar fashion with both electrodes showed similar damage, leading to the conclusion that the neural damage has more to do with the stimulation process than with the electrochemical reactions [20]. Furthermore, introducing a local anaesthetic to block neural action potentials prevented damage to

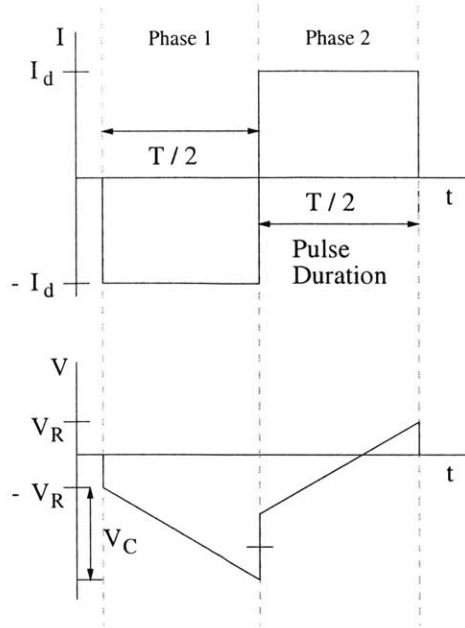


Figure 1-6: Stimulation Current and Electrode Voltage Waveforms

The first and last steps in the back voltage waveform represent the drop across the resistance caused by the step of current. The middle step is twice the magnitude because the *change* in current is twice as large. The ramps represent the charging voltage across the capacitance induced by constant current.

neural tissue exposed to stimulus currents from platinum electrodes [1]. Therefore, the damage seems to be caused by the act of stimulating the neural tissue, and depends not only on the maximum charge density, but also on the charge delivered in one phase of the biphasic current pulse [21]. Deeper analysis showed that the damage depends on the product of charge and charge density, with products over $80 \mu C^2/cm^2$ per phase often causing observable damage, and $32 \mu C^2/cm^2$ per phase proposed as a conservative safe limit [31].

For an electrode with a diameter of $400 \mu m$, a commonly used size in this project, the damage threshold for charge density would be $252 \mu C/cm^2$, and the conservative safe limit would be $159 \mu C/cm^2$.

However, these data come from experiments using a wide range of neural tissue and electrode sizes (from far smaller than our electrodes to far larger). None of the data is from retinal stimulation, and none is from electrodes of the area we are using. Therefore it must be stated that safe levels of charge or charge density for electrical

retinal stimulation are not yet known.

1.2.4 Stimulation Charge Thresholds

The level of electrical charge required to elicit a visual percept, or phosphene, in a human subject was determined by a number of acute human trials.

Acute Human Trials

Six trials have been performed to date, in which a human volunteer, typically blind as a result of retinitis pigmentosa, underwent electrical retinal stimulation for a few hours and verbally reported visual percepts. One volunteer was normally sighted, but had cancer in the eye socket, requiring surgical removal of the eye. In trials 1, 3, 4, 5, and 6, a microfabricated electrode was surgically introduced and placed very near the retina. A battery-powered portable current source stimulation system [15] provided a sequence of biphasic current pulses like the current waveform shown in Figure 1-6. In the last four of these trials (one sighted and three blind subjects), repeatable visual percepts were reported.

Perception Thresholds

The threshold for perception was determined by sequentially increasing current levels (for a fixed electrode geometry, electrode location on the retina, stimulus pulse width, and pulse width repetition frequency) past the level at which perception is first achieved, then decreasing the current levels until no perception is achieved. For several good electrodes (some electrodes happen to end up over severely diseased tissue, or do not conduct properly), the threshold is the lowest current level at which perception is achieved in at least 50% of the stimulations. These data, converted to threshold charge densities, are plotted for four subjects in Figure 1-7.

As can be seen in the figure, the threshold charge levels for the two blind patients with whom multiple pulse widths were used are somewhat consistent for most pulse widths, but rise for the 16ms width. This is likely due to a minimum stimulation

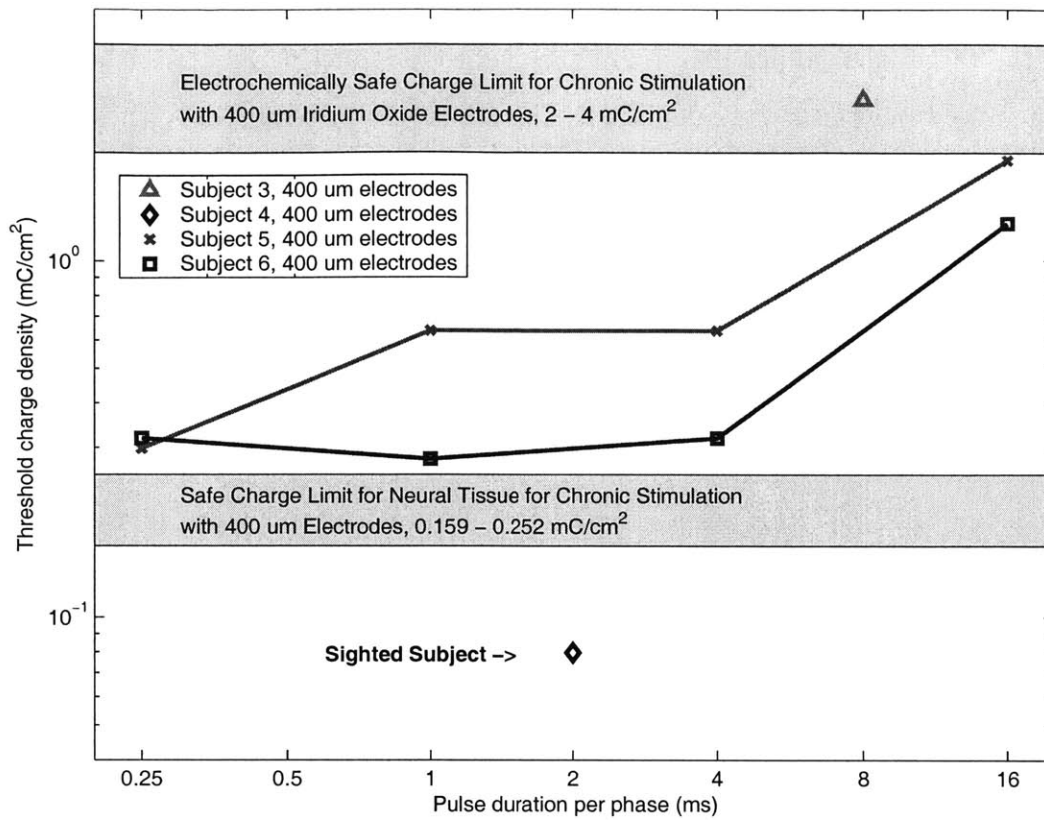


Figure 1-7: Perception Threshold Charge Densities

The thresholds for four subjects are shown. Subjects, 3, 5, and 6 had retinitis pigmentosa, and were blind, capable only of distinguishing bright light from darkness. Subject 4 was normally sighted. Also shown in the upper bar are the charge density limits set by electrochemical reactions in iridium oxide electrodes and in the lower bar, the charge density limits set by neural damage from chronic stimulation.

current required for perception. Note also that the threshold charge density for the sighted subject is far lower than for the blind subjects. From these data, for a proposed pulse phase duration of 5 ms, a target charge of 1 μC was chosen, giving a charge density of 0.8 mC/cm². This target value is higher than the threshold values at 4 ms for subjects 5 and 6, and should be sufficient to reach perception threshold in a patient using the implant daily and having time to learn how to use it.

1.2.5 Electrode Power Requirements

With the target charge per phase and the electrode model parameters, the approximate required power can be calculated. The actual power required by the electrodes depends highly on the architecture of the stimulator, and will be discussed in depth in Chapter 2.

To estimate the energy required to drive the electrodes, let us examine the power consumed by the series resistance during one stimulation cycle. If we assume the simple RC model shown in Figure 1-5B, then the capacitance simply stores energy temporarily, burning no average power. Taking the target resistance value of $1.1 K\Omega$, target charge of $1 \mu C$, the pulse waveform of Figure 1-6, and pulse phase duration of $5 ms$, we calculate that the power delivered to the resistance during stimulation is $P = (200 \mu A)^2 (1.1 K\Omega) = 44 \mu W$. At the maximum frequency of $100 Hz$, the stimulation current is constantly flowing, in one direction or the other, through every electrode. A 100 electrode array would then consume $4.4 mW$. This, however, is the power consumed only *within* the electrodes; additional power is consumed within any current-delivery circuitry. We estimate the total required power for 100 electrodes stimulated at $100 Hz$ (including power lost in the eye, stimulation circuitry, and power transmission receiver coil, described later) to be $10 mW$.

1.3 Power and Stimulation System

In addition to the usual challenges of designing an implant for the body, three special challenges exist for a retinal prosthesis: the eye is constantly in motion with respect to the head, the eye is a relatively small organ, and the retina is extremely delicate tissue only tenuously held to the back of the eye. The first challenge precludes the placement of a large battery pack elsewhere in the body, with wires running to the eye. Therefore all electronics must be inside the eye or attached to the outside of the eye, and power and data must be transmitted wirelessly to the implant. The second and third challenges make difficult the design of a power receiver, be it a coil for collecting magnetic fields or a photodiode array for collecting light. The data receiver

is a more straightforward problem, and is not examined in this thesis.

1.3.1 Power Delivery Methods

The two common methods for delivering remote power to an implant are laser and radio-frequency magnetic fields. The latter is far more common, but in the optically transparent environment of the eye, the former is also an option.

Laser

Collecting power inside the eye from an external laser via an array of photodiodes was considered early in the project, and revisited in Erich Caulfield's master's thesis [7]. If the laser can be aimed directly at the photodiode array, this is a reasonable method of power delivery, but it has several drawbacks. With a fixed laser, the patient must align the laser with the photodiode array by properly positioning his eye, effectively looking at the same point in space all the time. This drawback can be overcome by adaptively aiming the laser to hit the array, but it must be assumed that the laser will sometimes inadvertently hit the retina directly, and therefore the laser power must be low enough not to damage the retina. This limits the power that can be transferred to the implant to 15 mW for a 1 cm^2 photodiode array. However, an additional 85 mW will be dissipated within the photodiode array or absorbed by the retina.

Magnetic Fields

Much work has been done on coupled coil power transmission [19], [34], [8]. If a primary coil outside the eye generates an AC magnetic field, a secondary coil inside or attached to the eye can couple that magnetic field to generate power. This solves the problem of aiming the laser, since the magnetic field need not be perfectly "aimed", but a new problem arises: the magnetic field amplitude decreases with distance from the primary coil. Furthermore, just as the laser intensity is limited by the potential for retinal damage, so too is the magnetic field amplitude limited by the potential for induced heating in the eye. This heating, examined briefly in Section 2.2.4 and found

to be very small, is assumed to be primarily resistive heating from electrical currents induced by the magnetic field. It maybe the case that the electric field around the coil causes additional resistive heating. Furthermore, the mechanisms of heat transfer from the eye (thought to be largely via the massive retinal blood flow) are not studied here. Additional heating may come from the I^2R losses in the secondary coil itself. This is especially true if the coil system must be matched to the load for maximal power transfer, burning half of the total power within the coil. This issue is studied further in Sections 2.2.5 and 2.2.6. This is a serious problem for very small coils, but, as will be explained in Section 2.2.5, the coil can be made much larger and placed on the outside surface of the eye, allowing for the potential delivery of over $100mW$ with a load matched to the coil impedance. While this much power would likely be damaging to the eye, the *potential* to deliver this much power means that the coil can be underloaded, so that the implant receives the power it requires to drive the electrodes, but the power consumed by the coil amounts to only a few percent of the implant power.

1.3.2 Receiver Coil Placement

Four potential locations of the secondary receiver coil are shown in Figure 1-8. The most convenient from a design perspective would be against the retina (Figure 1-8 A). This would allow the coil, electronics, and electrode to be implanted as one package. However, the fragility of the retina precludes placement of a thick or heavy implant, limiting the possible power for this location to a maximum of $2.3mW$. Placing the coil in the anterior chamber (Figure 1-8 B) allows $10mW$ to be delivered, but is surgically difficult. An attractive alternative is to place the coil on the outside surface of the eye, either under the conjunctiva on the front of the eye (Figure 1-8 C), or on the temporal side of the eye (Figure 1-8 D). These two locations each allow over $100mW$ of power to be theoretically delivered, with a coil size that is comfortable for the surgeons.

The temporal location was chosen partly for its power-delivering capabilities, but also for ease of surgical access, allowing the coil and electronics to be attached to the

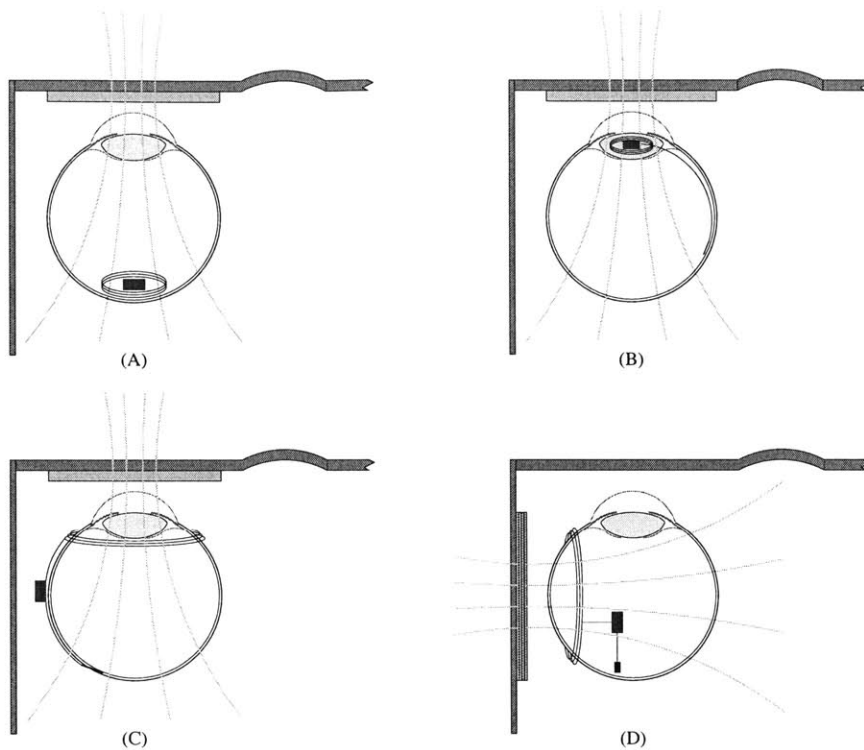


Figure 1-8: Possible Retinal Implant Locations

The secondary coil may be placed (A) against the retina, (B) in the lens capsule of the anterior region, (C) under the conjunctiva outside the eye, or (D) outside the eye on the temporal side.

sclera on the outside of the eye while an electrode array extends through a flap in the sclera to the subretinal space.

1.3.3 High-Level System Description

The system described in this thesis can be broken down into three main parts: the coupled coils delivering the power, the rectifiers setting up the power supplies, and the system for driving current through the electrodes. (Figure 1-9) The rectifiers further consist of a standard rectifier which establishes the power supply for the chip, and a synchronous rectifier which establishes several voltages on the energy storage elements necessary for driving the electrode. The task of driving current to the electrodes is handled by measurement and control circuitry and a switching network, as directed

by a few preset patterns of stimulation, since data transmission is not included in this thesis.

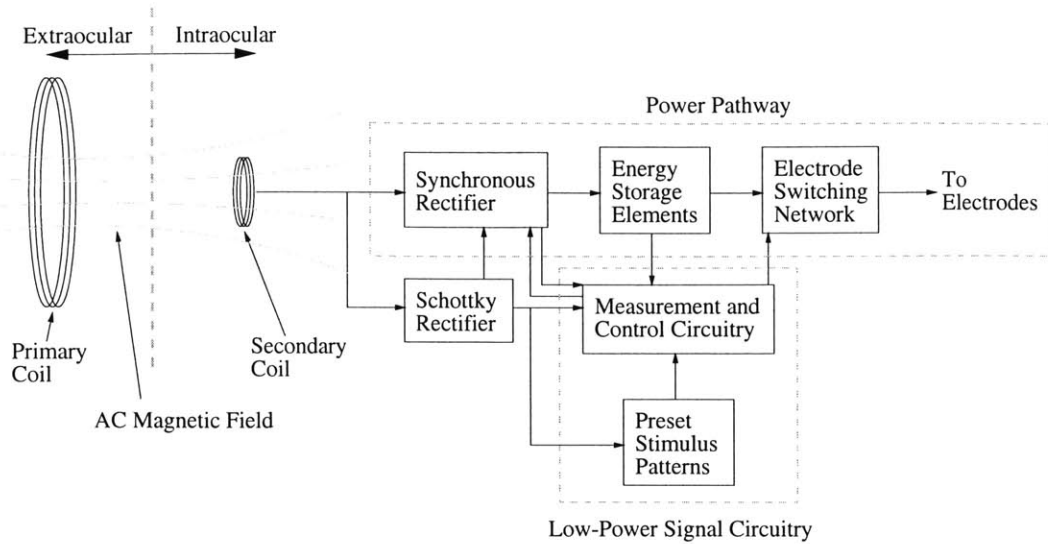


Figure 1-9: System Block Diagram

While each of these three parts performs an independent task, the design of each part depends substantially on the performance and requirements of the other parts, making the overall design somewhat complicated. These three main subdivisions are discussed further in Chapters 3, 4, and 5.

1.4 Related Work

The predecessors to our human trials were performed by Mark Humayun, Eugene de Juan, et. al., at Johns Hopkins University, now relocated to the University of Southern California. Some of the circuitry for their proposed implant has been developed by Wentai Liu's group at North Carolina State University. Philip Troyk at the Illinois Institute of Technology has explored power transmitters for cortical implants. Additionally, Second Sight, a for-profit company founded by Alfred Mann, is attempting to develop a commercial retinal implant.

In addition, one of the main concepts presented in this thesis for lower-power tissue stimulation was developed independently for adiabatic switching in CMOS circuits by Svensson and Koller [33].

1.5 Thesis Outline

Chapter 2 – Background Theory and Possible Architectures

Describes the critical problems facing the implant power system design, examines various current drive methods and compares the power consumed, explores energy recycling from electrodes, compares possible rectification architectures, and examines effects of absorbed magnetic field power and loading on the secondary coil.

Chapter 3 – Coupled Coils

Describes the design of the primary and secondary coils and their construction, describes the primary coil driver, and gives results of coil system testing.

Chapter 4 – Power System

Describes the implementation of the power supply rectifier, electrode driver rectifier, and integrated circuit components of the rectifiers. Discusses circuit and layout issues, and gives results of power system testing.

Chapter 5 – Electrode Drive

Describes the design of the switched-voltage architecture, including state machine and control circuitry design and implementation, and gives results of electrode drive system testing.

Chapter 6 – System Integration

Describes the system integration and testing.

Chapter 7 – Conclusions and Future Work

Summarizes the contributions in this thesis, and suggests areas where work still needs to be done.

Appendix A

Includes mathematical derivations and MATLAB scripts.

Chapter 2

Background Theory and Possible Architectures

This chapter is meant to span the gap between the general system descriptions in Chapter 1 and the specific implementation details in Chapters 3, 4, and 5 by describing the system design process. This organization is made necessary by the unusual interdependence of the three areas of focus in this work.

2.1 Methods of Electrical Neural Stimulation

Since neural stimulation through the iridium oxide electrodes is the primary goal, let us examine possible methods of stimulation, and the power consumed.

2.1.1 Constant Current Drive - The Optimal Charging Waveform

The typical method of stimulating neural tissue is driving a constant current through the electrodes for a fixed time, then driving the reverse current for the same time, as shown in Figure 1-6, repeated for convenience in Figure 2-1. For the simple series RC electrode model shown in Figure 1-5B, this results in the magnitude of the capacitance voltage ramping up and then back to zero, while the voltage across the resistance has

the same waveform shape as the current. The two in series have the total voltage shown in Figure 2-1 Examining only the first phase of constant current, and recalling from Section 1.2.5 that we need to deliver $1 \mu C$ of charge in $5 ms$, we see that *any* current waveform which integrates to $1 \mu C$ at the end of $5 ms$ is acceptable. However, we wish to minimize the total energy consumed during this time. The energy stored in the capacitance at the end of the $5 ms$ is an unavoidable constraint, but the energy dissipated in the resistance depends on the current waveform.




Figure 2-1: Stimulation Current and Electrode Voltage Waveforms

We prove in Appendix Section A.1 that the constant current waveform often used in stimulators is in fact the optimal waveform for minimizing power consumed within the electrode resistance [39]. Intuitively this makes sense, because to deliver the target level of charge during phase 1 with any other current waveform would require more current at some times and less current at other times, but since $P = I^2 R$, the power increase during the times of higher current, and therefore higher resistive voltage drop, exceeds the power decrease during times of lower current. For example, if we drive with twice the current for half of phase 1, and no current for the other half, we would burn 4 times the power for half the time, consuming twice the total energy. Therefore the constant current waveform minimizes the energy consumed *within* the

electrode resistance during the task of delivering the target charge to the electrode capacitance in the target time.

2.1.2 Electrode Power and Energy




Figure 2-2: Electrode Power Waveform

The power waveform of an electrode, viewed as separate resistive and capacitive components (left) and as input and output power (right).

A close examination of the second phase of the current and voltage waveforms in Figure 2-1 reveals power being supplied by the electrode. The electrode's power, the product of the current and voltage shown in Figure 2-1, is shown in Figure 2-2 from 2 different perspectives. In the waveform on the left, the area with horizontal lines shows the total energy consumed in the electrode resistance. This area is the minimum energy required to drive the electrode. The first triangular area with vertical lines shows the energy delivered to the capacitance from the source, the second triangular area with vertical lines shows the energy returned to the source from the capacitance, and the triangles have equal area – all of the energy supplied to the capacitance is subsequently removed. Over the first phase, the source supplies energy to both the resistance and the capacitance. During the second phase, at any given moment, the source supplies the *difference* between the power required by the resistance, I^2R , and the power being returned by the capacitance. For some time, at the beginning of the second phase, the capacitive power exceeds the resistive, so the power supplied by the source is negative, i.e. the electrode is *supplying* power to the source. This is shown more clearly in the waveform on the right. The red area with horizontal lines shows

the energy supplied from the source to the electrode, while the green area with cross-hatching shows the energy supplied from the electrode to the source. The net energy supplied to the electrodes is the red/horizontal area minus the green/cross-hatched area, which equals I^2RT .

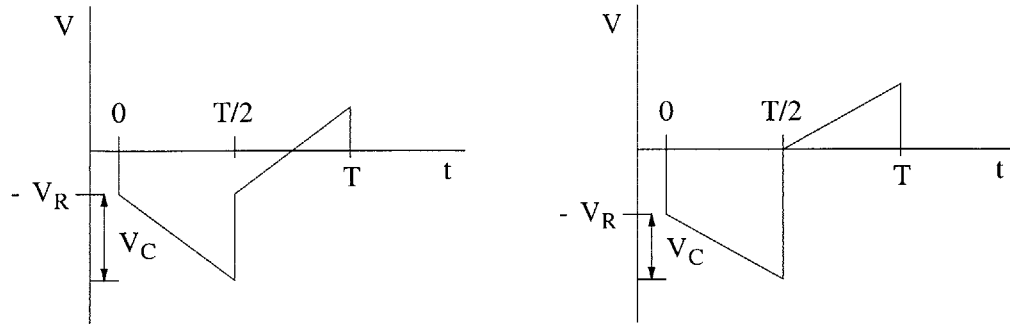


Figure 2-3: Back Voltage Waveforms for 2 Special Cases

Figure 2-3 shows two special cases. In the figure on the left, which we will call case 1, the second phase is half negative and half positive. The electrode supplies power to the source during the first half of the second phase, and the source supplies power to the electrode during the second half, and therefore no net energy is recovered from the electrode (or delivered to the electrode) during the second phase. Put another way, the electrode resistance consumes all of the energy stored in the electrode capacitance and no additional energy over the second phase. The figure on the right shows the more extreme case 2, in which no instantaneous power is ever recovered from the electrode. At the start of the second phase, the instantaneous power being returned by the capacitance is entirely consumed by the resistance.

Case 1 occurs if $V_C = 2V_R$, or $IT/2C = 2IR$, or:

$$RC = \frac{T}{4} \quad (2.1)$$

In case 2, $V_C = V_R$, so:

$$RC = \frac{T}{2} \quad (2.2)$$

If the net energy supplied by the source to the electrode during one biphasic pulse pair is I^2RT , then the average power delivered during stimulations with frequency

f is $I^2 R T f$. Since the charge threshold required for stimulation varies only slightly with pulse duration, we will take it as a constant, $Q = IT/2$. So the average power becomes $4Q^2 R f/T$, where Q and R are constants. If the electrodes are driven at the highest possible rate, given the pulse duration or, alternatively, if the pulse duration is set to minimize power losses given a stimulation frequency, then $f = 1/T$, and the average power is:

$$P_{avg} = 4Q^2 R \frac{f}{T} = 4Q^2 R f^2 \quad (2.3)$$

The first thing to notice from the above discussion is that delivering the required charge in a shorter period of time is always less efficient. Second, it is worth repeating the statement in Section 2.1.1 that constant current drive into the electrodes minimizes the power consumed within the electrode. Third, Equation 2.3 represents the absolute minimum power necessary to deliver the required charge within the required time (or with the required repetition frequency). In order to stimulate with these parameters using anywhere near this power value, the rest of the system, including the means for transferring power into the implant and means for transmitting current to the electrodes, must take negligible power compared to the electrode power.

2.1.3 The Canonical Stimulation Parameters

Before we examine methods for delivering the stimulus charge, it is necessary to put forth a consistent set of stimulation parameters to facilitate comparison between methods. This will be referred to as the canonical stimulation parameters, and refers to 100 electrodes, each modeled as a $1.1 K\Omega$ resistor in series with a $1 \mu F$ capacitor, with each electrode receiving $-1 \mu C$ of charge in $5 ms$, followed by $+1 \mu C$ of charge in $5 ms$, repeated continually at 100 Hz.

2.1.4 An Inefficient Constant Current Drive Method

While constant current drive minimizes the power dissipation within the electrode resistance, there are many circuit methods for generating constant or nearly constant

currents, and the amount of power consumed in the circuitry can vary drastically between methods.




Figure 2-4: Dual Supply Current Source Drive

$$V_{ss} = -V_{dd}; \quad V_{dd} \geq 2.5 \text{ V}$$

$$P = (2.5 \text{ V})(200 \mu\text{A})(100) = 50 \text{ mW}$$

The typical method of generating constant currents, shown in Figure 2-4, is to control the impedance of one or more transistors to maintain constant current flow from the DC voltage supply to the electrode. This version shows the simplest and most efficient circuit capable of driving the required current waveform into the electrodes. Most such circuits will include cascode transistors and use power supply voltages of two or more times this value. Yet even with this most efficient example of this circuit, constant stimulation through 100 electrodes requires 50 *mW* of power, whereas Section 1.2.5 showed that the minimum possible power required to drive 100 electrodes was 4.4 *mW*. The difference in power is consumed by the current source transistors. This traditional method in its most efficient form requires an order of magnitude greater power than the minimum possible.

2.1.5 Inter-Electrode Energy Recycling

One method for reducing the total consumed power involves recycling some of the energy from the electrode capacitance.

The Simplest Method

A charged electrode may be connected directly to an uncharged electrode, thereby recovering half of the charge and 1/4 of the energy from the charged electrode, as shown in Figure 2-5. However, if this method is applied to the transistor-based current source discussed above, the total power for 100 electrodes driven at 100 Hz may be reduced from 50 *mW* to 37.5 *mW*, still far above the theoretical minimum of 4.4 *mW*. Furthermore, the stimulation pulse width is limited by the RC time constant of the electrode, so very short pulse widths are not possible with this method.




Figure 2-5: Inter-Electrode Energy Recycling

While this method may not be the best solution, it shows an implementation of the principle discussed in Section 2.1.2, that a charged electrode may be used to supply power during its discharge phase.

More Efficient Methods

A simple improvement to the above method is slightly precharging the uncharged electrode. This reduces the recycled charge, but recycles the charge at a higher voltage, increasing the overall recycled energy. The optimal level of precharging is 1/3 of the target charge. This recycles only 1/3 of the charge from the charged electrode, but recycles 1/3 of its energy. This method is compared with the previous

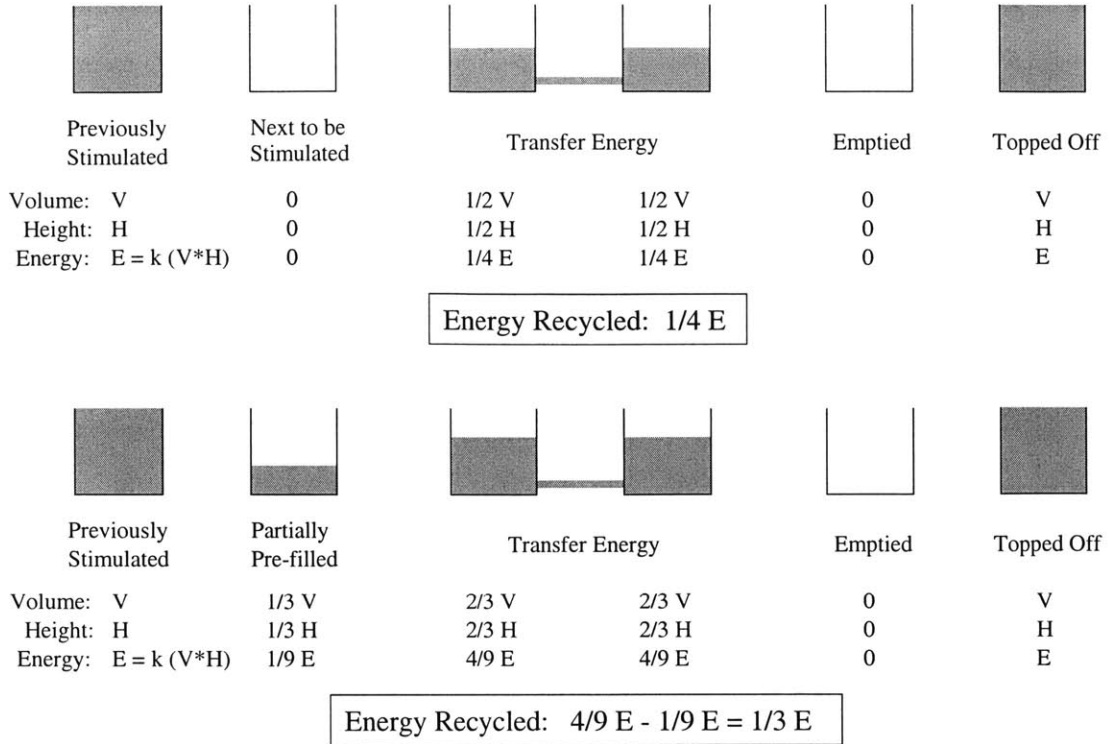


Figure 2-6: Electrode-to-Electrode Recycling - Hydraulic Analogy

By precharging the uncharged electrode, the recycled energy is increased from 1/4 to 1/3 of the energy from the charged electrode (filled bucket).

method, using a hydraulic analogy, in Figure 2-6. In this analogy, a small volume of water V at height H in the bucket has energy $E = gmh = g\rho VH$ for water with mass density per unit volume ρ . If the bucket has vertical walls and a bottom area A , then $V = AH$. Thus the total energy in a bucket filled to height H is:

$$E = \int_0^H g\rho V dH = \int_0^H g\rho AH dH = \frac{1}{2}g\rho AH^2 = \frac{1}{2}g\rho VH$$

so the total energy in a bucket is proportional to the product of the volume of water V and the height of the water H .

More energy may be recovered by using more than one charged electrode to drive the new electrode. Charge may be transferred efficiently by connecting the new electrode sequentially to an array of previously driven electrodes, each having slightly more charge remaining than the one before. For example, the new electrode is first

connected to a mostly discharged electrode, with 10% of the target charge. Approximately half of that charge is shared, and then the new electrode is connected to an electrode with 20% of the target charge. This process continues until the new electrode is nearly charged, at which point the remaining required charge is supplied from a current source.

This method is efficient because it keeps the voltage across the electrode resistances small and not greatly varying, and includes no other circuit resistances (except for the switches, which have negligible resistance compared to the electrode resistance). This is the same as saying that the current through the electrode resistance, and therefore the drive current for the electrode, is not greatly varying. So we can begin to approximate the ideal current waveform discussed above by switching to a large number of previously driven electrodes with varying residual charge.

Unfortunately, this method requires extensive controls to manage a large number of electrodes at varying levels of discharge. Furthermore, the available recycled energy supply is extremely data-dependent: the availability of “supply” electrodes depends on which electrodes have been driven at what times in the past. This method may be improved simply by not using previously charged electrodes as the energy storage mechanism, but by using separate capacitors. This is the approach taken in this thesis.

2.1.6 The Optimal Drive Voltage Waveform

Since we have a good estimate of the electrode impedance, we can generate the required current waveform by driving the electrodes with a voltage source which exactly follows the back voltage waveform shown in Figure 2-1. Performing this function with a perfect voltage source is the method for driving electrodes with the minimum possible power, with no power loss except that in the electrode resistance.

Note that, according to Section 2.1.2, there can be cases in which no energy, nor even instantaneous power, is returned from the electrode capacitance to the source. Even in these cases (when RC is larger than $T/4$ or $T/2$), the back voltage is still the optimal drive waveform. Using a time-varying voltage source that is identical to the

back voltage waveform is *always* the optimal means to deliver power to the electrodes with minimum consumed power.

This optimal waveform is approximately what is achieved in the electrode energy recycling method described above. The new electrode was switched to a series of electrodes with sequentially greater charge (capacitive voltage). The ideal solution is to generate the step-ramp back voltage waveform precisely. While generating such a waveform exactly takes a fair amount of control complexity, approximations to this waveform make implementation simple.

Ramping Capacitor Voltage

The back voltage waveform in Figure 2-1 is a series of steps and ramps. Steps are easily obtained by switching the electrodes from one voltage to another, and a ramp may be obtained by supplying a capacitor with a constant current. In this method, the capacitor voltage starts at $-V_R$, and the electrode is switched from its return to the capacitor, causing the first voltage step. The capacitor voltage is then ramped downward by controlled rectification from the AC power source, as shown in Figure 2-7, while supplying approximately constant negative current to the electrode. When the capacitor voltage reaches $-V_C - V_R$, the electrode is switched to another capacitor whose voltage is $-V_C + V_R$. This capacitor is then ramped in a positive direction, also by controlled rectification from the AC supply. When the voltage of this second capacitor reaches V_R , the electrode is switched back to its return, completing the waveform. The voltage waveform and resultant current waveform for this method are shown in Figure 2-8, and are a reasonable approximation of the ideal values.

This stimulation method seems like it should yield the lowest possible power dissipation for any real implementation of the required current stimulus, but there are some drawbacks.

The first ramping capacitor ends the first phase with a voltage of $-V_C - V_R$, and the second capacitor ends the second phase with a voltage of V_R . If stimulation frequencies are high, then the electrodes must immediately be switched to a capacitor waiting at $-V_R$. The best way to do this is to introduce a third ramping capacitor.

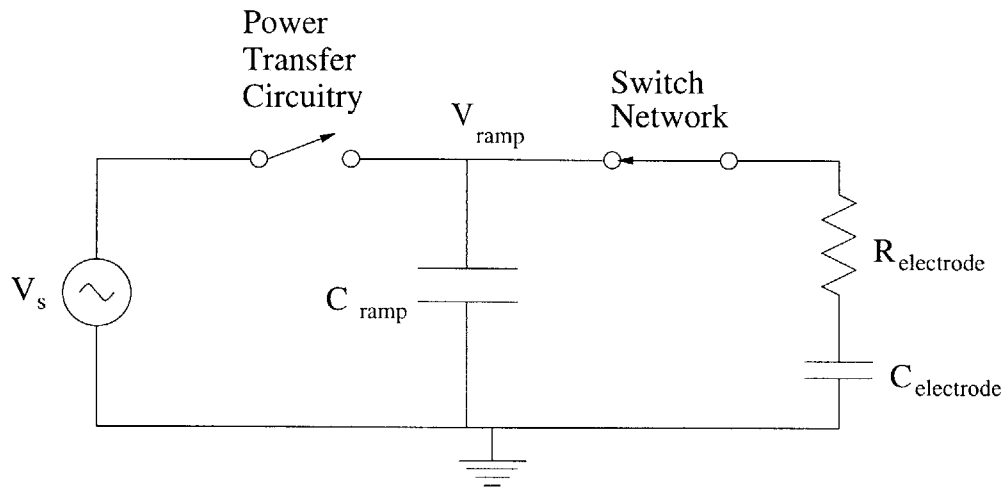


Figure 2-7: Ramping Capacitor Drive

In this circuit, the capacitor voltage is a ramp pattern, to simulate the back voltage ramp seen in driving an electrode with a constant current.




Figure 2-8: Voltage and Current Waveforms for Ramping Capacitor Drive

The voltage and current in this method closely approximate the ideal values (dashed lines).

At the end of a capacitor's ramp, it has two pulse widths of time to leak back into the AC source and change its voltage by $2V_R$ to be ready to start ramping the opposite direction. For example, if the electrodes are first connected to capacitor C_1 ramping down, then C_2 ramping up, for the next biphasic current pulse pair they would be connected to C_3 ramping down, then C_1 ramping up, followed by C_2 down and C_3 up, etc. This would allow high repetition frequencies and minimize the voltage swing necessary to reset a capacitor for the next ramp. Of course, not all of the charge removed from the capacitor is lost. In the transition from $-V_C - V_R$ to $-V_C + V_R$, as well as from V_R to the return voltage (ground in Figure 2-7), the charge may be returned to the AC source. In reality, the Q of this source is expected to be moderate (at most 16), and it may not be worth the added complexity to recycle the extra energy.

Another drawback is the circuit complexity of this implementation. Some reference step-ramp voltage must be generated, and used to compare the capacitor voltages for the purpose of controlling the rectifier. Furthermore, this step-ramp reference waveform must be scalable in both step size and ramp slope in order to accommodate changes in desired stimulus current amplitude.

This implementation is far from impossible, but its power savings in this case were not expected to make up for its complexity.

Multiple Capacitor Bank

A system for approximating the step-ramp voltage which is slightly coarser than the ramping capacitor, but simpler to implement, is a multiple capacitor bank. In this system, several capacitors are charged to various voltages and held constant. The electrodes are then switched sequentially to a series of increasingly negative voltages, then to increasingly positive voltages, to simulate the step-ramp waveform. This circuit idea, with 3 voltage steps, is shown in Figure 2-9, and its waveforms are shown in Figure 2-10.

This circuit includes no ramping control voltage, only a string of DC voltages, which are easier to implement, and may be scaled together to control the delivered

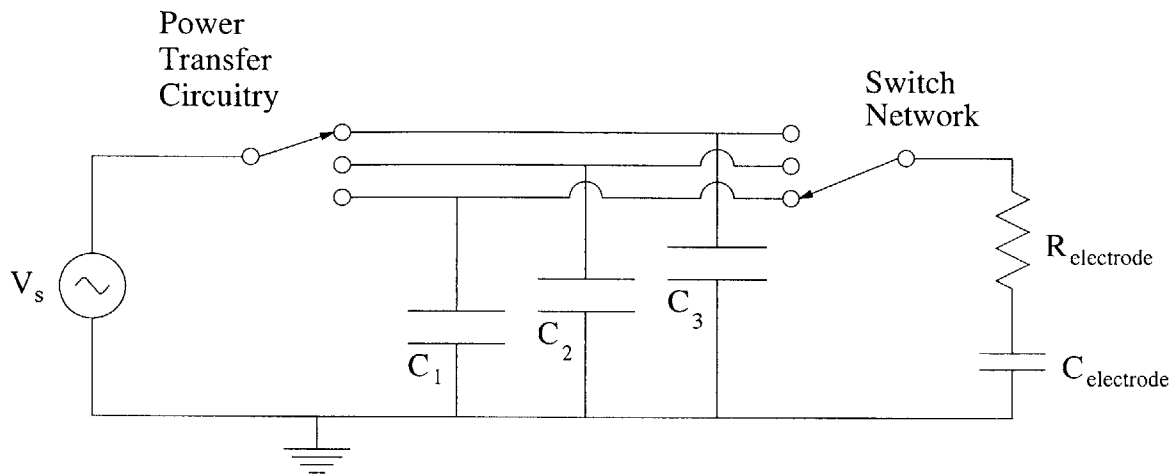


Figure 2-9: Multiple Capacitor Bank

In this circuit, the electrode is switched sequentially to capacitors of increasing voltage magnitude, approximating the desired ramp waveform.




Figure 2-10: Voltage and Current Waveforms for Multiple Capacitor Bank

The voltage and current in this method are coarse approximations of the ideal values (dashed lines).

charge.

2.1.7 Optimal Capacitor Number and Voltage

If a bank of capacitors at varying voltages is used to drive the electrodes, the number of capacitors and their DC voltages must be determined. For control simplicity, the electrodes are assumed to be connected to each capacitor for an equal amount of time, so if 5 voltage steps are used during the 5 *ms* phase, then the electrodes are connected to each capacitor for 1 *ms*.

Simple MATLAB scripts (included in Appendix Section A.2) were used to calculate the total energy consumed by an electrode. The supply capacitors are assumed to be constant voltage sources, and the charge delivered to the electrode during the connection time is calculated. For example, for a previously uncharged electrode with resistance R and capacitance C connected to voltage source V for time T , the charge delivered, Q , is equal to $CV(1 - e^{-T/RC})$, and the energy delivered is QV . The sum of the energy from each source, including the sources delivering negative energy, is the total net energy for one stimulation.

The capacitor voltages were set to be equal steps, with the exception of the lowest and highest voltages, which were set to whatever voltage is required to reach the threshold charge in the fixed time. Furthermore, an anodic bias, as discussed in Section 1.2.1, is added to the electrodes in the form of one voltage step. The electrodes sit at the bias, and the first voltage step in the stimulation is to 0 volts. The voltage waveforms for a 3-step, 4-step, and 5-step implementation are shown in Figures 2-11, 2-13, and 2-15, respectively, and the resultant currents are shown in Figures 2-12, 2-14, and 2-16.

It was found that a 5-step implementation takes 481 *nJ* per stimulation, a 4-step takes 506 *nJ*, and a 3-step takes 580 *nJ*. This corresponds to 9%, 15%, and 32%, respectively, above the minimum required power. The 4-step implementation was chosen for 2 main reasons: it seems like a point of diminishing returns, with a small savings for going from 4 steps to 5, and because 4 steps is a good number for digital implementation in a finite state machine. Five steps requires either 3 state bits per

phase or 5 state bits for the biphasic pulse, and much more combinational logic.




Figure 2-11: Three-Step Voltage Waveforms (Calculated)
 Real and ideal voltage waveforms for a 3-step system with 4 capacitors. This system requires 580 nJ per stimulation.




Figure 2-12: Three-Step Current Waveforms (Calculated)
 Real and ideal current waveforms for a 3-step system.




Figure 2-13: Four-Step Voltage Waveforms (Calculated)
 Real and ideal voltage waveforms for a 4-step system with 5 capacitors. This system requires 506 nJ per stimulation.




Figure 2-14: Four-Step Current Waveforms (Calculated)
 Real and ideal current waveforms for a 4-step system.




Figure 2-15: Five-Step Voltage Waveforms (Calculated)
 Real and ideal voltage waveforms for a 5-step system with 6 capacitors. This system requires 481 nJ per stimulation.




Figure 2-16: Five-Step Current Waveforms (Calculated)
 Real and ideal current waveforms for a 5-step system.

2.1.8 Stimulation Power Comparison

To give a firmer reasoning for the architecture choice, we show a normalized power calculation for each method. Assume that 100 electrodes each require $1 \mu C$ of charge, delivered by approximately $200 \mu A$ over $5 ms$, followed by current of opposite sign for the same duration. These stimulation pulses are repeated at $100 Hz$, the maximum rate for these pulse durations. The electrodes are assumed to consist of $1.1 K\Omega$ of resistance in series with $1 \mu F$ of capacitance.

As stated previously, the traditional transistor current source requires at least $50 mW$ to deliver the required charge, while the theoretical minimum required power is $4.4 mW$. Table 2.1 lists power consumption for other topologies, including the multiple capacitor bank.

Drive Method	Power Consumed
Transistor Current Source	$\geq 50.0 mW$
- with Electrode-to-Electrode Recycling	$37.5 mW$
- with Recycling and Pre-Charge	$33.3 mW$
Capacitor Bank, 3-steps	$5.8 mW$
Capacitor Bank, 4-steps	$5.1 mW$
Capacitor Bank, 5-steps	$4.8 mW$
Ideal Step-Ramp Voltage Source	$4.4 mW$

Table 2.1: Power Consumption for Various Stimulation Methods
Methods are normalized for 100 electrodes ($R = 1.1 K\Omega$, $C = 1 \mu F$) driven to $1 \mu C$ in $5 ms$, discharged for $5 ms$, repeated at $100 Hz$.

Table 2.1 makes clear the magnitude of the benefit of using a capacitor bank, and Section 2.1.7 justifies the selection of 4 steps for this design.

2.1.9 Electrode Drive Tests in Saline

To verify that the currents calculated above for a capacitor bank electrode drive are realistic, electrodes were connected to an arbitrary function generator and placed in a saline solution (0.9% isotonic sodium chloride irrigation solution). A simple current sense amplifier served as the return path, sensing current through the electrode.




Figure 2-17: Test Circuit for Electrode Drive Tests in Saline

The Agilent 33250A arbitrary voltage waveform generator drives current into the electrodes, through the saline bath and back through a local return electrode, held at virtual ground by the op amp. The op amp is an LT1355 and $R_{sense} = 10\text{ K}\Omega$.




Figure 2-18: Test System for Electrode Drive Tests in Saline

Figure 2-17 shows the saline test setup. The function generator drives the electrodes with either the multiple-step waveform or the ideal step-ramp waveform. The current passes from the electrode, through the saline, and back via a local return on the array. The op amp holds the return at virtual ground, and forces the return current through the sense resistor, so that $V_{out} = -I_{drive}R_{sense}$. A photograph of the test system is shown in Figure 2-18.

Figure 2-19 shows the electrode array used for the saline drive experiments. The polyimide array includes a 4x5 grid of 400 μm IrOx electrodes, interspersed with a 4x5 grid of 100 μm IrOx electrodes, with substantial IrOx return area surrounding the electrodes. The electrodes are in the rightmost part of the array, and cover a rectangular area that is approximately 2 mm x 2.5 mm. The large area on the left contains the contact pads for attachment to a PC board.




Figure 2-19: Array with 40 Electrodes

This electrode array was used in voltage drive tests in saline. This is the same type of array used some of our human surgical trials. Shown at about 1.4X actual size. (The vertical line near the center of the figure is the juncture of 2 glass slides holding down the array.)




Figure 2-20: Four-Step Voltage Waveforms (Measured)




Figure 2-21: Four-Step Current Waveforms (Measured)




Figure 2-22: Step-Ramp Voltage Waveforms (Measured)




Figure 2-23: Step-Ramp Current Waveforms (Measured)

This figure shows overlaid plots of 5 different electrodes on one array, each independently driven at different times by the same voltage waveform. Note the invariance in impedance across electrodes.

Figures 2-20 and 2-21 show the voltage and current, respectively, for the 4 step saline experiments. Figures 2-22 and 2-23 show the voltage and current, respectively, for the step-ramp saline experiments.

2.2 Power Recovery and Rectification

As discussed in Section 1.3, the implant must be powered by AC magnetic fields delivered from outside the body. The coil designs and the power circuitry design are closely related: the electrode requirements determine the power needed from the secondary coil, the circuitry determines the voltage needed from the secondary coil, and the coil resistance and inductance determine the maximum current that can be drawn from the coil.

2.2.1 Possible Rectifier Architectures

Several rectifier architectures were considered for the control circuitry power supply, and are shown in Figures 2-24, 2-25, and 2-26. For a V_{dd} of 3.5 V, the first rectifier requires peak V_{oc} from the coil to be approximately 4 volts, the last two require approximately 2 volts.




Figure 2-24: Full Bridge Rectifier

Aside from the required open-circuit voltage, the primary difference between these three rectifiers is the direct connection to the coil. The full bridge rectifier in Figure 2-24 has no direct connection to the AC side, meaning that at any time, one side of the AC source can be connected to DC power and one side to ground, or both sides of the AC source can be floating within the DC rails, held only by back-biased diodes.




Figure 2-25: Voltage Doubling Rectifier




Figure 2-26: Dual Half-Wave Rectifier

The doubler [18] in Figure 2-25 has one node of the AC source tied to DC ground, and has one node oscillating to $\pm V$ and one node oscillating between V_{dd} and ground. The dual half-wave rectifier in Figure 2-26 has one side of the AC source tied to the midpoint between the DC rails, and has one node oscillating between V_{dd} and ground.

The control circuitry will include comparators monitoring the AC source for both clock extraction and for synchronous rectification to charge the capacitor bank. In the full bridge rectifier, any comparators looking at one or both sides of the AC source must be able to function perfectly across their whole power supply range, which is difficult to do at low power. The other two rectifiers give a direct connection to the AC source, which makes comparisons easier. The dual half-wave rectifier has its connection to the AC source at mid-rail, which makes the comparison far easier. Furthermore, this mid-rail voltage serves as a convenient current return node for the electrodes, making the dual half-wave rectifier a very attractive option.

2.2.2 Synchronous Rectifier

The AC voltage on the coil must be rectified not only for the V_{dd} power supply for the chip, but also for the capacitor bank which will drive the electrodes, consuming far more power. Since this rectifier will handle the majority of the power being delivered from the coil, it should be as efficient as possible.




Figure 2-27: Charging a Negative-Voltage Capacitor from the AC Voltage
When the magnitude of the AC voltage exceeds that of the capacitor voltage, the rectifier switch is turned on for a short time, charging the capacitor.

A synchronous rectifier allows control of the turn-on and turn-off points, allowing us to limit the voltage drop across the rectifier MOSFET switch. Figure 2-27 shows the AC waveform and the negative voltage on one of the capacitors. The rectifier turns on when the magnitude of the AC voltage exceeds the magnitude of the capacitor voltage, and stays on for a short time, during which charge is delivered to the capacitor.

Figure 2-28 shows a simple version of a synchronous rectifier. The comparator monitors the AC waveform and the capacitor voltage, and trips the one-shot circuit when the AC voltage exceeds the capacitor voltage. The one-shot circuit turns on the rectifier switch for a short time. The pulse duration of the one shot may be adjusted based on the data to charge for longer times, and therefore less efficiently, during times of heavy electrode loading on this capacitor.

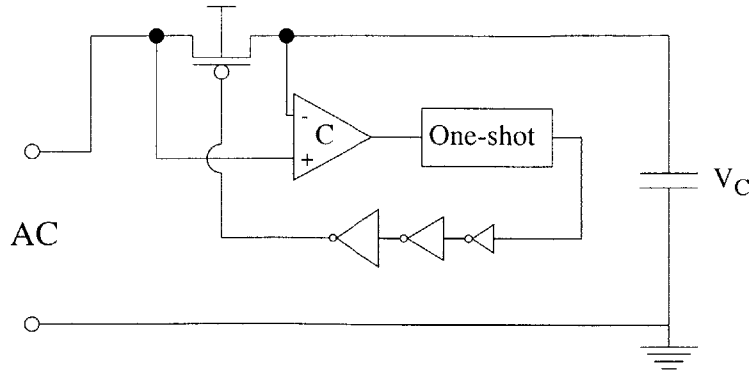


Figure 2-28: Synchronous Rectifier With Comparator and One-Shot
 When the AC voltage exceeds the capacitor voltage, the comparator trips the one-shot, which turns on the switch for a fixed time.

2.2.3 Magnetic Field Calculations

The magnitude of the magnetic field generated by the primary coil must be decided, and several factors (recommended field limits, regularly used field strengths, power losses within the coil, power absorbed in the tissue, etc.) influence the decision.

The American National Standards Institute (ANSI) recommends (in ANSI C95.1) occupational limits of $\frac{16.3}{f} MA/m$, RMS for magnetic fields in the range of frequencies of interest (100 KHz to 100 MHz) [7], though existing medical devices which employ magnetic fields regularly exceed these levels by a factor of approximately 3 to 8 [9]. Furthermore, FCC regulations limit electromagnetic interference (EMI) at some distance from the device, except at prescribed frequency bands set aside for industry, scientific and medical (ISM) purposes. The power transmission frequency was chosen to be 125 KHz because it falls within the ISM bands, allows reasonably strong ANSI field limits (184 A/m, peak), and allows any comparators monitoring the AC waveform to operate relatively slowly, and therefore at low power.

The magnitude of the AC magnetic field along the axis through the center of the primary coil, shown in Figure 2-29, is [11]:

$$H = \frac{N_p I}{2} \frac{r_p^2}{(z^2 + r_p^2)^{\frac{3}{2}}} \quad (2.4)$$




Figure 2-29: Magnetic Field Intensity vs. Axial Displacement From Primary
 Note that, for this 38 mm diameter primary, the field 3 cm from the coil, which corresponds to the posterior retina for an epi-retinal coil, is 15% of that at the center of the coil. The field 1 cm from the coil, which corresponds to a coil in the lens capsule, under the conjunctiva, or on the temporal side of the eye, is 70% of that at the center of the coil.

where H is the magnetic field intensity at a distance z along the axis of a primary coil of average radius r_p , with N_p turns, carrying current I . The primary coil radius is assumed not to vary significantly from r_p .

Off-axis Magnetic Field

Equation 2.4 gives a simple expression for the field along the axis of the primary coil. In the calculations to follow, we would like to assume that this axial expression is the field strength throughout the secondary coil. In order to verify this, the true field strength is calculated numerically as a function of radial displacement. It is assumed that the primary and secondary coils are coaxial and separated by a fixed distance, and that the field is generated in free space, and is unchanged by the presence of the

secondary coil.

The Biot-Savart expression for the ϕ -directed magnetic field, $d\vec{H}$, at distance r and angle θ from a small piece of wire of length ds carrying current i is:

$$d\vec{H} = \frac{1}{4\pi} \frac{i d\vec{s} \times \hat{r}}{r^2} = \frac{1}{4\pi} \frac{i ds \sin(\theta)}{r^2} \hat{\phi} \quad (2.5)$$

This expression is deceptively simple, but calculating the angles can be tricky. Assume that our primary coil has N_p loops of wire carrying current i , and that all loops are coplanar and lie within $r_p \pm \delta$ from the center of the coil, with $\delta \ll r_p$. We wish to calculate the magnetic flux through our secondary coil (coaxial to the primary and separated by distance z). To do that, we require the component of the magnetic field that is normal to the disc enclosed by the secondary for every point on that disc, from every point on the primary coil.




Figure 2-30: Coil Alignment (L) and Angle Calculations (R)
Solid lines are in the plane of the primary coil, while dashed lines are out of the plane.

Figure 2-30 shows the coil alignment on the left and the angle calculations on the right. The figure on the right shows the primary coil and the point in space, P , which falls within the disc of the secondary coil. The measurements of point P with respect to the center of the secondary are projected onto the plane of the primary coil for convenience. Point P sits at a distance r_1 ($r_1 < r_s$) from the center of the

secondary, at angle α , as shown. The point at which radius r_p intersects the primary coil is the segment of the coil, $d\vec{s}$, which is being examined. From Equation 2.5, we need expressions for r^2 and $\sin(\theta)$. As can be seen in the figure, $r^2 = d^2 + z^2$ and $d^2 = x^2 + y^2$, while $x = r_1 \sin(\alpha)$ and $y = r_p - r_1 \cos(\alpha)$. So:

$$r^2 = z^2 + r_1^2 + r_p^2 - 2r_p r_1 \cos(\alpha) \quad (2.6)$$

The angle θ may be calculated from the triangle formed by r , r' , and the tangent to the coil, with length x in the triangle. In this triangle, r' and x form a right angle, so:

$$\sin(\theta) = \frac{r'}{r} = \frac{\sqrt{z^2 + y^2}}{r} \quad (2.7)$$

Equation 2.5 can now be solved for $d\vec{H}$ by multiplying by N_p and substituting $r_p d\alpha$ for ds . However, this gives the $\hat{\phi}$ -directed field at that point in space, not the field normal to the disc formed by the secondary coil. We want $d\vec{H}_\perp$, or $d\vec{H} \cos(\beta)$. First note that $d\vec{H}$ must be perpendicular to $d\vec{s}$, so its reflection on the primary coil plane points along segment y , not d . So $\cos(\beta) = y/r'$, and the normal field is:

$$d\vec{H}_\perp = \frac{N_p i}{4\pi} \frac{r_p d\vec{\alpha} \sin(\theta)}{r^2} \cos(\beta) \quad (2.8)$$

The perpendicular fields at point P are numerically integrated from every point on the primary coil, and then calculated for every point P along a radius of the secondary coil. Figure 2-31 shows the result. For the same values, Equation 2.4 gives 501.5 A/m at the center of the secondary, which agrees. Also, notice the vertical scale in the figure. The mean field strength through the secondary coil is 494.5 A/m , so using Equation 2.4 for the whole secondary gives an error of only 1.4%.

The MATLAB script which does the numerical integration and generates Figure 2-31 is included in Appendix Section A.3.




Figure 2-31: Magnetic Field Intensity vs. Radial Displacement From Center of Secondary

2.2.4 Magnetic Field-Induced Heating

In addition to providing power to the implant, the magnetic fields also deliver power to the eye itself in the form of induced heating. This heating is potentially damaging to the retinal tissue, and must be considered in deciding the field strength for the implant.

The applied magnetic field induces electric fields and currents in the tissue, causing Joule heating. Some approximations can make the calculation of this power quite simple. Figure 2-32 shows the simplified model. The magnetic field from the secondary coil is assumed to be \hat{z} directed, uniform and of value equal to that at the center of the coil. We know by examining Figure 2-29 that this is a very conservative approximation, since the field strength will fall off with axial displacement into the eye.

Faraday's Law for this case is:

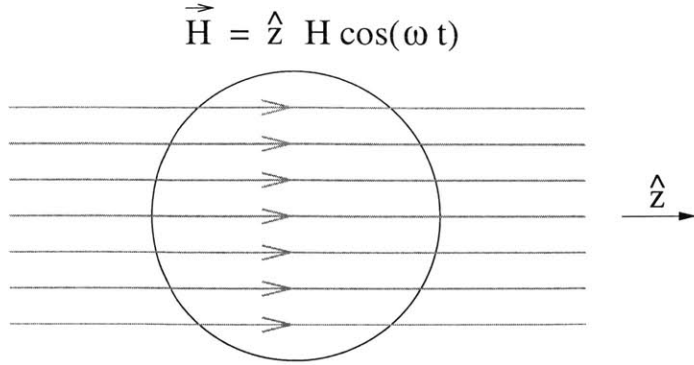


Figure 2-32: Magnetic Fields in the Eye

$$\vec{\nabla} \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} = \hat{z} \mu_0 \omega H \sin(\omega t) \quad (2.9)$$

By symmetry, \vec{E} cannot be z dependent or ϕ dependent. By the right-hand rule, \vec{E} is $\hat{\phi}$ directed. If the curl of \vec{E} is \hat{z} directed, then \vec{E} is $\hat{\phi}$ directed and r dependent:

$$\hat{z} \frac{1}{r} \frac{\partial(rE_\phi)}{\partial r} = \hat{z} \mu_0 \omega H \sin(\omega t) \quad (2.10)$$

The solution to Equation 2.10 is:

$$E_\phi = E_0 r \quad (2.11)$$

$$2E_0 = \mu_0 \omega H \sin(\omega t) \quad (2.12)$$

$$\vec{E} = \hat{\phi} \frac{\mu_0 \omega H}{2} r \sin(\omega t) \quad (2.13)$$

Now that we have the induced electric field, it is straightforward to find the Joule heating power from the field. Figure 2-33 shows the definitions for the volume integration required to calculate the power.

The power can be calculated from the product of the voltage and current:




Figure 2-33: Calculating the Power Dissipated by the Electric field

$$\vec{J} = \sigma \vec{E} \quad (2.14)$$

$$I = \vec{J} \cdot d\vec{A} \quad V = \vec{E} \cdot d\vec{s} \quad (2.15)$$

$$P = \sigma E^2 \cdot d(\text{Volume}) \quad (2.16)$$

$$P_{eye}(t) = \int_{-r_{eye}}^{r_{eye}} dz \int_0^{r'} dr \int_0^{2\pi} r d\phi \sigma E(r)^2 \quad (2.17)$$

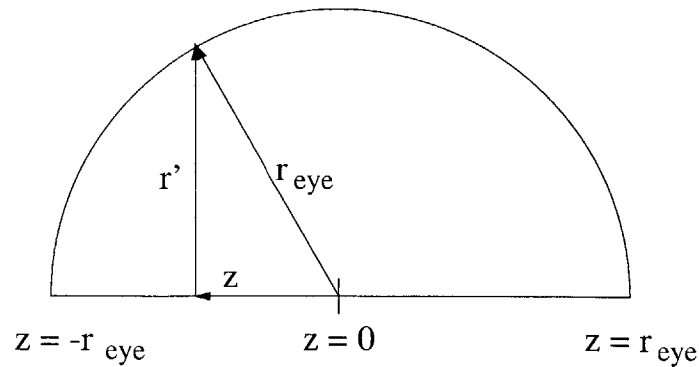


Figure 2-34: Definitions for power volume integral

Figure 2-34 shows the definitions for the limits in Equation 2.17. From the figure, it can be seen that:

$$r' = \sqrt{r_{eye}^2 - z^2} \quad (2.18)$$

Which gives a relatively simple integration yielding:

$$P_{eye}(t) = \frac{2}{15} \pi \sigma \mu_0^2 \omega^2 H^2 \sin^2(\omega t) r_{eye}^5 \quad (2.19)$$

$$\bar{P}_{eye} = \frac{1}{15} \pi \sigma \mu_0^2 \omega^2 H^2 r_{eye}^5 \quad (2.20)$$

The eye of an average adult is 23.5 mm in diameter, so for the suggested field strength in [7] for 125 kHz , the average field-induced power burned in the eye is $2.8\ \mu\text{W}$. For 8 times the recommended ANSI limit, the field-induced power is $178\ \mu\text{W}$. As previously stated, these are thought to be conservative estimates, and the true magnetically induced power losses should be lower, but there are also magnetically induced losses throughout the rest of the head, and there may be additional conductive power losses due to the electric field of the wires in the primary coil. So while these calculations are helpful in getting a feel for the non-circuit power losses, they should not be relied upon without further study.

2.2.5 Secondary Coil Parameters

The secondary coil power delivery limits presented in Sections 1.3.1 and 1.3.2 require explanation. We can assume that the coupling coefficient between the primary and secondary coils is fairly small, so the secondary provides negligible loading to the primary. This means we can turn on the transmitter circuit and assume that it generates the predicted AC magnetic fields, regardless of the existence of the secondary.

For an impedance matched load on the coil, the average power to the load (ignoring non-Ohmic losses and higher-order Ohmic losses in the coil, such as skin effect, proximity effect, and increased resistive losses due to resonance between the inductance and unmodeled parasitic capacitance) would be:

$$\bar{P}_{max} = \frac{1}{8} \frac{V_{oc}^2}{R_{coil}} \quad (2.21)$$

$$R_{coil} = \frac{\rho L}{A_{wire}} = \frac{(\rho)(2\pi\bar{r}N)}{\left(\frac{\alpha A}{N}\right)} = \frac{2\pi\rho\bar{r}N^2}{\alpha A} \quad (2.22)$$

where V_{oc} is the peak (not rms) value of the open-circuit voltage of the coil, R_{coil} is the series resistance of the coil, ρ is the resistivity of the metal in the coil, N is the number of turns in the secondary coil, \bar{r} is the average radius of the coil, α is the “packing coefficient” of the wire, or the area of conductor divided by A , the total cross-sectional area through the implant (see Figure 2-35). However, as will be shown later in this section, impedance matching does not lead to minimal power dissipation inside the eye.




Figure 2-35: Secondary Coil Cross-Section

Shown are the definitions of the variables pertaining to coil size for a typical secondary coil. The packing coefficient, α , is the ratio of conductor area to A : $\alpha = \frac{N\pi r_{wire}^2}{A}$

Assuming that the magnetic field strength throughout the secondary coil is equal to the field strength at its center (as verified above), the open-circuit voltage generated is:

$$V_{oc} = \frac{d\lambda}{dt} \approx \omega\mu_0 H \bar{A}_H N \approx 2\pi^2 f \mu_0 H \bar{r}^2 N \quad (2.23)$$

Now we have all of the elements of Equation 2.21, but one approximation makes the math far easier [38]:

$$\bar{r}^2 \approx \bar{r}^2 \quad (2.24)$$

$$P \approx \pi^3 (\mu_0 f H)^2 \frac{\alpha A}{4\rho} \bar{r}^3 \quad (2.25)$$

From Figure 2-35 we can see that an upper packing bound is obtained from $r_{wire} = 0.5\sqrt{A/N}$ so $\alpha = \pi/4$. This does not take into account wire insulation and packing nonidealities, which will be significant. In reality, α is some constant which expresses the packing density of the wiring, with weak dependence on wire size and number of turns.

Note then, that, neglecting the influence of wire size on packing factor, the extracted power in Equation 2.25 is independent of the wire thickness and number of turns. If the ANSI standard described in Section 2.2.3 is taken as a hard limit, then the fH product in Equation 2.25 is a constant, and the power depends only on physical properties and dimensions of the coil in this model.

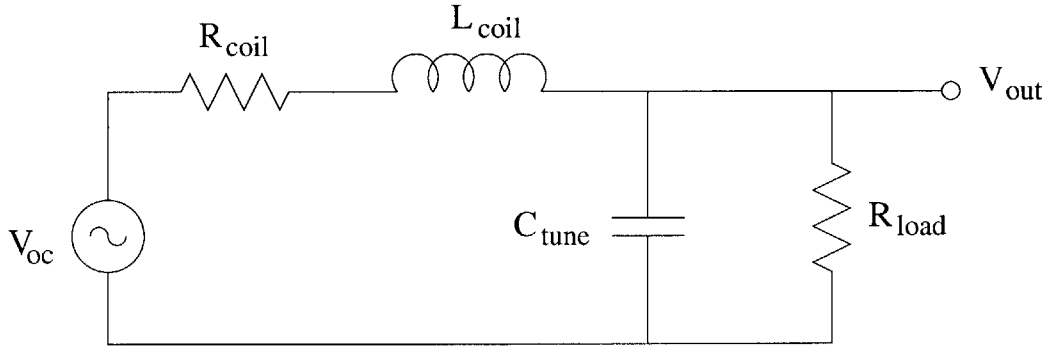


Figure 2-36: Simplified Coil Model

Assume a very simple model of the coil and electrode load, as in Figure 2-36. The shunt capacitance boosts V_{out} by resonance, at a cost of increased power dissipated in R_{coil} . V_{oc} is defined above in Equation 2.23, R_{coil} is defined above in Equation 2.22, and L_{coil} is defined below in Equation 2.26. The capacitance is whatever is necessary to make V_{out} equal to the target value of just over 2 volts. R_{load} and C_{tune} will be discussed further in Section 2.2.6.

$$L_{coil} = \frac{\lambda}{I} \approx \frac{\pi \bar{r}^2 N \mu_0 H}{I} \approx \frac{\pi \mu_0 N^2 \bar{r}^4}{2(t^2 + \bar{r}^2)^{3/2}} \approx \frac{\pi \mu_0 N^2 \bar{r}}{2} \quad (2.26)$$

The last approximation in L_{coil} is valid because the thickness is significantly less

than the average radius. However, as described in Section 3.1.1, this equation is far less accurate than empirical inductance calculation formulae.

Coil Placement and Related Constraints

In Figure 1-8, four coil locations are examined. The primary difference among these locations is the size restriction, and therefore limit on available power, at each location. Table 2.2 examines the size constraints (outer diameter, inner diameter, thickness), dictated by biological constraints, for each location. In any location where inner diameter is not constrained, it is set to 0.4 times the outer diameter. Any turns wrapped inside this limit decrease the maximum available power by adding significant resistance, but little additional flux linkage. The table also lists the maximum possible power (extractable with a matched load) from Equation 2.25.

Location	Outer Diameter	Inner Diameter	Thickness	Max. Power
retina	10 <i>mm</i>	4 <i>mm</i>	100 μ <i>m</i>	2.3 <i>mW</i>
anterior	6.5 <i>mm</i>	2.6 <i>mm</i>	2 <i>mm</i>	10.9 <i>mW</i>
conjunctiva	16 <i>mm</i>	13 <i>mm</i>	1.5 <i>mm</i>	165 <i>mW</i>
temporal	12 <i>mm</i>	6 <i>mm</i>	2 <i>mm</i>	105 <i>mW</i>

Table 2.2: Coil Parameters at Various Locations

The thickness and outer and inner diameters allowed for a coil at each location are listed in this table. Also listed is the maximum possible power to a matched load at the ANSI magnetic field limit.

Non-matched Loading

As alluded to above, the optimal loading for this application is not a matched load. The power dissipation for the implant is relatively fixed at 10 *mW* (conservative estimate for a 4-step capacitor bank, allowing for some inefficiencies and overhead power loss), and any extra power burned in the eye should be minimized. Matching the load to the coil resistance means that the latter burns 10 *mW*, and that, in addition to any heating power of the applied magnetic field, is waste power, and potentially damaging to the eye.

A much better solution is to underload the coil, extracting the necessary power for the implant while burning a small fraction of that power in the coil resistance. This requires a coil with the capability of delivering far more power than is needed. The conjunctival or temporal coils listed in Table 2.2 and shown in Figure 1-8 meet that requirement. As stated in Section 1.3.2, the temporal location also allows easy surgical access and so was chosen as the secondary coil location.

2.2.6 Coil Resonance and Loading

If we set a minimum required value for V_{out} in Figure 2-36 and a minimum required value for the corresponding P_{out} into the load resistance, we can vary other parameters (magnetic field, number of turns, tuning capacitance, load) to minimize the wasted power. This method becomes much easier by first converting the parallel load RC to a series RC, then by using the geometric methods described in [41], [40] and summarized below.

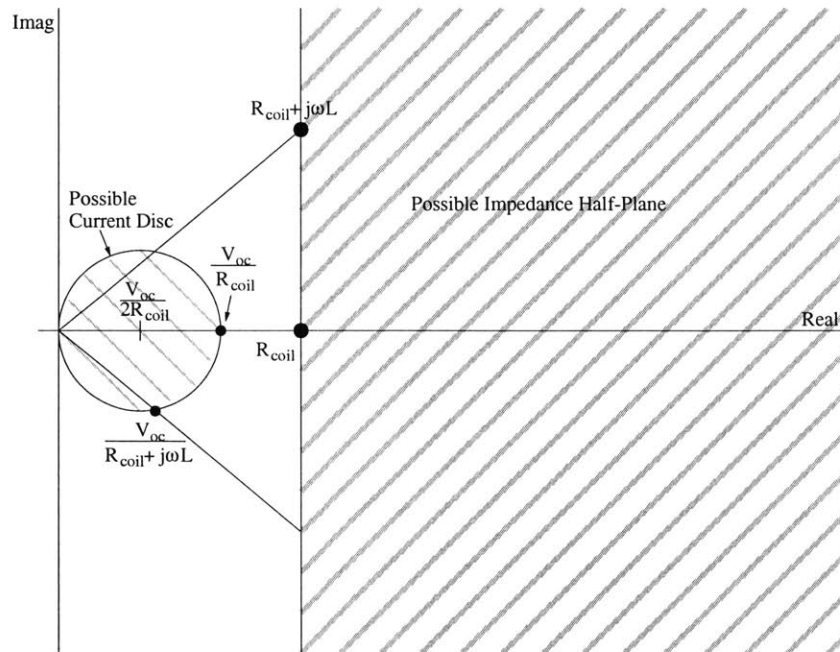


Figure 2-37: Impedance to Current Mapping

The source impedance shown in Figure 2-36 is plotted, and all possible load impedances added to it are shown in the possible impedance half-plane. These impedances are transformed into the superimposed current plane, as described in [40].

The complex source and load impedances in Figure 2-36 are transformed into the complex current plane (assuming a given V_{oc}), as shown in Figure 2-37. Since the real part of the total impedance must be at least as large as the series coil resistance, and the imaginary part could be anything, the impedance is a half-plane to the right of the coil resistance. This half-plane maps to the inside of a circle in the complex current plane, showing all possible currents from V_{oc} . The center of this circle is the matched load point, where maximum power is delivered to the load. The perimeter of the circle delivers zero power to the load. A circle concentric to the current circle has as its perimeter the current which gives the minimum necessary output power. The voltage constraint is introduced by a circle with its center on the current circle at the point of short-circuit source current and radius $V_{min}/\sqrt{R_{coil}^2 + \omega^2 L^2}$. The perimeter of this circle is the current that gives the minimum output voltage, and points outside the circle give greater voltage. The power and voltage constraint figures are combined and reproduced in Figure 2-38.




Figure 2-38: Minimum Power and Voltage Geometry

On the complex plane from Figure 2-37, the minimum power and voltage requirements are superimposed. This figure has been adjusted so that the leftmost intersection of the voltage and power discs lies on the real axis, yielding the optimal operating point.




Figure 2-39: Geometric View of Coil Loading

This figure shows the possible current disc at left in blue, passing through the origin, possible power disc in red, just to the right of the current disc, and insufficient voltage disc in green, concave to the left. Dotted red circles and dashed green circles show 50% above and below the minimum values, and dashed red vertical lines show power values of 10, 15, and 20 mW.

In the complex current plane, for given V_{oc} , power is simply proportional to the real part of the current, i.e., the x axis. So the point of minimum power is the farthest left intersection of the real axis, voltage circle, and power circle. Figure 2-39 shows this geometric picture for $P_{out} = 10mW$ and $V_{out} = 2V$. The operating point uses parallel load resistance and capacitance values of 200Ω and $820pF$, and uses $11.6mW$ of total power, and requires only 3.5 times the ANSI field recommendation. Appendix Section A.4 includes the MATLAB script which generated Figure 2-39.

While this method gives a nice overview of the resonance and loading issue, it has some drawbacks in reality. The load will not be linear and cannot easily be made to look linear. The power supply rectifier will conduct a spike of current at the positive and negative peaks of the AC voltage waveform. Those spikes will be

flattened and spread somewhat by the AC side inductance, but the rectifier will look far from linear. The synchronous rectifier that will charge the capacitor bank will also draw non-sinusoidal currents. As the AC voltage exceeds the voltage of the capacitor being charged, the current draw will increase, then will stop when the rectifier switch is turned off. For these reasons, the tuning capacitance in Figure 2-36 will be made larger, sliding the intersection point in Figure 2-39 further up the power circle. Notice that the power consumption does not increase greatly for small excursions from the real axis along the power circle, but has the potential to increase drastically for larger excursions. Increasing this capacitance desensitizes the output voltage to nonlinearities in the loading, but has a cost in total consumed power.

The following three chapters describe the circuit implementation of the three main components of this system.

Chapter 3

Coupled Coils

This chapter continues the issues addressed in Chapter 2 regarding coupled coil power transmission to the implant. The two primary concerns are power losses in the secondary coil and ensuring that the implant is supplied with sufficient power at the proper voltages. The primary coil and transmitter are of secondary concern, and are designed to be functional, not necessarily optimal.

3.1 Secondary Coil Design

Since the coil placement and V_{dd} rectifier architecture have been decided in Chapter 2, the design for the secondary coil will now be finalized. The limits given by the surgeon for a coil on the temporal side of the eye are 12 *mm* outer diameter, 6 *mm* inner diameter, and 1 – 2 *mm* thickness, where 1 *mm* is a comfortable value, and 2 *mm* is allowable if that thickness is necessary.

The original design was for a foil wound coil: a thin layer of insulated metal foil wrapped to form a one-layer “pancake” coil with a very good packing coefficient and reduced proximity effects [17]. The coil was designed to be the maximum allowable diameter, and the maximum buildable thickness. To reach the required voltage, the coil would include 60 turns, meaning the combined thickness of metal and insulation in one layer would be 100 μm . As it turns out, this design, shown in Figure 3-1, is very difficult to fabricate, due to the difficulty of making thin ridges of metal 1 *mm*

high, or even $100\ \mu\text{m}$ high. Attempts were made to cut a strip of foil tape ($1\ \text{mm} \times 2\ \text{m}$), spray several coats of conformal insulating coating spray on one side, and wrap the coil around a form by hand. All such attempts failed, yielding coils with lower than expected resistance and very low inductance, presumably due to shorts through cracks in the conformal coating after it was wrapped.




Figure 3-1: Coil Cross-Section

Shown are the definitions of the variables pertaining to secondary coil size. The packing coefficient, α , is the ratio of conductor area to A , or, for conductor and insulator of the same h , $\alpha = \frac{t_{\text{metal}}}{t_{\text{insulation}} + t_{\text{metal}}}$

In the final design, the foil coil was abandoned for a more traditional wire-wound coil. In addition, the surgical constraint of placing a wide annulus on a sphere was considered. One solution is to design the coil to have the dimensions listed above, but to have the same curvature as the eye. A far simpler solution is to thin the annulus down from a width of $3\ \text{mm}$ to less than $2\ \text{mm}$. This is possible because there is far more than enough power available.

The final design for the coil called for 60 turns of 36 AWG Belden 8058 insulated copper magnet wire (copper diameter $127\ \mu\text{m}$, max. outer diameter $147\ \mu\text{m}$, resistance $414.85\ \Omega/\text{Kft}$). The coil has an outer diameter of $12\ \text{mm}$, an inner diameter of $8.6\ \text{mm}$, a mean radius of $5.1\ \text{mm}$, and a total thickness of $700\ \mu\text{m}$. This allows 5 turns per radial layer, and 12 radial layers. Of course, this takes a very optimistic view of the precision of wire placement in a hand-wound coil. Section 3.2 describes the construction of the secondary coil.

3.1.1 Calculated Inductance

Equation 2.26 gives the coil inductance as calculated approximately from first principles. Unfortunately, this equation gives an error in excess of 30 % for this coil. This is because the equation assumes that the entire field within the coil takes the value of the field at the coil center. In fact, in the plane of the coil, the field grows monotonically with radial distance from the coil center, and grows quite dramatically near the wires of the coil. The inductance of a coil can only be accurately calculated with empirical formulae. [10] Equation 2.26 predicts $36.24 \mu H$, while the empirical formula predicts $54.03 \mu H$. In Section 3.2, these numbers are recalculated for the actual coil dimensions, and the empirical formula comes much closer to the measured value.

3.1.2 Calculated Resistance

Copper wire of 36 AWG has resistance measuring $414 \frac{\Omega}{Kft}$. Since the coil, as designed, uses $1.92 m$ of wire, that amounts to 2.6Ω , not including skin effect and proximity effect. The skin depth at $125 KHz$ is $\delta = \sqrt{\frac{2}{\omega\mu\sigma}}$, or $186 \mu m$, for $\sigma = 1/17n\Omega m$ and μ the same as vacuum. Since the skin depth is greater than the diameter of the wire, the skin effect is ignored. However, skin effect may be a concern for higher-frequency harmonics, and proximity effect may be a concern. Therefore the impedance of the final constructed coil was carefully measured.

3.2 Secondary Coil Construction

Construction of a coil with 60 turns of wire requires some sort of form or bobbin on which to wrap the turns. A form was designed to fit the turns as described above, with a little bit of extra room left for errors in hand-winding. Forms were cut from a Delrin rod on a lathe. The outer diameter was machined down to $13 mm$ on the lathe, and a hole approximately $2.2 mm$ was drilled down the center for later attachment of the form to a structure using a nylon 4-40 screw. A groove measuring $900 \mu m$ instead of the planned $700 \mu m$ was cut into the rod to an inner radius of $4.75 mm$

(calculated from the circumference after machining). The forms were cut apart, and the final pieces measured approximately 4 mm in total thickness. The secondary coil was wound on the form by hand. The first layer of windings stayed aligned and flat, but it was impossible to control the placement of all other windings. Some side to side pattern was maintained, but it was far from precise. Still, the 60 windings fit within the groove.






Figure 3-2: The Secondary Coil




Figure 3-3: Secondary Coil Windings

Figure 3-2 shows, at left, a diagram of the coil form measurements, at center, a

drawing of the coil form, and at right, a photograph of the secondary coil. Figure 3-3 shows a magnified view of the secondary coil windings, in which you can see the imperfect wire windings and the grooves in the plastic from machining the Delrin rod to the proper outer diameter.

From the measured secondary coil inner and outer diameters, the inductance and resistance may be recalculated. The mean radius is approximately 5.5 mm , and the inductance, from the empirical formulae in [10], is $60.13\ \mu\text{H}$. The total wire length, including 2 leads, each measuring one circumference, is 2.14 m , and therefore the calculated resistance is $2.91\ \Omega$. The measured secondary coil inductance and resistance at 120 KHz were $58.1\ \mu\text{H}$ and $3.2\ \Omega$, respectively. The DC resistance, measured on a Keithley 197 digital multimeter, was $2.95\ \Omega$.

A better understanding of the secondary coil impedance may be gained by examining the impedance magnitude and phase over a range of frequencies of interest. The secondary coil impedance was measured on an Agilent 4284A precision LCR meter, at frequencies from 1 KHz to 1 MHz , and the results are plotted in Figure 3-4.




Figure 3-4: Secondary Coil Impedance Magnitude and Phase

The extracted resistance and inductance are plotted in Figure 3-5, and show remarkable consistency in the inductance measurement, but a large frequency dependence in the resistance. This effect results from some combination of skin effect and proximity effect, and will cause increased I^2R losses within the coil for the highly distorted switch currents drawn by the rectifier.




Figure 3-5: Secondary Coil Resistance and Inductance

The inductance value is very constant across 3 orders of magnitude, but the resistance suffers from skin depth and/or proximity effects. At 125 KHz , the resistance is 3.20Ω , but the third, fifth, and seventh harmonics see resistances of 5.36Ω , 9.44Ω , and 15.4Ω , respectively.

3.3 Primary Coil Design and Construction

The primary coil design and construction closely paralleled that of the secondary coil. Recall from Equation 2.4 that if the diameter of the primary coil is too small, the field strength drops off quickly with axial displacement. But the coil must not be so large as to become unwieldy. A primary coil diameter of 36 mm was chosen to meet both constraints.

The number of turns was decided by the switch parameters in the class E driver of Section 3.4. From Equation 2.4, with all other parameters fixed, the magnetic field H is determined by NI , the product of the number of turns and the current through the turns in the primary coil. So I can be considered to be proportional to $1/N$. If too few turns are used, the switch will need to carry too much current. But the peak voltage across the coil is proportional to the product of L and I , its inductance and current. Since L is proportional to N^2 , and I is proportional to $1/N$, then V is proportional to N . If too many turns are used, the voltage across the transistor drain will be too high. With the number of turns set to 45, the peak coil current is about 1 A and the peak coil voltage is about 120 V, which yields less than 50 V across the transistor drain.

A Delrin form was machined in the same manner as the form for the secondary coil, and wrapped with 45 turns of 30 AWG Belden 8055 copper magnet wire (254.6 μm copper diameter, 284 μm max. outer diameter, 103.19 $\frac{\Omega}{Kft}$). The final measured outer diameter of the primary coil was approximately 39 mm, and the inner diameter was 35 mm, giving a mean radius of 18.5 mm. Figure 3-6 shows at left, the primary coil form design with intended measurements, at center, a drawing of the primary coil form, and at right, a photograph of the final wound primary coil.




Figure 3-6: The Primary Coil

The designed primary coil wire length and resistance were 5.09 m and $1.72\ \Omega$, respectively, while the final calculated values, including one circumference for each lead, were 5.46 m and $1.85\ \Omega$. The measured resistance was $1.86\ \Omega$.

The inductance calculated from Equation 2.26 using the designed parameters was $71.9\ \mu\text{H}$. Using an empirical formula, the same parameters yield $142.5\ \mu\text{H}$. From the final measured parameters, the empirical formula yields $143.1\ \mu\text{H}$. The measured inductance was $153.5\ \mu\text{H}$.

3.4 Primary Coil Class E Driver

To generate the AC magnetic field which delivers power to the implant, a moderate AC current ($\sim 0.5 - 0.75\text{ A}$) is required through the primary coil, which has a large inductance. This then generates a large back voltage ($\sim 60 - 90\text{ V}$), which requires at least a first-pass design of an intelligent coil driver.

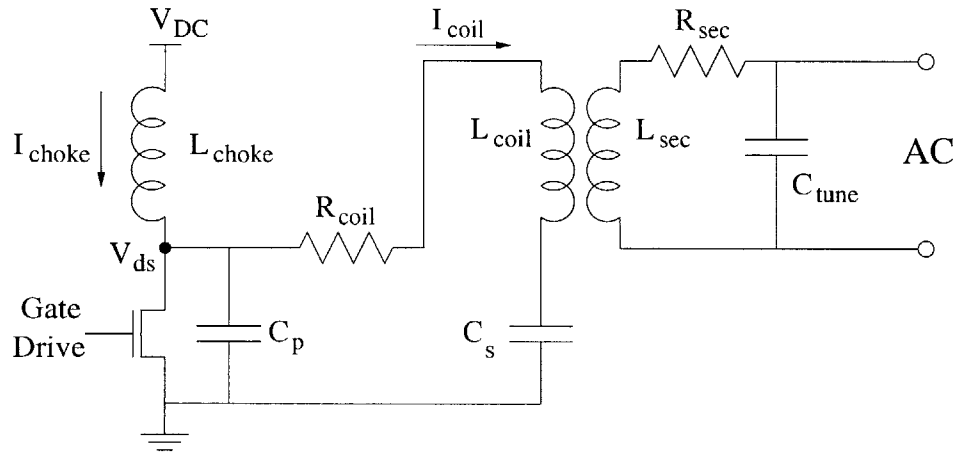


Figure 3-7: Class E Driver

The class E amplifier, shown in Figure 3-7, uses the switch to drive the primary coil at its resonance frequency with the combination of series and parallel capacitors [32]. When the switch is off, the coil inductance resonates with C_p and C_s in series. When the switch is on, L_{coil} resonates only with C_s . Inductor L_{choke} supplies current to make up for energy lost to the resistances and to the secondary load.

Specifically, when the switch is on, the choke fluxes up while the coil current drops to zero, then begins to reverse, flowing into the transistor switch. When the switch

turns off, the coil and choke currents previously flowing into it immediately begin to charge C_p , causing a steep rise in drain voltage. As the drain voltage rises to several times V_{DC} , the choke current decreases, and the coil current drops toward zero. The coil current then reverses, drawing charge from C_p , and bringing the drain voltage down. For true class E operation, the switch should turn on when V_{ds} is equal to zero, ideally with zero slope to minimize the effects of small timing errors. Note that when the switch is off and $\frac{dV_{ds}}{dt}$ is zero, I_{coil} is equal to I_{choke} , not zero as is often reported. For a resonant system of very high Q, I_{choke} becomes much less than the peak value of I_{coil} . The peak current through the coil is nearly linearly related to V_{DC} , so a simple variable power supply can be used to control the magnetic field strength, with V_{DC} ranging from 2 to 6 volts.

If the loading from the secondary is changing significantly, feedback should be used to ensure class E operation [34, 43]. Feedback controllers for the class E driver are being developed by Mariana Markova in our lab, but for this thesis, the class E driver is run open loop, with a simple timer circuit driving the gate. The timing circuit uses a National Semiconductor LM555 timer to generate a simple astable multivibrator, as shown in Figure 3-8. The final class E circuit with 555 timer is shown in the photograph in Figure 3-9.

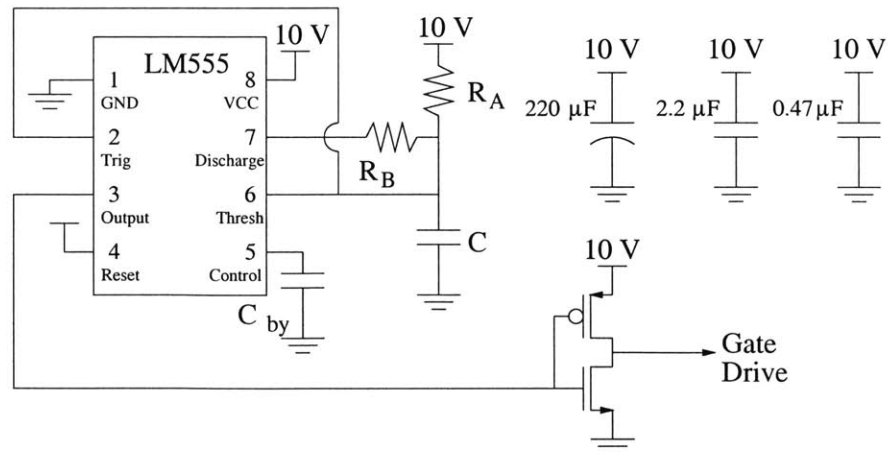


Figure 3-8: Class E Timing Controller

Since the alignment between the primary and secondary coils will vary constantly in an implant attached to the eye, and since the number of electrodes in use, and

Component	Value/Part	Component	Value/Part
C_p	63.9 nF polypropylene	C_s	12 nF polypropylene
R_{coil}	1.86 Ω	L_{coil}	153.5 μH
NFET	ZVN4310A	PMOS	ZVP2106A
Power NFET	IRFZ24N	L_{choke}	500 μH J.W. Miller ferrite hash RF choke
R_{sec}	2.95 Ω	L_{sec}	63.3 μH
C_{tune}	0 - 1 nF		
R_A	414 Ω (1 K Ω trim pot)	R_B	162 Ω (1 K Ω trim pot)
C	12.5 nF ceramic	C_{by}	0.01 μF ceramic

Table 3.1: Class E Driver Components




Figure 3-9: Class E Driver Circuit

therefore power draw from the coil, will also vary, additional feedback should be used to ensure that the secondary coil flux is always the required value. To implement this feedback, information about V_{dd} and the voltage on the bank of capacitors should be sent to the driver via back telemetry, and V_{DC} should be adjusted to maintain the required voltages. Again, for this thesis, this feedback is not implemented.

3.5 Coil Alignment Jig

In order to keep the coils aligned for experiments, a non-magnetic alignment jig was designed and built. The primary coil is affixed to the structure, while the secondary coil is affixed to a positioning block. The secondary is coaxial with the primary, but the positioning block may be moved to any axial displacement by a lead screw. The lead screw runs down a channel cut through the base of the structure, and the positioning block is T-shaped, fitting into the channel.

The structure is made of polycarbonate, the coil forms and positioning block are made of delrin, and the mounting screws, lead screw, and terminating nut are nylon. The lead screw knob is mostly plastic, with a small amount of brass, but it is far from the primary coil, so the metal should not influence the field. The alignment jig is shown at approximately 60% scale in Figure 3-10. A photograph of the jig is shown in Figure 3-11.

3.6 Coil System Testing

The coil coupling was tested with the system shown in Figure 3-12. This circuit uses a dual half-wave rectifier to establish the power supply, and a variable load resistance. These measurements are not meant to carefully model the switching function of the synchronous rectifier, but simply to provide a power drain from the secondary coil. In all of these measurements, C_{tune} was not explicitly added.

In the first test, the secondary coil was held by the coil jig to be coaxial to the primary coil, with a 15 *mm* displacement between their centers. The voltage supply




Figure 3-10: Coil Alignment Jig




Figure 3-11: Photograph of the Coil Alignment Jig

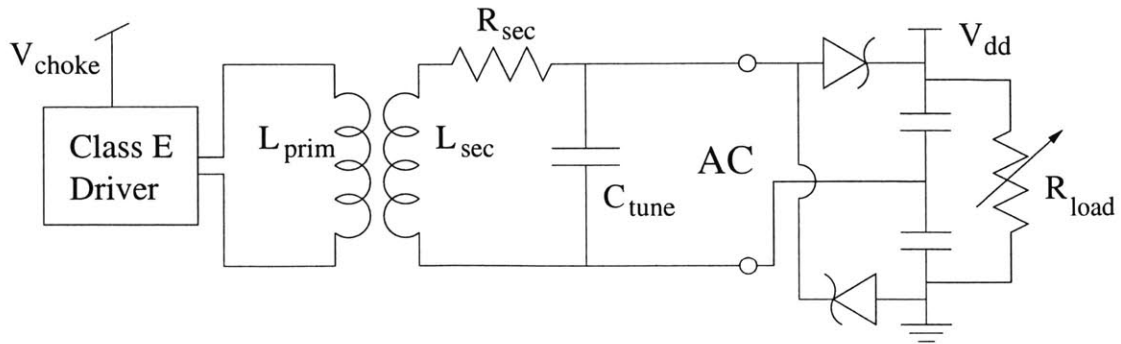


Figure 3-12: Coil Coupling Test Circuit

for the class E choke was stepped from 1 V to 6 V by steps of 0.25 V. This should step the output magnetic field in a monotonic manner. The output field was calculated from the primary coil current, which in turn was calculated from the peak drain voltage on the class E transistor. Since the charge on the drain capacitance is the integral of the coil current,

$$Q = \int I \sin(\omega t) = -\frac{I}{\omega} \cos(\omega t) \quad (3.1)$$

So the drain capacitance charge at the end of 1/4 cycle is I/ω , and the coil current is:

$$I_{coil} = V_{ds} C \omega \quad (3.2)$$

Recall that the magnetic field strength at the center of the primary coil is;

$$H = \frac{NI}{2r_p} = \frac{NV_{ds}C\omega}{2r_p} \quad (3.3)$$

Figure 3-13 shows a plot of magnetic field strength at the center of the primary coil (calculated from the peak drain voltage of the class E switch) as a function of class E choke supply voltage. This measurement was repeated with nine different load resistances on the secondary, from 500 Ω to 94 K Ω . As assumed previously, the primary coil is practically immune to loading on the secondary due to the low coupling coefficient.

Figure 3-14 shows the measured rectified power supply voltage on the secondary side as a function of the same class E choke supply voltage. The nine load resistance values, listed in the figure caption, substantially affected the secondary power supply voltage. For example, a 1 K Ω load represents just over 12 mW of delivered power at the target 3.5 V supply. This requires 5.5 V on the choke, generating 1040 A/m at the center of the primary, and 487 A/m at the center of the secondary, or, respectively, 5.6 and 2.6 times the ANSI suggested limits at our frequency.

Examining the V_{dd} supply voltages in Figure 3-14 more carefully, we may extract a rough linear approximation to the rectifier's Thevenin resistance. Figure 3-15 shows




Figure 3-13: Magnetic Field at Center of Primary
 Note the invariance of the primary field to changes in the secondary load.




Figure 3-14: Rectified Power Supply Voltage vs. Class E Choke Supply
 Measured for nine load resistance values: 500 Ω , 1 K Ω , 1.5 K Ω , 2 K Ω , 2.5 K Ω , 3 K Ω , 4 K Ω , 5 K Ω , 94 K Ω .

V_{dd} plotted against R_{load} for two values of class E choke voltage, 3 and 6 volts. Just above those plots are shown the output voltage of a DC voltage supply with a series source resistance under the same load, calculated with:

$$V_{out} = V_{source} \frac{R_{load}}{R_{source} + R_{load}} \tag{3.4}$$

In Figure 3-15, the plot from the 6 V choke voltage is compared with the plot of a 5.55 V DC voltage source with 370 Ω of series source resistance, while the plot from the 3 V choke voltage is compared with a plot of a 2.65 V source with 390 Ω of resistance. However, if the lightest load (94 K Ω) is ignored, the source resistances drop to 300 Ω . This makes sense, since the small-signal resistance of the diodes making up the rectifier becomes smaller under heavier loading.




Figure 3-15: Equivalent Rectifier Thevenin Resistance Plots

Finally, the axial displacement between the coils was varied, and the results plotted in Figure 3-16. The secondary load was set to the negligible value of 94 K Ω , and curves of secondary supply voltage versus displacement were plotted for the same values of choke voltage used in the previous tests.




Figure 3-16: Rectified Power Supply Voltage vs. Axial Coil Displacement

These tests show that the class E driver and coupled coil system are sufficient to develop the required voltage for the chip power supply and to deliver the required power to drive the electrodes.

Chapter 4

Power System

As discussed in Sections 2.2.1 and 2.2.2, power must be delivered from the secondary coil to storage capacitors. In the case of the V_{dd} rectifier, the voltage is fixed and the power level is low, so a simple Schottky diode rectifier will do. In the case of the multiple capacitor bank or a more general intermediate distribution network, the storage capacitors handle most of the power in the system, so they must receive power from the coil in an efficient manner. This chapter describes in great detail the implementation of these rectifiers, specifically the integrated synchronous rectifier.

4.1 Power Supply Rectifier

The dual half wave rectifier of Figure 2-26 was chosen for the reasons described in Section 2.2.1. The Schottky diode chosen for this rectifier was the SD-103AW in a SOD-123 surface mount package. The power supply consists of two $3.3\ \mu F$ ceramic chip capacitors in an 0805 surface mount package, as shown on the PC board in Figure 4-1. The midpoint of these two capacitors is labeled V_{mid} .

4.2 Capacitor Bank Rectifier

The synchronous rectifier used to charge the capacitor banks was implemented on an integrated circuit due to its complexity. As shown in Figure 2-27, the rectifier




Figure 4-1: Power Supply Rectifier Implementation

This board includes the entire secondary side of the system, but this photograph shows the V_{dd} Schottky rectifier diodes and capacitors.

must turn on a switch between the storage capacitor and the AC waveform as soon as the AC waveform magnitude exceeds the magnitude of the capacitor voltage. The capacitor charges through this switch for a short time, then the switch is turned off. This process is repeated until the capacitor voltage reaches its target reference voltage, and then is not repeated until the capacitor voltage drops below the reference (i.e., as the capacitor supplies charge to the electrodes). Alternatively, the switch may be turned on when the AC voltage exceeds the *reference* (as opposed to capacitor) voltage. This has the same effect if the circuit is operating correctly in the steady state (where $V_{cap} \approx V_{ref}$), but changes the circuit operation if the capacitor voltage falls far below the reference. In this case, the capacitor charges more quickly when the switch turns on, since the voltage difference is higher, but the charging is also less efficient for the same reason.

Two major architectures were examined for this rectifier design, one controlled entirely by timing, and one controlled entirely by voltage. The final design is somewhat of a hybrid, but is much closer to voltage-based rectification.

4.2.1 Timing-Based Rectifier

If the AC waveform on the secondary coil is assumed to have no harmonic distortion and to maintain a constant magnitude, then the voltage of the waveform is known as long as the phase is known. Therefore, to turn on the switch when the AC waveform reaches a certain voltage, it is only necessary to know the AC magnitude and frequency, and to count cycles of a clock that is much faster than, and phase-locked to, the AC waveform.

This architecture avoids the offsets and delays inherent in a continuous-time voltage comparator, but requires a phase-locked loop or delay-locked loop, which adds substantial complexity. In addition, it is very difficult to maintain the AC waveform at a constant magnitude, since any misalignment of the coils changes the coupling. The V_{dd} telemetry feedback method described in Section 3.4 would help reduce these effects, but likely will not have sufficient bandwidth to eliminate them completely. Further, the constant switching of the AC waveform to large capacitors will introduce harmonic distortion, resulting in significant deviations from the expected voltage.

4.2.2 Voltage-Based Rectifier

Since the control of the rectifier switches ultimately depends on the AC voltage and capacitor (or reference) voltage, it would seem that the best rectifier implementation would monitor those voltages. One such implementation could include two continuous-time voltage comparators. One compares the AC waveform and the capacitor voltage, turning the switch on when the AC waveform exceeds the capacitor voltage. The second comparator monitors the capacitor and reference voltages, turning the switch off when the capacitor voltage exceeds the reference.

This architecture has a few drawbacks. It requires two continuous-time comparators which must be fast and accurate, thus consuming significant power. Additionally, these comparators are simply high-gain amplifiers, so if the inputs are sufficiently close together, the output could be at some midway voltage, burning short-circuit power in a downstream digital stage. This would occur in the case of the second comparator,

which monitors two voltages which are practically DC. Another drawback occurs if the capacitor does not reach the reference voltage within one cycle. The AC waveform remains somewhat clamped to the capacitor voltage. This means that the peak of the AC voltage will be greatly reduced, and the V_{dd} rectifier may not turn on, causing the power supply to sag.

4.2.3 Synchronous Rectifier Architecture

The architecture chosen for the rectifier is a hybrid of those described above, and is shown in the block diagram in Figure 4-2. The figure numbers above blocks in this figure refer ahead to where the blocks are expanded in greater detail. Recall from Section 2.1.7 that the current waveform used always begins with a pulse of negative current followed by a charge-balanced pulse of positive current, that four steps are used for each electrode stimulation current phase, and that the electrodes are biased to some anodic voltage, more positive than the electrode return voltage. Since the electrode return is connected to V_{mid} , let us call the anodic bias V_{p1} . So the electrode normally sits at V_{p1} , and during the negative current phase, is switched to four increasingly negative voltages. The first such voltage is V_{mid} itself. The next three voltage sources are capacitors at voltages labeled V_{n1} , V_{n2} , and V_{n3} in sequence. During the positive current phase, the electrode is switched to voltages V_{n1} , V_{mid} , V_{p1} , and V_{p2} in sequence. The electrode is then open-circuited, and a weak current source, described in Section 5.2.3, ensures that it returns to the bias voltage at V_{p1} . This discussion makes clear the need for five storage capacitors, whose voltages are controlled to sit at five reference voltages, two above V_{mid} and three below.

As a functional overview of the system shown in Figure 4-2, a clock is first derived by comparing the AC secondary voltage and V_{mid} . Five reference voltages are generated, and the five capacitor voltages are compared with those on the rising clock edge to find out which capacitors need to be charged. At most one capacitor is charged during a single clock phase; capacitors with voltages above V_{mid} are charged during the positive clock phase when the AC voltage is above V_{mid} , and capacitors with voltages below V_{mid} are charged during the negative clock phase. Once a capacitor

is selected for charging, it is compared with the AC secondary coil voltage. In the case of a “p” capacitor (above V_{mid}), the rectifier switch is turned on when the AC voltage becomes higher than the capacitor voltage. The opposite is true for an “n” capacitor (below V_{mid}). The rectifier switch stays on for a duration determined by control circuitry, depending on how many electrodes are loading the capacitor being charged.

Now let us examine the parts of Figure 4-2 which perform the functions just discussed.

The reference generator block in the upper left contains two different voltage reference circuits and a number of voltage followers. At left is a V_T based reference circuit, which generates voltages nbias, referenced to ground, and pbias, referenced to V_{dd} . However, since the current return electrode will connect to V_{mid} , the five reference voltages must be referenced to V_{mid} . So a current-source based voltage reference generates those voltages, using nbias and pbias as bias voltages for the current sources and V_{mid} as the reference node. These circuits are described in Section 4.3.1. All of these voltages are buffered by voltage followers, described in the same section. Nodes pbias and nbias have MOS capacitors to V_{dd} and ground, respectively, and the five reference voltages have MOS capacitors to V_{mid} (not shown, due to space constraints).

The clock circuitry at the bottom of the figure, described in Section 4.3.3, begins with a clock extractor. This block includes a comparator which is high when the AC voltage is higher than V_{mid} , as well as some clock buffers and a simple one-shot circuit used to generate the clocking signals for the clocked comparators. The one-shot signal is the input for the clock generator block, which generates a pair of non-overlapping clocks, a pair of overlapping clocks, and two other clock signals for the clocked comparators. In addition, the clocked comparators need 2 bias clocks, which normally sit at ground and V_{dd} , but jump for a brief time to nbias and pbias, respectively. This function is provided by the clocked followers.

The clock waveforms are shown in Figure 4-3. As shown in the figure, the main chip clock, clk, is derived directly from the AC waveform and V_{mid} , as described in the text. Signal clk140 is a 140 ns wide one-shot output, and generates all of the

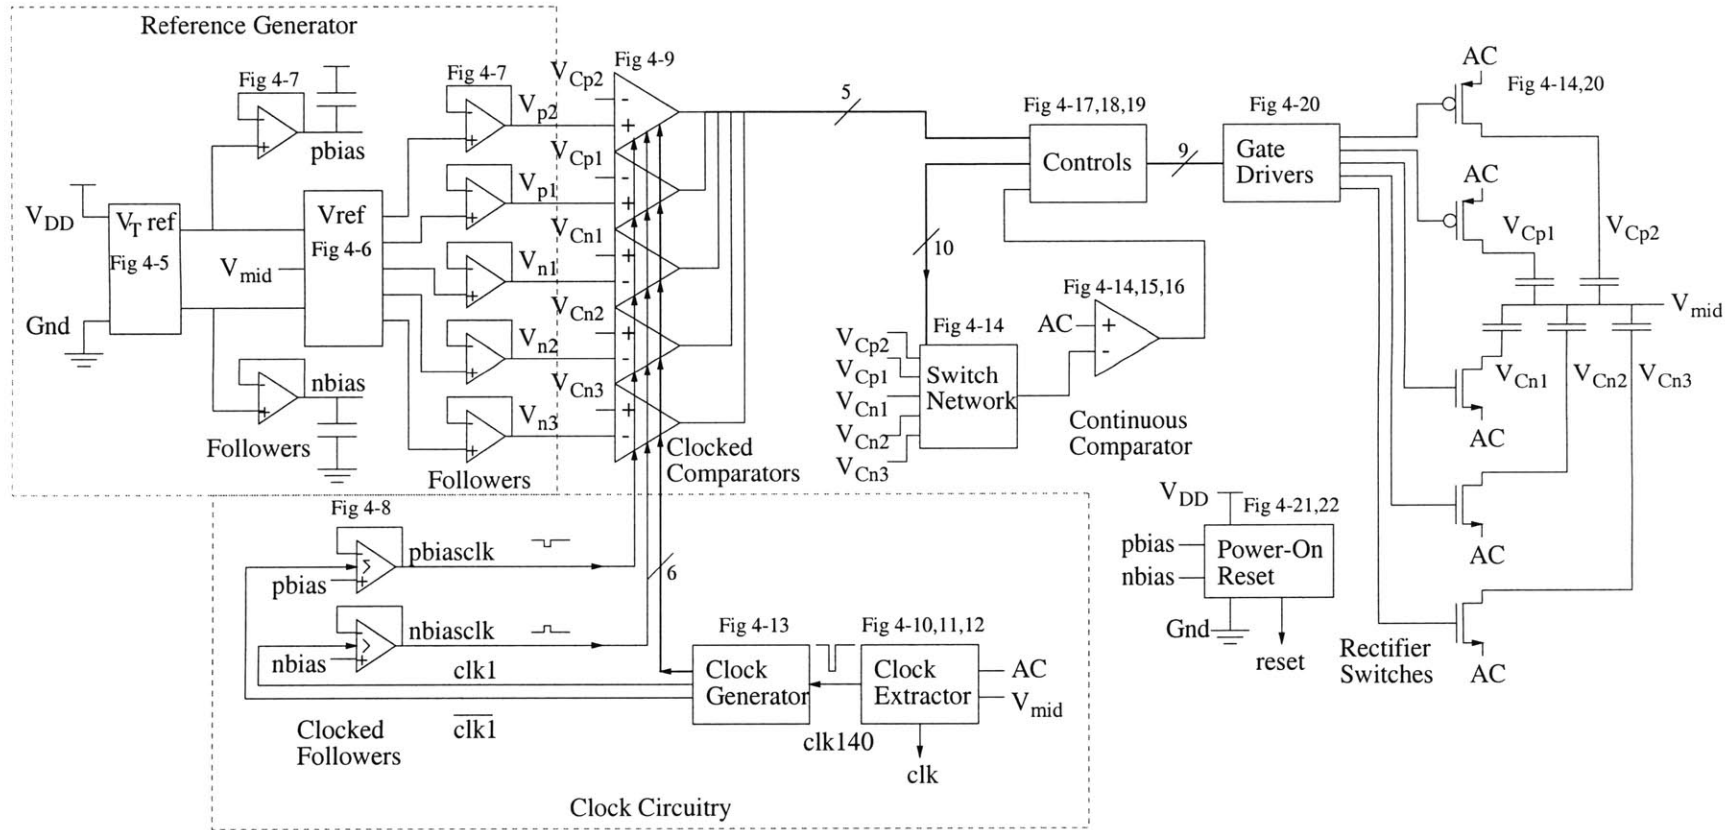


Figure 4-2: Synchronous Rectifier Block Diagram

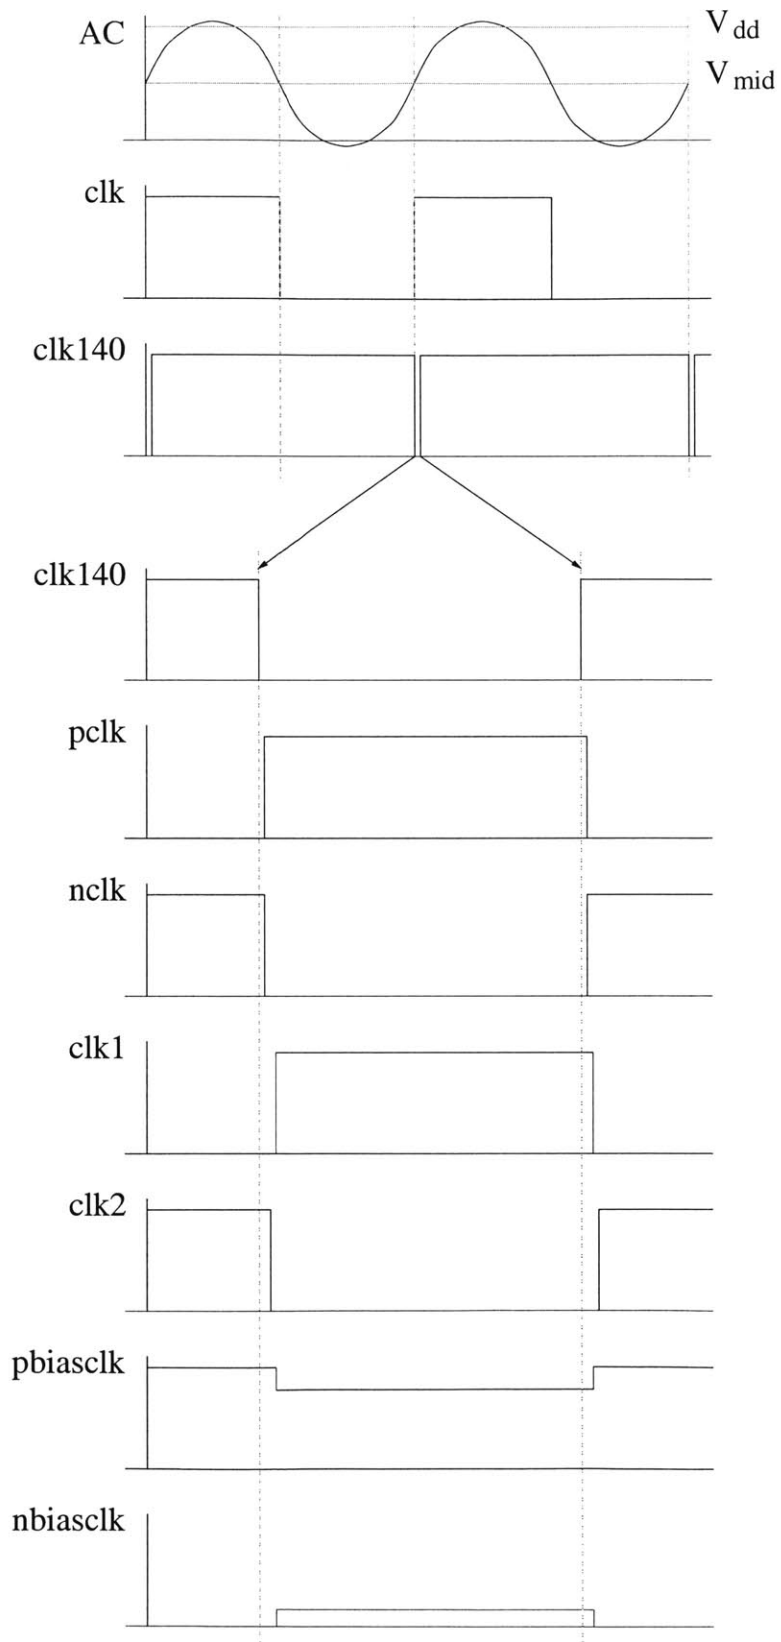


Figure 4-3: Clock Timing Diagrams

clocked comparator signals. In the fourth trace, clk140 is expanded to give a better view of the timing, and that time scale is continued in the lower traces. Signals pclk and nclk are derived from clk140, and are specifically synchronized to change together. Signals clk1 and clk2 are non-overlapping clocks, slightly delayed from pclk and nclk. Signals $\overline{\text{clk1}}$ and $\overline{\text{clk2}}$, not shown, are simply the inverse of clk1 and clk2, and form overlapping clocks.

In the middle of the figure, the five reference voltages are compared with the five final capacitor voltages in a clocked comparator, described in Section 4.3.2, which is clocked at the rising zero crossing of the AC waveform. The clocked comparator outputs are used by the controls block, described in Sections 4.3.5 and 4.3.6, to determine which capacitors need more charge. The core of the synchronous rectifier is the continuous comparator, described in Section 4.3.4, which monitors one of the capacitor voltages at a time and switches when the AC secondary coil voltage exceeds that capacitor voltage. This signal, too, is used by the controls, determining when to turn on the appropriate rectifier switch. The controls turn on the switch, via the gate drivers, described in Section 4.3.7, for a short duration, depending on the number of electrodes loading that capacitor. The rectifier switch delivers charge from the secondary coil to the storage capacitor to maintain its voltage at the appropriate reference voltage.

Digital controls described in Chapter 5 determine which capacitor is charged on any given cycle of the secondary AC voltage. In addition, these controls also determine the duration for which the rectifier is turned on. Figure 4-4 shows an example timing diagram. In this example, capacitor C_{p1} is loaded with a small number of electrodes, C_{n1} is loaded with several electrodes, C_{n2} is heavily loaded, and C_{p2} and C_{n3} are unloaded. During the first positive phase, capacitor C_{p1} is charged for a short time, since it is lightly loaded. Since C_{p2} is unloaded (and is assumed to have been charged previously), it needs no charge, so C_{p1} is also charged on the second positive phase. In the first negative phase, C_{n1} is charged for a longer time, since it is moderately loaded. The second negative phase shows a more complicated charging scheme. C_{n2} is heavily loaded, so it is charged for the whole time the AC voltage is greater than the

capacitor voltage, with a brief reprieve in the middle to ensure that the V_{dd} rectifier turns on.




Figure 4-4: Synchronous Rectifier Control Function

4.3 Capacitor Bank Rectifier Implementation

4.3.1 Voltage Reference

The voltage reference circuit begins with a simple and well-known V_T -dependent reference, as shown in Figure 4-5. This circuit was introduced to me by Luke Theogarajan, and I chose to use it here instead of a bandgap circuit because it is simple and does not require bipolar transistors. In this reference, MN1, in its linear region, serves as a resistor, which has a V_T drop across it. Since the current mirror formed by MP1 and MP2 forces the currents through MN1 and MN2 to be equal, the linear I-V curve from the resistor and the parabolic I-V curve from transistor MN2 intersect at only two stable operating points for the node labeled nbias: at voltage, current pairs (0,0) and $(V_T, V_T/R)$. The network in the right half of the circuit ensures that node nbias reaches the non-zero operating point. At turn-on, node V_{cap} is at ground and the output of the inverter formed by MP3, MP4, and MN6 is V_{dd} . MN5 pulls down on node nbias, starting current flow in the mirror, while MP5 injects a small current into the capacitor and into node nbias via MN4. As V_{cap} exceeds V_T , the inverter turns off the

startup circuitry, leaving nbias and pbias to settle to their desired values. Voltages nbias and pbias will be used throughout the circuitry in this chapter.

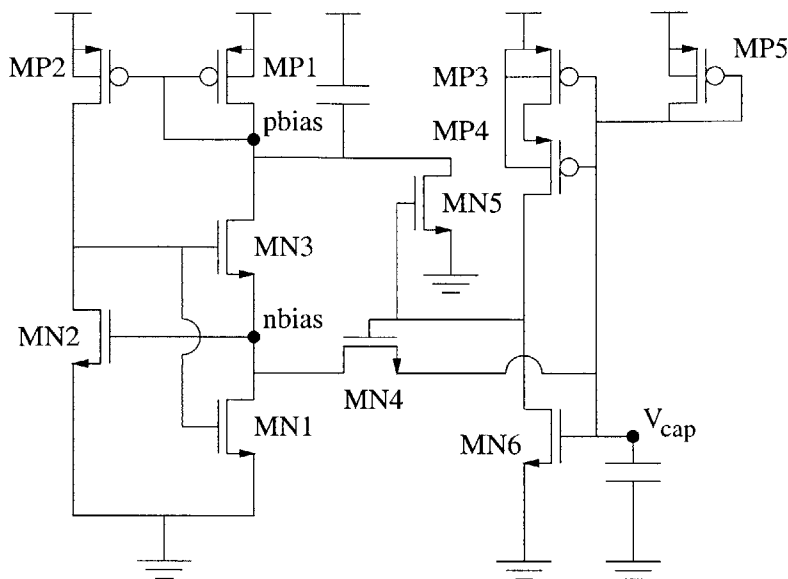


Figure 4-5: V_T -dependent Voltage Reference

The reference voltages for the rectifier are produced by running constant current through a string of resistors, as shown in Figure 4-6. Transistors MP1 and MN1 form the current sources, and three additional sources can be added to adjust the reference voltages. The relative widths of transistors MN1, MN2, MN3, and MN4 are 2, 1, 2, and 4, respectively, and the PMOS widths are the same. This allows between 0 and 7 steps of current to be added to the baseline 2 steps.

In circuit simulations, the power consumed by the entire voltage reference circuit, including the V_T -based reference, ranged from $5.9 \mu W$ to $15.9 \mu W$, depending upon the reference voltage level.

These reference voltages, as well as nodes pbias and nbias, must then be buffered to prevent interference from other circuitry. Figure 4-7 shows the follower circuit used to buffer V_{p1} , V_{p2} , and pbias. It is a well-known circuit using a differential pair, whose output is mirrored to a wide-swing output stage. Voltages below V_{mid} are buffered by the complementary circuit. In circuit simulations, each of these buffer circuits operates with less than $3 \mu W$ of power.

Finally, the clocked comparator described in Section 4.3.2 requires two clocked bias

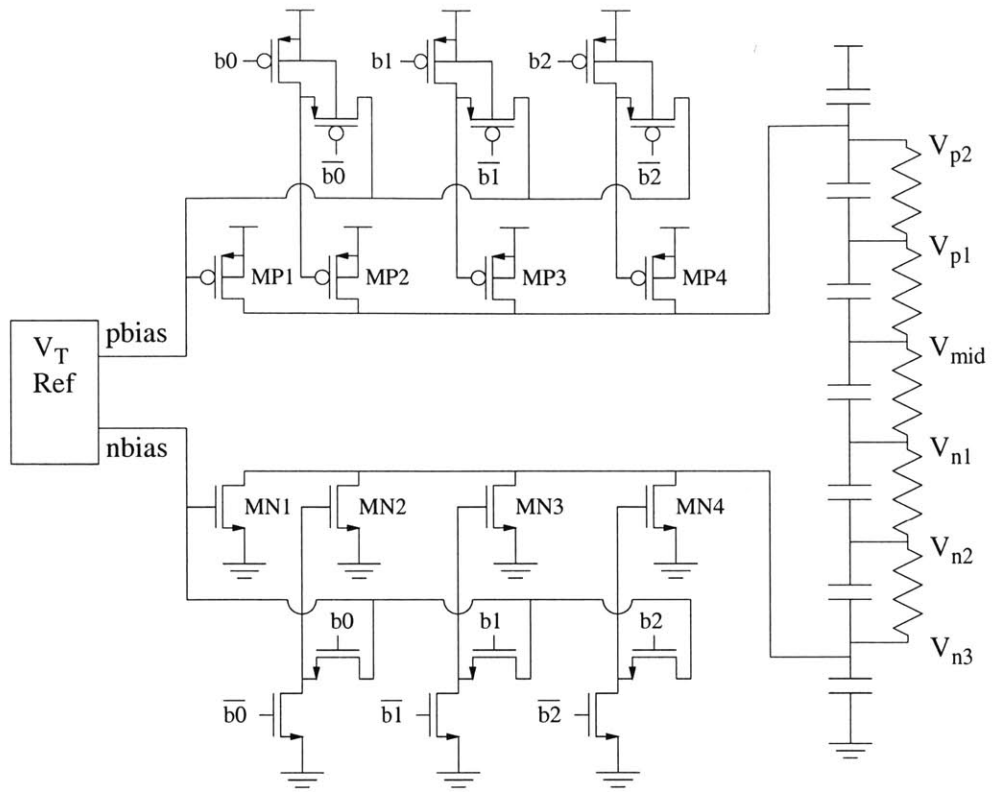


Figure 4-6: Voltage References




Figure 4-7: PMOS Input Voltage Follower

voltages. These voltages must sit at V_{dd} and ground, and when clocked by a pulse, must jump to pbias and nbias, respectively. The circuit in Figure 4-8, shown to me by Luke Theogarajan, supplies clocked pbias, and a complementary circuit supplies clocked nbias. When $\overline{\text{clk1}}$ is high, the NMOS current source for the differential pair is off, while the PMOS pair is turned fully on, bringing V_{out} high. This turns off any PMOS device whose gate is connected to V_{out} . When $\overline{\text{clk1}}$ is low, the PMOS pair forms a mirror and the differential pair receives current from the NMOS current source. The circuit now works as a voltage follower, and V_{out} is equal to pbias. Circuit simulations showed that the pbias and nbias clocked buffers use less than $42 \mu W$ and $35 \mu W$, respectively, or less than $0.75 \mu W$ and $0.7 \mu W$ when the duty cycle of $140 \text{ ns} / 8 \mu s$ is considered. The generation of $\overline{\text{clk1}}$ is described below in Section 4.3.3.

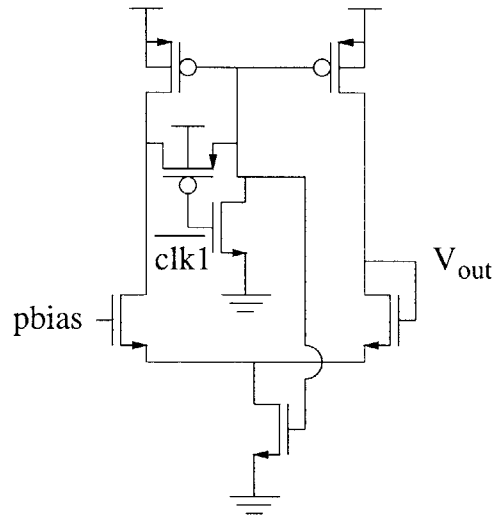


Figure 4-8: Clocked Voltage Follower

4.3.2 Clocked Comparators

The reference voltages are DC values, and the capacitor voltages are being controlled to match them carefully. Therefore the comparator which monitors these two voltages must be able to resolve a small difference between 2 practically DC voltages. In addition, any continuous comparator monitoring the two DC, nearly equal voltages would at some point have its output voltage in the middle region between V_{dd} and ground, causing significant short-circuit current in any digital logic connected to it.

Therefore, a clocked comparator was chosen to monitor these signals.

Figure 4-9 shows the implementation of this clocked comparator. This circuit was developed jointly by Luke Theogarajan and myself. Transistors MN1 and MP1 form one inverter, MN2 and MP2 form another, and their outputs and inputs are cross-coupled. Each inverter has two different current source power supplies. Looking only at the left half-circuit, the inverter (when in its linear region) receives a low, consistent level of bias current from transistors MN3 and MP3, and many times that level of current when transistors MN5 and MP5 are turned on (nCLK high and pCLK low). The “bowtie” symbols are a common representation of a passgate, an NFET and PFET in parallel.

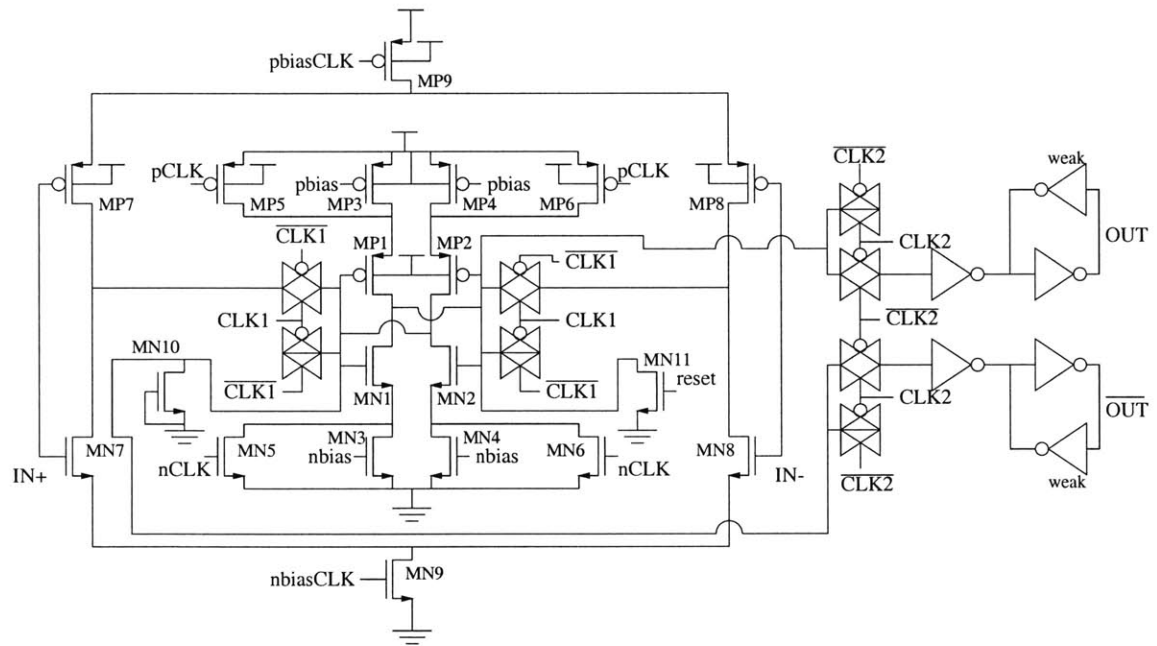


Figure 4-9: Clocked Comparator

The reference voltage and capacitance voltage are connected to IN+ and IN-, such that a high comparator output means that the capacitor needs charge. That is, references V_{p2} and V_{p1} are connected to IN+ on their respective comparators, and capacitor voltages V_{Cp2} and V_{Cp1} are connected to IN-, so that the comparator output is high when the reference voltage is above the capacitor voltage, meaning the capacitor requires more charge. The “n” reference voltages are thus connected to IN-, and the “n” capacitor voltages, to IN+. During comparator operation, the

cross-coupled inverters are weakened (pCLK high, nCLK low), the input passgates are turned on (CLK1 high), the output passgates are turned off (CLK2 low), and the PMOS and NMOS differential pairs are biased with clocked follower circuits like the one shown in Figure 4-8. The input differential pairs take control of the cross-coupled inverters while the output latches hold their previous state. At the end of the sample phase, the inverters are returned to full power (nCLK high, pCLK low), the differential pairs are unpowered (pbiasCLK high, nbiasCLK low), the input passgate turns off and the output passgate turns on (CLK1 low, CLK2 high). The cross-coupled inverters amplify the small differential input and drive their outputs to the rails, and when the output passgates turn on, the state of the main latch is copied to the output latches.

Transistor MN11 allows the main latch to be reset to a known state, while MN10 provides a dummy device to balance any parasitic effects of MN11. The passgates each have a half-sized dummy passgate on the more sensitive side. The dummy passgate is clocked exactly opposite the real passgate, providing first-order charge injection cancellation. The half-sized dummy passgate was added only to one side in order to save layout area, but in retrospect, the security of having solid charge injection immunity would have been worth the additional area. This idea is examined further in Section 7.2.

In circuit simulations, this clocked comparator consumed 11 pC of charge in one 140 ns comparison. With the 3.5 V supply and an 8 μs clock, this amounts to 4.8 μW , and there are five such comparators in the system.

4.3.3 Clock Circuitry

In order to control the clocked comparator, and to clock all of the digital controls for this rectifier and for the electrode switching network to be presented in Chapter 5, a base clock must be generated. This was done with a simple continuous comparator monitoring the AC input power waveform and the midpoint voltage between V_{dd} and ground, named V_{mid} .

Clock Extractor

The comparator used to extract the base clock need not be able to compare over a wide range, since V_{mid} should remain at mid-rail, but the comparator ideally should be perfectly fast, with infinite gain. While these specifications sound impossible, we can take advantage of the fairly sinusoidal shape of the AC input (especially as it crosses V_{mid}) to meet the specifications with a low-power comparator.

Figure 4-10 shows the clock extractor, which uses a predictive circuit before the comparator. This is taken from Alex MeVay's work [22], with modifications to make the circuit bidirectional. Functionally, this circuit compensates for the bandwidth and gain limitations of its comparator by calculating the *slope* of the incoming AC waveform and subtracting voltage proportional to the slope from the DC value. Thus, a quickly rising AC signal will reduce the DC voltage seen by the comparator. The equal voltage point seen by the comparator will occur earlier in time than the true equal voltage point, and if this is carefully tuned, the comparator output will switch exactly when the real voltages are equal.

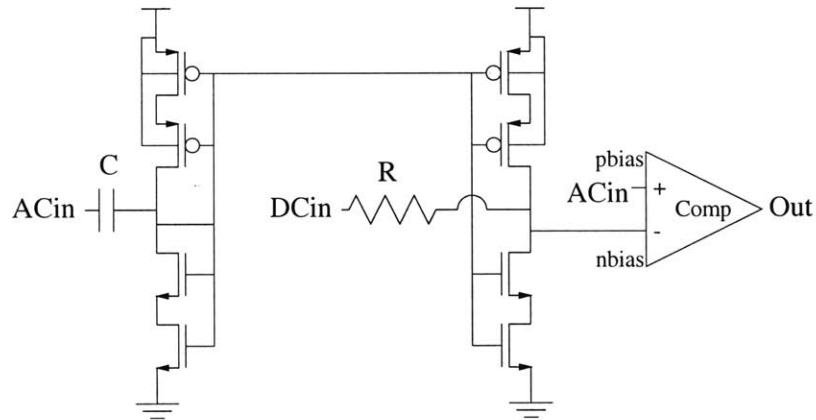


Figure 4-10: Clock Extractor

The transistors in Figure 4-10 form PMOS and NMOS current mirrors. The input capacitor current is the derivative of the AC voltage, and that current feeds into the mirror. The mirror current passes through the resistor, yielding a voltage drop between the DC voltage and the comparator input. A rising AC voltage generates a current into the NMOS mirror, causing current to be drawn from the DC input

through the resistor into the other side of the NMOS mirror, reducing the effective DC voltage. The 500 fF capacitance was made with a poly1-poly2 capacitor, and the $125\text{ K}\Omega$ resistance was made with an nwell resistor.

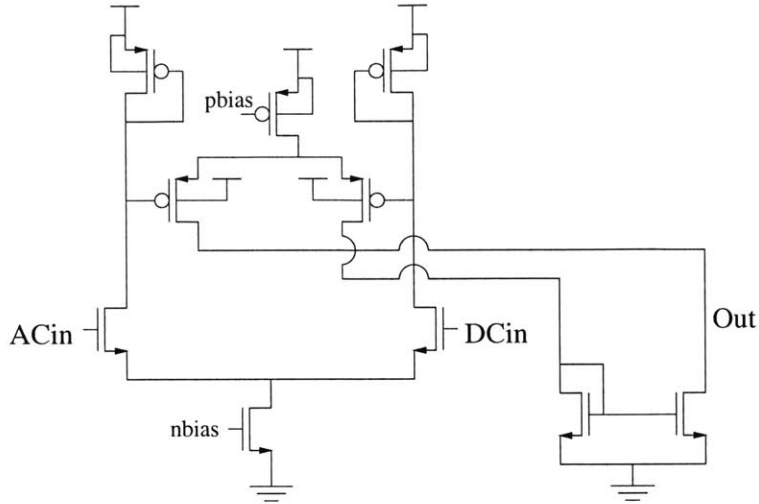


Figure 4-11: Clock Extractor Comparator

Figure 4-11 shows the implementation of the comparator from Figure 4-10. It consists of an NMOS differential pair with PMOS diode loads, followed by a PMOS differential pair with an NMOS mirror load.

Simulations showed that the clock extractor circuit consumes approximately $22\ \mu\text{W}$ of power.

Clock Generators

The previous paragraphs describe extraction of the base clock signal from the 125 KHz sine wave on the secondary coil. However, Section 4.3.2 makes clear the need for several clocks with different timing sequences.

Figure 4-12 shows the circuit which distributes the clock generated by the comparator. The device shown in the middle of the figure is a one-shot circuit, which produces a low pulse 140 ns wide when its input is a rising edge. This circuit is discussed in Section 4.3.5. The 140 ns clock is used to derive the clocks for the clocked comparator, as shown in Figure 4-13. Signals pCLK and nCLK are the first to change




Figure 4-12: Clock Distribution

(weakening or strengthening the main inverter latch), and then CLK1 and CLK2 are generated with a standard non-overlapping clock circuit.

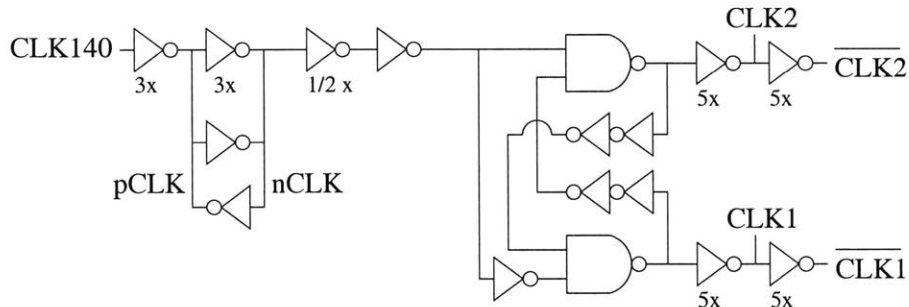


Figure 4-13: Clock Generator for the Clocked Comparator in Figure 4-9

4.3.4 Continuous Comparator

The Switch Network, Continuous Comparator, and Rectifier Switches blocks in Figure 4-2 are expanded in Figure 4-14. One of the five storage capacitor voltages is compared to the incoming AC sinusoid. The comparator turns on the appropriate MOSFET switch, via the gate drive, for a fixed amount of time, depending on some control parameters. The passgates were controlled by a state machine and turned on in a non-overlapping fashion. A continuous-time comparator is necessary in this case because the AC voltage is changing on a much more rapid time scale than the clock, so a clocked comparator could not be used for this application.

The continuous-time comparator shown in Figure 4-14 is implemented as shown in Figure 4-15 with a somewhat low power Bazes amplifier (described below) with a modified version of MeVay's differentiating front end similar to that used in the clock

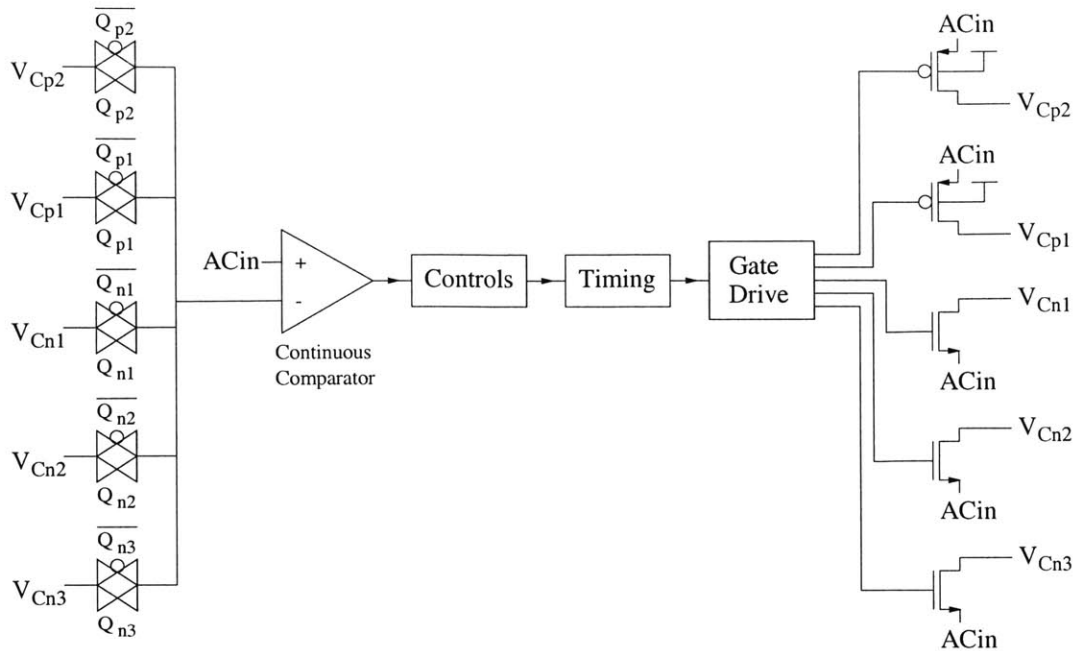


Figure 4-14: Rectifier Block

extractor (Figure 4-10), but with additional current cancellation mirrors to prevent current from being drawn from the storage capacitors. A rising AC input voltage injects current into MN1 via the input capacitance. This current is mirrored to MN2 and drawn from the resistor, lowering the voltage at the negative terminal of the amplifier. If that were the complete circuit, it would draw current from whatever was connected to node DCin. In this case, that current is cancelled to first order by mirroring the current in MN2 to MN3, MP3, and MP5, which injects the same current into node DCin. This method of first-order current cancellation was conceived by Luke Theogarajan.

Figure 4-16 shows the amplifier used in the continuous-time comparator. This amplifier is what Bazes [3] called the very wide common mode range differential amplifier (VCDA). This amplifier is fully complementary due to the combined PMOS and NMOS differential inputs, and functions over the entire supply range, thus it may be used with any of the storage capacitors. It is self-biased from the left column of transistors, divorcing its design from that of the bias network.

To understand the operation of this amplifier, let's look only at the NMOS input stage. The output currents of this differential pair are fed to a folded cascode load,

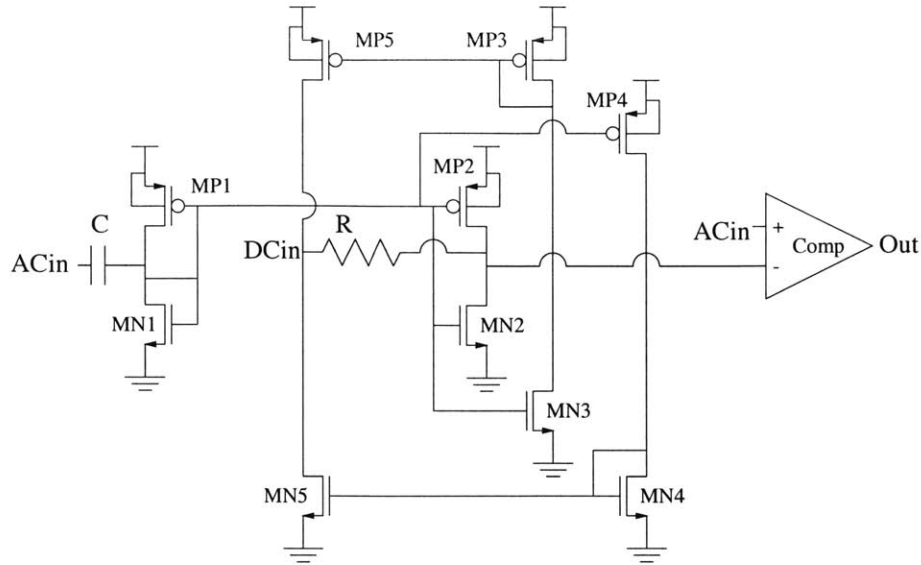


Figure 4-15: Continuous Comparator

with MP6 and MP7 serving as the cascode transistors for the input pair, MP4 and MP5 serving as bias current sources, and MN4 through MN7 serving as a cascode current mirror load. The PMOS differential pair works in a complementary fashion, with MN6 and MN7 serving as cascode transistors for the input pair, MN4 and MN5 giving bias current, and MP4 through MP7 serving as the cascode mirror load.

Initial design estimations showed that the amplifier would consume approximately $10 \mu A$ of current, or $35 \mu W$ of the $200 \mu W$ power budget, which is why Figure 4-14 shows only one such amplifier shared across the five capacitor channels. In the end, simulations reveal that this amplifier consumes approximately $54 \mu W$. However, this comparator is the most critical component of the rectifier system, and so it is reasonable for its power consumption to be rather high.

4.3.5 Rectifier Timing

When the comparator switches, the rectifier switch must be turned on for a certain time. As discussed in Section 4.2.3, the rectification time depends on the number of electrodes loading each capacitor. Figure 4-4 shows two different examples of switch timing waveforms, and this section describes the circuits which implement them. A

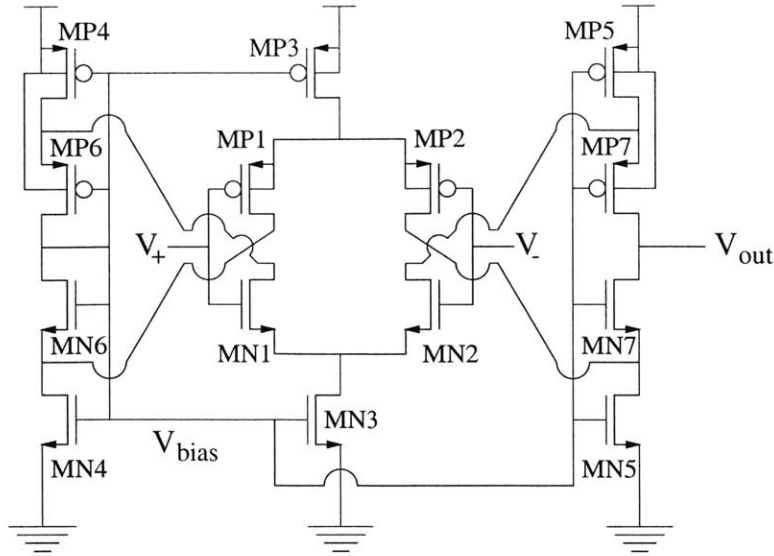


Figure 4-16: Bazes Very Wide Common Mode Range Differential Amplifier

simple pulse of fixed time is generated by a one-shot circuit, while the dual pulse with the short delay in the middle is created by a more complicated circuit, which includes a clock subdivider.

One-Shot Circuits

A typical one-shot circuit is shown in Figure 4-17, and consists of a nand gate and some inverting delay chain (such as the three inverters shown). The major drawback to the circuit as shown in the figure is that the inverter delay is typically on the order of 1 ns for our $1.5\text{ }\mu\text{m}$ CMOS process. The period of our AC waveform is $8\text{ }\mu\text{s}$, and we may wish to have delays of several hundred nanoseconds. It is inefficient in both area and power to simply string hundreds of minimum sized inverters together, so a more efficient solution must be found.




Figure 4-17: Typical One-Shot Circuit

An inverter with very long devices will have a much slower rise time, but that

has potentially devastating effects if used improperly. For example, if an inverter with PMOS and NMOS lengths that are 100 times the minimum size is placed right before a minimum sized inverter, the gate of the minimum sized inverter rises very slowly. This means the minimum sized inverter spends a significant amount of time with both NMOS and PMOS turned on, conducting substantial short-circuit current. When multiplied by many implementations of the one-shot circuit, this can amount to a large amount of wasted power. A better solution is to size the sequence of inverters by exponentially decreasing length, in exactly the same way that a string of inverters is scaled up by exponentially increasing width. This solution was tested in simulation, for $\lambda = 0.8 \mu m$, with a string of 3 exponentially sized inverters surrounded by minimum sized devices, as shown in Table 4.1 for a sizing factor of 5.

PMOS (W/L)	7/2	7/250	7/50	7/10	7/2
NMOS (W/L)	3/2	3/250	3/50	3/10	3/2

Table 4.1: Transistor Sizes (in λ) for a Symmetric Inverter Delay Chain

This gives an approximately symmetric delay while limiting short-circuit currents. Note, however, that the one-shot circuit need not have symmetric delay in its inverter string. In fact, the recovery time of the one-shot for a high to low input transition is limited by the inverter string delay. An asymmetric delay reduces the recovery time. To achieve asymmetric delay, alternating PMOS and NMOS devices are sized long, while all others are minimum sized, as shown in Table 4.2.

PMOS (W/L)	7/2	7/250	7/2	7/10	7/2
NMOS (W/L)	3/2	3/2	3/50	3/2	3/2

Table 4.2: Transistor Sizes (in λ) for an Asymmetric Inverter Delay Chain

For a rising input signal, the output of the first minimum sized inverter falls, turning on in sequence all of the long devices, giving a long delay. A falling input signal turns on all of the minimum sized devices, giving a very short delay. An additional benefit comes from the trip point of the inverters. Take for example the trip point of the third inverter. With a weak NMOS and strong PMOS, its trip point

is above midrail. For a rising input, the output of the second inverter is low, and rising slowly due to the very weak PMOS. The delay from the second inverter to the third is now increased due to the increased trip point of the third inverter.

Thus, with asymmetric sizing, the recovery time of the one-shot is greatly decreased, the circuit area is decreased since only half of the devices need to be long, and the area is further decreased for a given delay by the shifting that occurs in the inverter trip point with asymmetric sizing. The delays for the one-shot circuits used in this system were implemented with this sizing method. However, this delay circuit could have been implemented more robustly with a current-starved inverter chain. Since a bias circuit was needed for other circuitry, it could have been used for current-limiting transistors, further reducing the area needed, and decreasing the process dependence of the delay. Asymmetric sizing would still have been employed to help reduce the one-shot recovery time, but the recovery time would be limited by the allowed inverter currents. This idea is examined further in Section 7.2.

In circuit simulations, a 200 ns one-shot circuit, driven by an 8 μ s clock, consumed less than 0.6 μ W.

Clock Subdivider

The dual pulse waveform shown in the last phase of Figure 4-4 is created in two steps. First, the one shot circuit is removed, and the switch is controlled by the comparator output, so it stays on while the AC voltage is greater than the capacitor voltage. Second, a separate circuit subdivides the clock, giving a short pulse when the sinusoid is approximately 10° from reaching its peak. This short pulse interrupts the rectifier switch, releasing the AC voltage from the capacitor load, allowing it to turn on the V_{dd} Schottky rectifier. The AC secondary coil voltage rises very quickly because the effect of dropping the coil current down to zero causes a very steep rise in coil voltage. At the end of the short pulse, the rectifier switch is turned on, continuing the capacitor charging.

The clock is subdivided by establishing a sawtooth waveform which ramps from ground to V_{mid} during one phase of the clock. A feedback circuit ensures that the

ramp exactly reaches V_{mid} at the end of the clock phase. A bias network generates a voltage that is approximately 33% of V_{mid} , and a comparator determines when the ramp crosses this bias voltage. This should occur just before the middle of the clock phase, or just before the peak of the AC voltage. The comparator triggers a one-shot with a 400 ns pulse width.

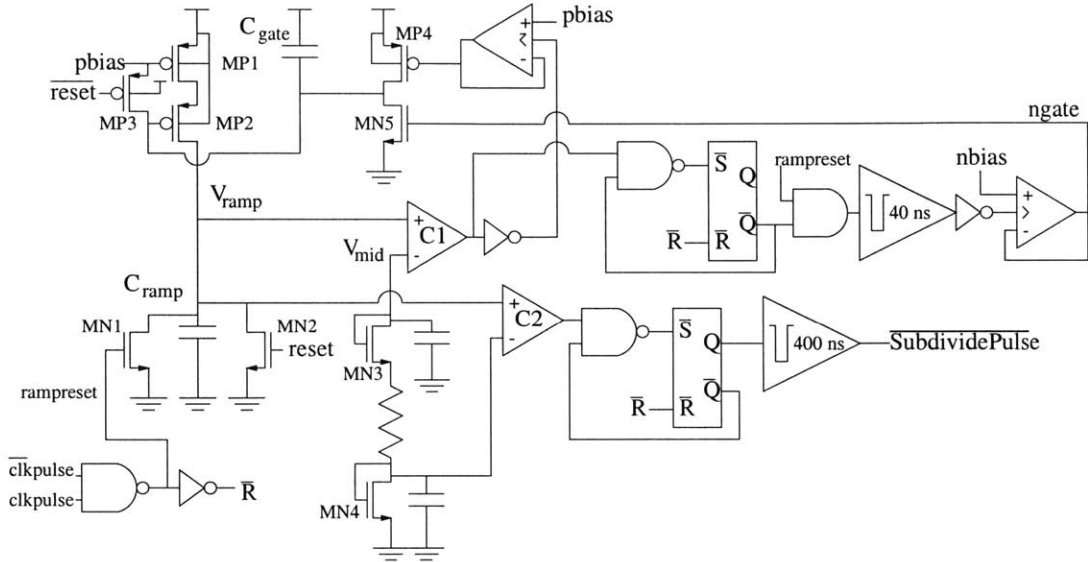


Figure 4-18: Clock Subdivider

This clock subdivider is implemented as shown in Figure 4-18, with comparator C1 implemented as shown in Figure 4-11, and comparator C2 implemented as the complementary circuit. Transistors MP1 and MP2 comprise a current source into capacitor C_{ramp} , generating voltage V_{ramp} . The gate of MP1 is at pbias, and MP2 is used to force MP1 into its linear region, thereby limiting its current. The gate of MP2 starts at pbias (having been set there by MP3 during reset), is held steady by capacitor C_{gate} , and is adjusted as needed by MP4 and MN5. At the beginning of a new clock phase, a one-shot pulse is generated, either clkpulse or $\overline{\text{clkpulse}}$. This pulse discharges C_{ramp} via MN1, and the sawtooth ramp begins. If V_{ramp} reaches V_{mid} before the end of the clock phase, comparator C1 turns on the pbias buffer, turning on MP4, raising the gate of MP2 to reduce the ramp current. The longer comparator C1 is on, the greater the reduction of the ramp current. If the end of the clock phase arrives and comparator C1 has not turned on, a 40 ns one-shot turns on the nbias

buffer, turning on MN5, lowering the gate of MP2 to raise the ramp current. The output is a 400 ns pulse derived from comparator C2, which monitors the sawtooth and the voltage on diode MN4, which is less than half V_{mid} . The resistor is 750 K Ω , implemented in an nwell. This output pulse begins before the peak of the AC voltage and extends through the peak, and is used to turn off the synchronous rectifier during this time.

4.3.6 Controls

A simple set of logic gates controls which capacitors receive charge during which phase of the AC waveform. The circuit receives as input the output of the five clocked comparators (signals CompP1, etc.). Obviously, if the AC voltage is positive (greater than V_{mid}), and C_{p1} needs charge while C_{p2} does not, then C_{p1} is charged during that phase. In cases where charge is needed on more than one capacitor, a register decides which is to be charged, and increments so that a different capacitor is charged the next time.

Figure 4-19 shows the implementation of the controller. State variables $\overline{Qp1}$, $\overline{Qp2}$, etc. are generated as sums of products with networks of NAND gates. First, in the upper left circuit, enable signal enP or enN is generated, depending on the clock phase, if the clock subdivider pulse and reset are off (high). The other two circuits on the left toggle the state of a register on the rising clock edge if charge is needed on more than one capacitor, changing state variables Qtp, Qtn0, and Qtn1. These state variables serve as tie-breakers when more than one capacitor needs charge during the same phase. The two positive state variables, Qp1 and Qp2, are defined in the center circuits, and the three negative state variables are defined in the circuits on the right.

$$Qp1 = CompP1 * enP * (\overline{CompP2} + CompP2 * Qtp) \quad (4.1)$$

$$Qp2 = CompP2 * enP * (\overline{CompP1} + CompP1 * \overline{Qtp}) \quad (4.2)$$

$$Qn1 = CompN1 * enN * (\overline{CompN2} * \overline{CompN3} +$$

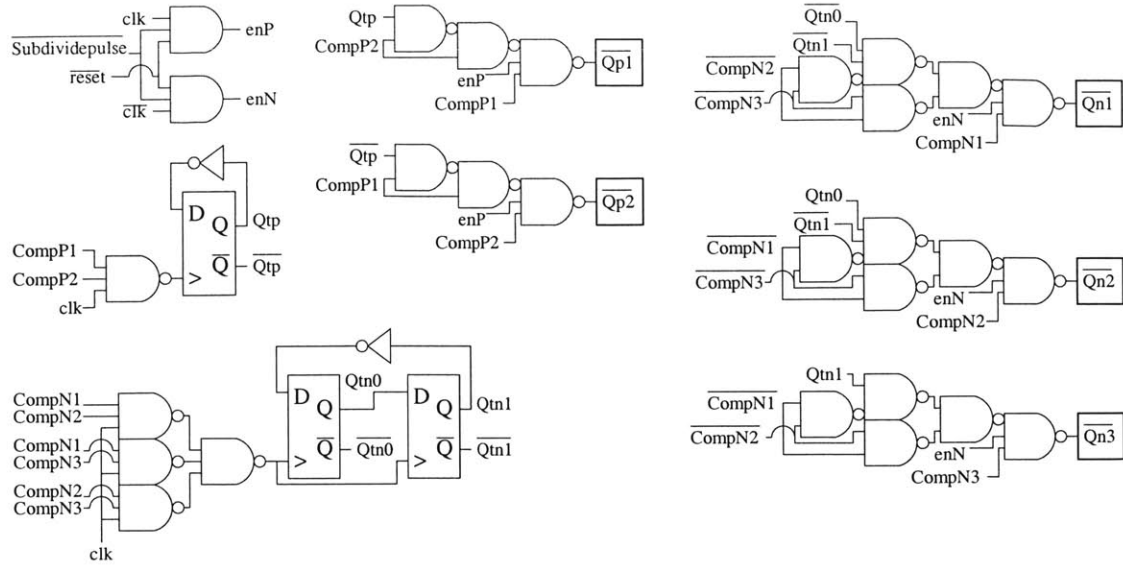


Figure 4-19: Capacitor Charging Controller

$$(\overline{\text{CompN2}} + \overline{\text{CompN3}}) * \overline{Qtn0} * \overline{Qtn1}) \quad (4.3)$$

$$Qn2 = \text{CompN2} * enN * (\overline{\text{CompN1}} * \overline{\text{CompN3}} + (\overline{\text{CompN1}} + \overline{\text{CompN3}}) * Qtn0 * \overline{Qtn1}) \quad (4.4)$$

$$Qn3 = \text{CompN3} * enN * (\overline{\text{CompN1}} * \overline{\text{CompN2}} + (\overline{\text{CompN1}} + \overline{\text{CompN2}}) * Qtn1) \quad (4.5)$$

The logic equations above (4.1 through 4.5) describe the state variables. For example, variable Qp1 turns on only if it is the correct clock phase (enP) and capacitor Cp1 needs charge (CompP1). In addition, either no other P capacitor needs charge ($\overline{\text{CompP2}}$), or another P capacitor needs charge, but the tie-breaking register says the charge should go to Cp1 (CompP2*Qtp).

4.3.7 Gate Drive and Switches

Figure 4-20 shows the control circuitry to drive the rectifier switches. The AND-OR-Invert (AOI) circuit at left determines when a switch turns on. When clk is high, one of the positive capacitors is being monitored by the comparator, and the comparator output goes high when the AC voltage exceeds the capacitor voltage. Therefore, the

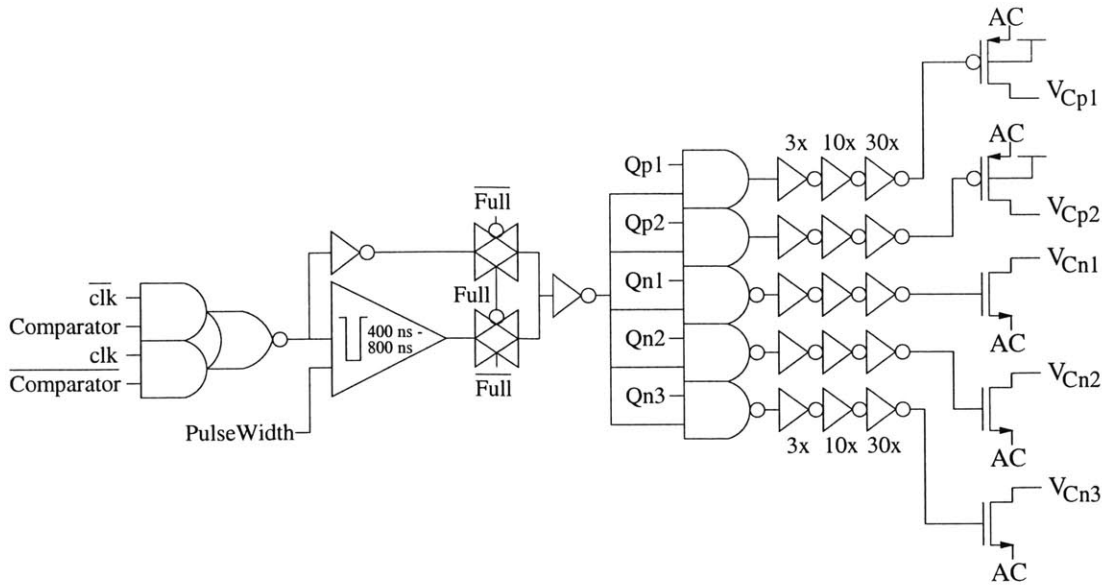


Figure 4-20: Gate Drive Circuitry

appropriate switch should turn on when clk and Comparator are both high. When a negative capacitor is being monitored, the comparator output goes low when the AC voltage dips below the capacitor voltage, so the switch should turn on when clk and Comparator are low. The signal Full determines whether the switch will be on for a fixed one-shot duration or for the entire time the comparator is high (in the case of a positive capacitor). Recall that even when Full is high and the switch is on while Comparator is high, the state variable (e.g., Q_{p1}) turns off for a brief time in response to the clock subdivider, in order to ensure turn-on of the V_{dd} rectifier. If Full is low, the switch turn-on duration is determined by the signal PulseWidth , which sets the one-shot to have a duration of either 400 ns or 800 ns . The series of NAND and AND gates ensures that the switch control signal from the comparator turns on only the appropriate switch, the switch whose state variable is on. The signal is then buffered through inverters of increasing size in order to drive the gates of the large rectifier transistors.

In addition to the rectifier switches, weak leakage paths were added to the storage capacitors. Since the gates of the rectifier transistors are set to V_{dd} (PMOS) or ground (NMOS) when the devices are turned off, subthreshold currents will flow when the AC voltage carries the PMOS sources above V_{dd} or the NMOS sources below ground.

This leakage current can drive the capacitor voltages beyond their target values if they are unloaded. The leakage transistors are operated in their linear resistive region, and are as follows: an NMOS from V_{mid} to V_{Cn1} , sized $3\lambda/30\lambda$, an NMOS from V_{mid} to V_{Cn2} , sized $3\lambda/80\lambda$, an NMOS from V_{mid} to V_{Cn3} , sized $3\lambda/160\lambda$, all with their gates tied to V_{dd} , a PMOS from V_{mid} to V_{Cp1} , sized $3\lambda/400\lambda$, with its gate at ground, and a PMOS from V_{mid} to V_{Cp2} , sized $3\lambda/1000\lambda$, with its gate at nbias. The NMOS rectifiers need stronger leakage paths because their gates are driven closer to threshold by the Schottky diode voltage drop, so they leak higher subthreshold currents.

4.3.8 Power On Reset

To ensure that the control circuitry initializes in the correct state when the implant is powered, a power on reset circuit was included. Figure 4-21 shows the circuit, which includes a diode string, MOS capacitor, and a comparator. The node labeled $V_{dd}/4$ is actually slightly greater than one-fourth of the power supply due to body effect in the PFETs nearer to ground. Furthermore, the voltage at that node rises more slowly than the power supply due to the MOS capacitor. When the voltage of the node labeled $V_{dd}/4$ exceeds nbias, the simple comparator flips state, bringing the chip out of reset mode. The comparator circuit is shown in Figure 4-22.

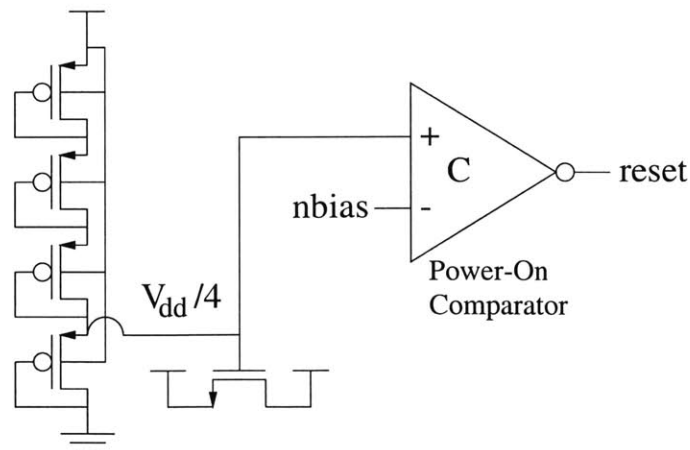


Figure 4-21: Power On Reset Circuitry

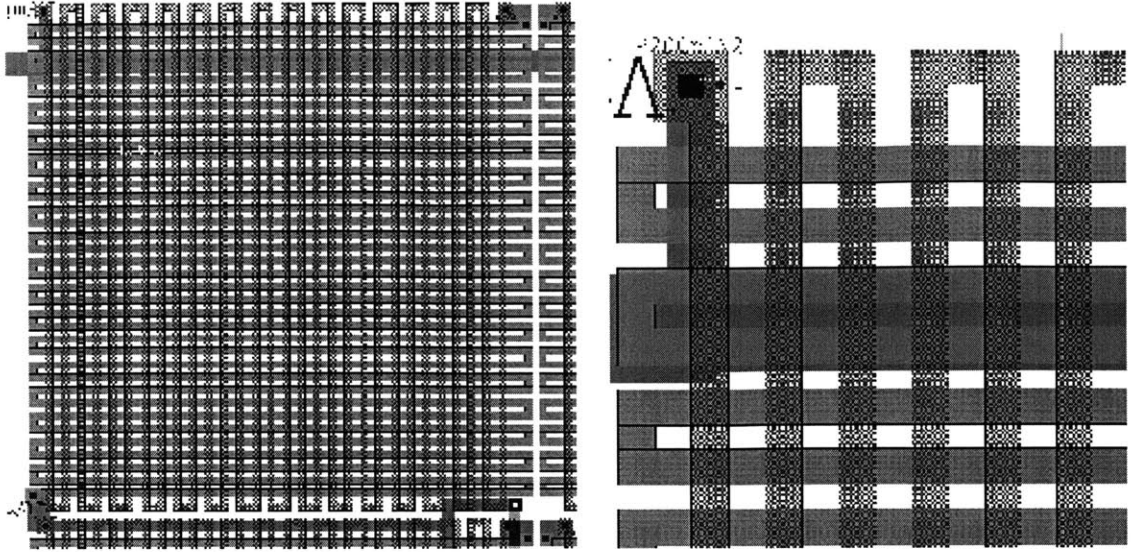


Figure 4-23: Poly-Poly2 Resistor Layout

At left, one resistor structure is shown, with the edges of three others. At right is a closer view of one corner of a resistor. The first polysilicon layer is red, and runs horizontally, the poly2 layer is grey and cross-hatched, and runs vertically. Ignore the upside-down V at the upper left of the right figure; it is simply part of a node label.

Clocked Comparator Layout




Figure 4-24: Clocked Comparator and Clock Generator Layout

Notice the five identical clocked comparators, as well as the clock generator on the right, abutting the rightmost comparator and sharing an nwell with it.

One significant layout error was discovered during chip testing. The clocked comparators described in Section 4.3.2 and shown in Figure 4-9 are reasonably sensitive circuits, but are controlled by several different clocks, generated by the circuit in Figure 4-13. The five clocked comparators were laid out in a row, but the clock generator circuit was placed right next to one of them, even sharing an nwell with part

of that comparator. In some of the chips tested, this particular comparator had an offset of over 500 mV (as compared to the other four comparators, with offsets well under 50 mV). However, at least one chip in each batch had less than 100 mV offset for the offending comparator, which was deemed acceptable. Figures 4-24 and 4-25 show the layout bug, and Figure 4-27 shows the correction, which was never sent for fabrication.




Figure 4-25: Closer View of Clocked Comparator and Clock Generator Layout
This view clearly shows the proximity problem and the shared nwell.




Figure 4-26: Die Photograph of Clocked Comparator and Clock Generator Layout




Figure 4-27: Corrected Clocked Comparator and Clock Generator Layout
The clock generator no longer shares an nwell with the rightmost comparator, and is isolated from it by substrate contacts.

4.4 Power System Testing

The power system chip was tested, and found to work satisfactorily. The clocked comparator layout problem, as discussed above, can be worked around by selecting chips with acceptably low offset on that reference channel. Figure 4-28 repeats the block diagram from Figure 4-2, but shows figure numbers for test data above the relevant nodes.

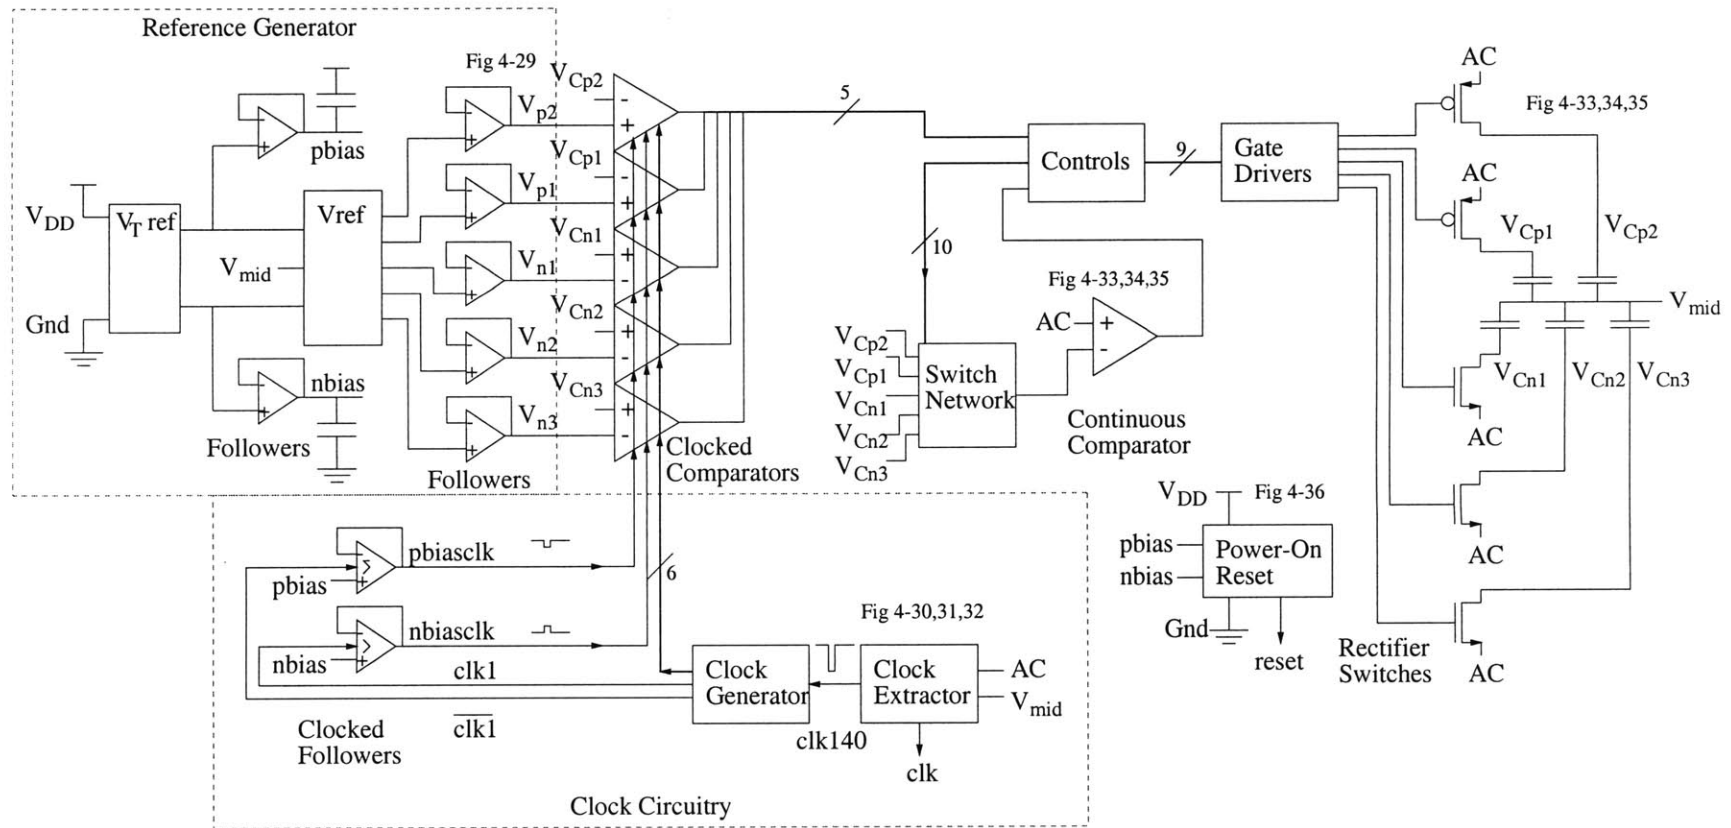


Figure 4-28: Synchronous Rectifier Testing Block Diagram

4.4.1 Voltage Reference Testing

The voltage reference circuit was tested by cycling through the eight reference levels. A simple 74HC393 counter was given a burst of eight clock pulses from a function generator, and three of its count bits were used as the reference level bits. Figure 4-29 shows the resulting output of the voltage reference circuit.




Figure 4-29: Reference Voltage Circuit Test Results

This plot shows the eight voltage levels of the five reference channels. Recall from Section 4.3.1 that instead of ranging from 0 to 7 steps of voltage, this circuit ranges from 2 to 9 steps. The center line is V_{mid} , lines above V_{mid} are V_{p1} and V_{p2} , respectively, and lines below are V_{n1} , V_{n2} , and V_{n3} .

4.4.2 Clock Extractor Testing

The output of the clock extractor comparator is shown in Figure 4-30, with close views of the rising and falling edges in Figures 4-31 and 4-32, respectively. Note that the predictive circuit employed in the clock comparator reduces the offset/delay to a very small value, roughly 30 ns.




Figure 4-30: Clock Extractor Test Results




Figure 4-31: Clock Extractor Test Results - Rising Edge

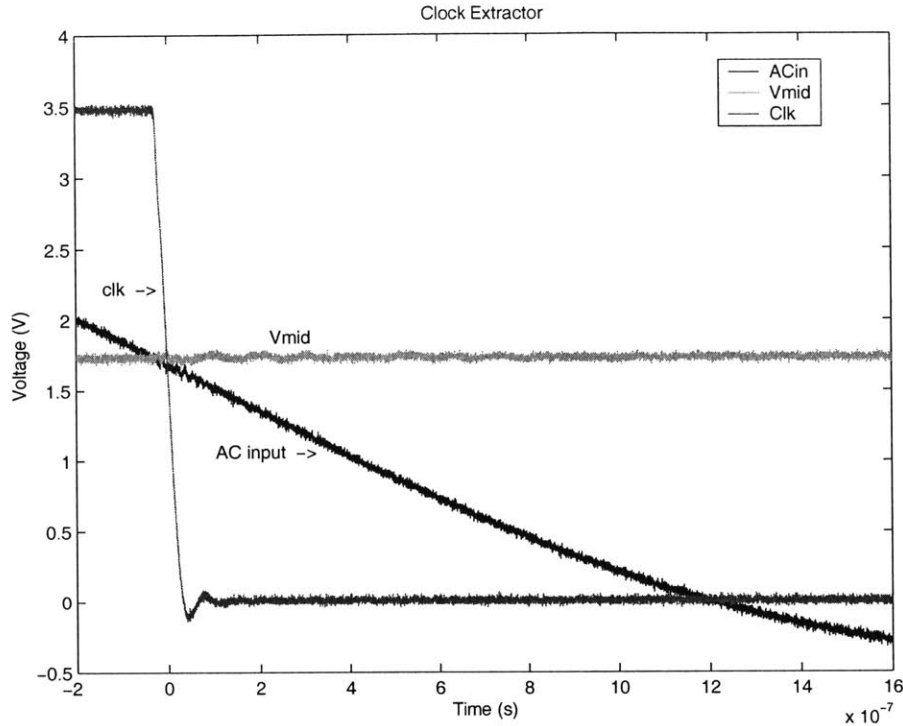


Figure 4-32: Clock Extractor Test Results - Falling Edge

4.4.3 Synchronous Rectifier Comparator Testing

Due to the limited number of package pins, the output of the synchronous rectifier comparator cannot be examined directly. However, information on its effect can be extracted from the distortion of the AC input waveform. Figure 4-33 shows this effect. The effective delay of the synchronous rectifier can be reported as no greater than the time between the crossover of the AC voltage and the capacitor voltage and the noticeable distortion of the AC voltage. This delay, taken from Figures 4-33, 4-34, and 4-35, ranges from 40 ns to 44 ns , with a mean of 41.3 ns .

These figures show four traces, the AC voltage, V_{mid} , V_{p2} , and V_{Cp2} . Note that V_{p2} and V_{Cp2} are practically equal. This shows the very small offset voltage of the clocked comparators. Note also the large oscillations on V_{mid} and V_{Cp2} . This is due to the inductance of the very long wires necessary for this test setup.




Figure 4-33: Synchronous Rectifier Test Results - 400 ns

The synchronous rectifier, with rectification time set to its smallest value, nominally 400 ns.

4.4.4 Synchronous Rectifier Timing Testing

The timing controls for the synchronous rectifier were extracted in the same way as the comparator delays, from the distortion effects on the AC voltage. Figure 4-33 shows rectification with a nominal duration of 400 ns. The actual measured duration was 668 ns. Figure 4-34 shows rectification with a nominal duration of 800 ns, and actual measured duration of 964 ns. The error in the shorter duration is likely due to variations in the process, resulting in a one-shot circuit that is 67% slower than designed. The second delay, then, should be 1.34 μ s. But it is so much longer than expected that it ran into the limitations set by the clock subdivider circuit. That is to say, the one-shot stayed on for the full 1.34 μ s, but the clock subdivider intervened, turning off the state variable, and therefore the rectifier. Were this circuit to be redesigned, current-starved inverters would be used in one-shot delays for better timing accuracy, and the clock subdivider circuit would turn off the state variables at a later time, closer to the peak of the AC voltage.




Figure 4-34: Synchronous Rectifier Test Results - 800 ns

The synchronous rectifier, with rectification time set to its larger value, nominally 800 ns.

Figure 4-35 shows rectification for the full duration that the AC voltage exceeds the capacitor voltage, except for a short time in the middle allowed for Schottky diode V_{dd} rectification. This middle turn-off time was 664 ns in duration, as expected from the above experiments.

4.4.5 Power On Reset Testing

The results of testing the power on reset circuitry are shown in Figure 4-36. Note that, as the power supply approaches its final voltage, reference voltage nbias settles quickly to its correct value, while the output of the diode divider, $V_{dd}/4$, is delayed by the capacitance on that node. The result is that the chip reset signal goes high for a short period, resetting the chip to a known starting state.

This test suffered from a startup problem with the Agilent E3646A power supply. It would ramp slowly (over a period of 20 ms) to 1 volt, then jump quickly (in about 2 ms) to 3.5 volts. However, as the figure shows, the power on reset circuit still




Figure 4-35: Synchronous Rectifier Test Results - Full Time Rectification
 The synchronous rectifier, set to rectify while the AC voltage is larger than the capacitor voltage, with a short pause to turn on the Schottky diode V_{dd} rectifier.




Figure 4-36: Power On Reset Test Results

worked.

4.4.6 Rectifier Chip Power

The total current used by the rectifier chip was measured, and the resulting chip power from V_{dd} , not including power drawn from the AC voltage, was calculated. Optimistic estimates from circuit simulations, not including the substantial digital logic, amount to $150 \mu W$. The measured current was $52 \mu A$, giving a chip power of $182 \mu W$, a very reasonable value.

Chapter 5

Electrode Drive

This chapter describes the section of the chip which handles the electrodes, switching the appropriate electrodes among the series of storage capacitors. Recall from Chapters 2 and 4 that each electrode is biased to sit at V_{p1} , and moves through four voltage steps for negative current stimulation and four voltage steps for positive current stimulation. Electrodes are switched to V_{mid} , V_{n1} , V_{n2} , and V_{n3} during the negative first phase, and to V_{n1} , V_{mid} , V_{p1} , and V_{p2} during the positive second phase. Each phase lasts 5 ms , so each step lasts 1.25 ms , and a full stimulation lasts 10 ms . The stimulations begin either immediately at the end of the previous stimulation, for a frequency of 100 Hz , or after a 10 ms delay following the end of the previous stimulation, for a frequency of 50 Hz .

Since this chapter is far more digital than Chapter 4, the circuits, at least those which work properly, will be described in a somewhat more cursory manner.

5.1 Pre-Programmed Test Patterns

Since data transmission is not included in this thesis project, test stimulation patterns were pre-programmed in the chip. Using two control bits, four patterns were defined, as described in Table 5.1.

Figure 5-1 shows the electrodes driven for the “X” test pattern, while Figure 5-2 shows the electrodes driven for the “MIT” pattern. The latter figure includes the

00	No Electrodes On
01	50 Hz, “X” Pattern
10	100 Hz, All 15 Electrodes
11	50 Hz, Repeating “MIT” Pattern

Table 5.1: Pre-Programmed Test Patterns

state bits for a simple grey code counter described in Section 5.2.2. The row and column numbers in the figures are somewhat vestigial, as this 3x5 electrode array is backward compatible with hardware used to drive a 10x10 array in the past. These 15 electrodes connect to what would be approximately the middle of the 10x10 array, and hence have columns numbered 4, 5, and 6.




Figure 5-1: “X” Test Pattern

These patterns are described in this first section because they determine in part the control circuitry. Knowing which electrodes need to be turned on in which situations makes the control circuitry easier to understand.

5.1.1 Electrode Pattern Groups

Notice in the above figures of groupings that certain electrodes are always on or off together. Electrodes 1 and 11, for example, are both on in the “X,” “M,” and “I” patterns, and off in the “T” pattern. Likewise, electrodes 5 and 15 are always on and off together, as are electrodes 2, 3, 4, 12, 13, and 14, and electrodes 6, 7, and 10. There are actually only six unique electrode commands, for electrodes 1, 2, 5, 6, 8, and 9.

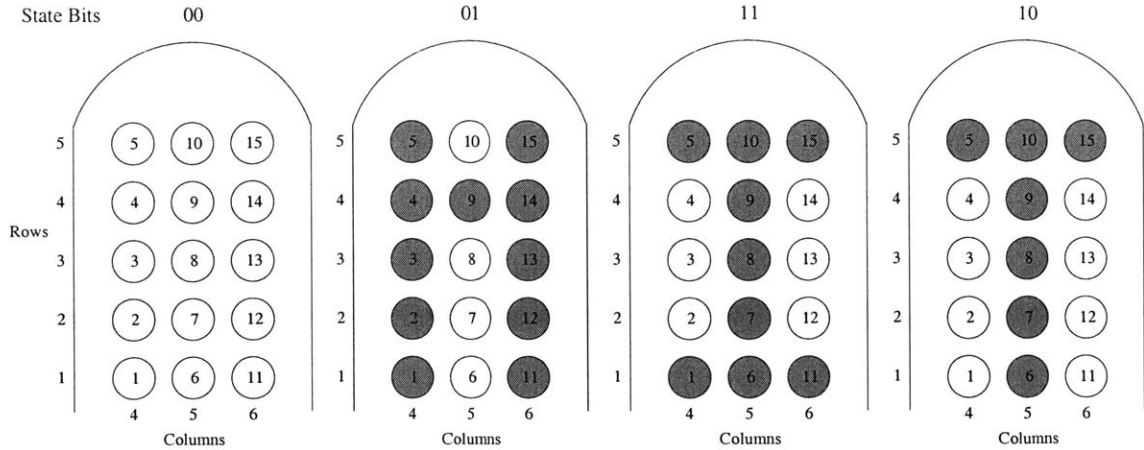


Figure 5-2: “MIT” Test Pattern

All other electrodes follow one of those control signals. In the actual implementation, the groupings described below increase the number of unique electrode groups to nine instead of six.

5.1.2 Electrode Drive Groups

Recall from Section 4.2.3 that, when an electrode is driven, it is switched from its bias voltage of V_{p1} to V_{mid} , V_{n1} , V_{n2} , and V_{n3} , in sequence, then in the other direction to V_{n1} , V_{mid} , V_{p1} , and V_{p2} , after which it is open-circuited to return to the V_{p1} bias via a very weak current source, described in Section 5.2.3.

If every electrode followed that sequence in phase, only one of the storage capacitors would be loaded at any given time, and it would be heavily loaded. A more efficient method is to group the electrodes and phase-shift the groups with respect to one another. In other words, when one group of electrodes is starting its negative phase, another group might be starting its positive phase, distributing the load across multiple capacitors. Four groups were chosen, but at 100 Hz , the highest possible stimulation frequency, the four groups merge into two. The four groups, labeled A through D, and their capacitor switching sequences, are outlined in Tables 5.2 and 5.3.

Using these electrode drive groupings, as well as the control groupings described in the previous section, we can define nine unique signals to drive the 15 electrodes.

Group	Electrodes
A	1, 11, 3, 13
B	2, 12, 6, 9
C	5, 15, 7, 10
D	4, 14, 8

Table 5.2: Electrode Drive Group List

State	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Group	Capacitor connection															
A	m	n1	n2	n3	n1	m	p1	p2	-	-	-	-	-	-	-	-
B	-	-	-	-	m	n1	n2	n3	n1	m	p1	p2	-	-	-	-
C	-	-	-	-	-	-	-	-	m	n1	n2	n3	n1	m	p1	p2
D	n1	m	p1	p2	-	-	-	-	-	-	-	-	m	n1	n2	n3

Table 5.3: Electrode Drive Grouping Capacitor Connections

The top line shows the state count, where each count is one-fourth of a stimulation phase, or 1.25 ms . When the stimulation frequency is 100 Hz , the first eight states merge with the last eight, i.e., states 0 and 8 are the same. In this case, groups A and C are the same, and B and D are the same. For brevity, m is used in this table to describe the V_{mid} node.

These signals control electrodes 1/11, 2/12, 3/13, 4/14, 5/15, 6, 7/10, 8, and 9.

Note that, in Table 5.3, no electrode drive groups are connected to the same sign (p or n) capacitor at the same time. This means that at no time are multiple capacitors of the same sign loaded by electrodes, making the synchronous rectification easier.

5.2 Digital Control System

Having discussed the capacitor connection sequence for the electrodes, as well as which electrodes turn on at which time, we now proceed to the circuitry which implements these patterns and connections.

Three main pieces comprise the digital control system, shown in Figure 5-3. The main state machine, in the center, is described in Section 5.2.1 and generates timing signals for both the digital controls and the analog synchronous rectifier. It divides the $8\text{ }\mu\text{s}$, 125 KHz clock signal down to a 1.25 ms clock used for electrode switching.

This clock steps a counter, which serves as the state count shown in Table 5.3. In addition, the main state machine determines the capacitor connections for each of the four electrode drive groups (A, B, C, D) at any given time. That is to say, if an electrode were in group C, for example, and if that electrode were supposed to be delivering stimulation current, signals from the main state machine would determine the capacitor (or V_{mid}) to which that electrode should be connected. The electrode state machine (Section 5.2.2) and MIT loop controller (Section 5.2.2) determine, for a given stimulation pattern, which electrodes should be receiving stimulation current. Its six output signals correspond to the six electrode pattern groups discussed in Section 5.1.1. The switch controller, described in Section 5.2.3, determines from the main state machine and electrode state machine which of the 90 (15 electrodes, 5 capacitors + V_{mid}) large switches should be turned on, and switches it on via inverter chain switch drivers. In addition, the switch controller includes the weak current sources which return open-circuited electrodes to their bias voltage.

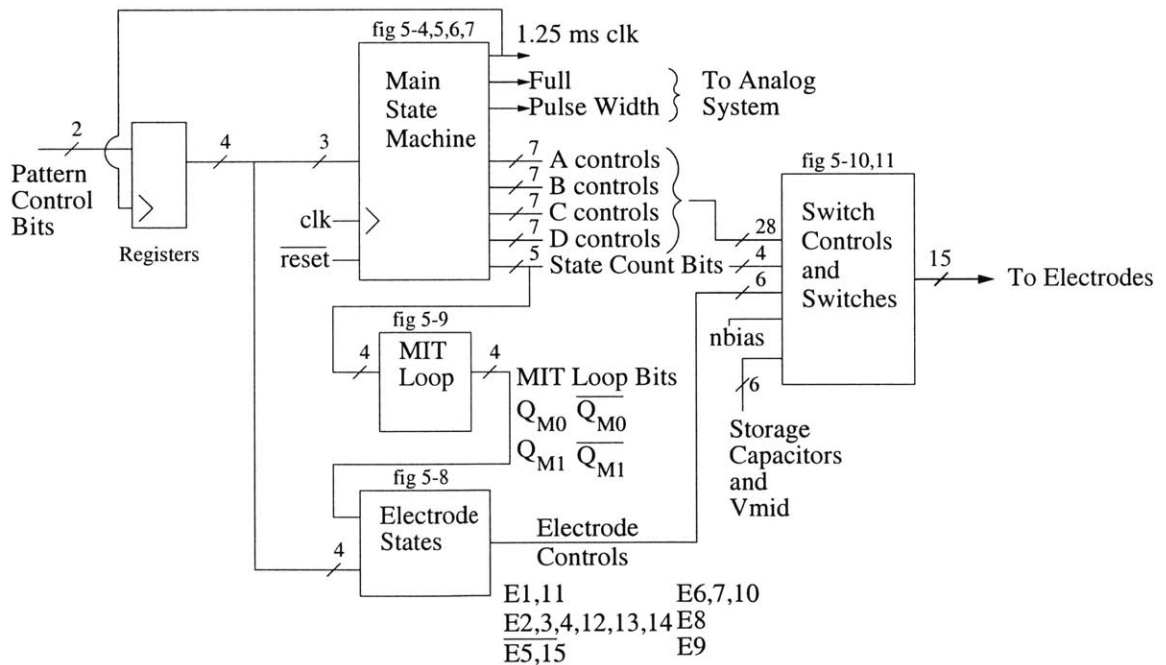


Figure 5-3: Digital Control System Block Diagram

5.2.1 Main State Machine

The main state machine is shown in Figure 5-4. The subcircuits are explained in greater detail in the following sections. The first important piece of this circuit is the clock divider, which divides the $8\ \mu s$ chip clock down to a $1.248\ ms$ clock. This clock, considered for simplicity to have a period of $1.25\ ms$, is high most of the time, with a low pulse of $32\ \mu s$. It is the equivalent of a counter's $\overline{\text{RippleCarryOut}}$ signal.

The $1.25\ ms$ clock then enters a 4-bit counter, which keeps track of the stimulation phase. The two lowest bits, Q8a and Q9a, determine which quarter of the $5\ ms$ stimulation phase the system is in. The third bit, Q10a, determines whether the system is in the positive or negative stimulation phase, and the fourth bit, Q11a, allows for a $10\ ms$ delay between stimuli if $50\ Hz$ stimulation is used. These four bits account for the 16 states in Table 5.3.

In that table, note that the electrode groups shift by four states at a time. In other words, in moving from group A to group B, we need to subtract four states. To implement this, it is easiest to subtract one from the third and fourth bits. A 2-bit subtracter takes very little circuitry, and is implemented with the small AND-OR-Invert (AOI) circuit, generating Q11b, and by inverting the third bit (i.e., $Q10b = \overline{Q10a}$). From here, $Q11c = \overline{Q11a}$, $Q11d = \overline{Q11b}$, $Q10c = Q10a$, and $Q10d = Q10b$. Thus the four state bits for each of the four electrode groups are easily extracted from the state bits for group A.

At the right of Figure 5-4, switch control subcircuits determine, based on the four electrode group state bits, which capacitor that group should connect to, based on Table 5.3. For example, if group A is in state 0 ($\overline{Q11a}\ \overline{Q10a}\ \overline{Q9a}\ \overline{Q8a}$), its electrodes (those which are turned on) should connect to V_{mid} , so signal Amid is high. When group B is in state 0 ($\overline{Q11b}\ \overline{Q10a}\ \overline{Q9a}\ \overline{Q8a}$, state 4 in the table), signal Bmid is high. The switch control also corrects state variable Q11 depending on whether stimulation is occurring at 50 or 100 Hz. At 50 Hz, Q11a toggles high and low as expected, but at 100 Hz, Q11a should always be low, so that the electrodes always connect to somewhere. Note in Table 5.1 that the lower order bit alone distinguishes between

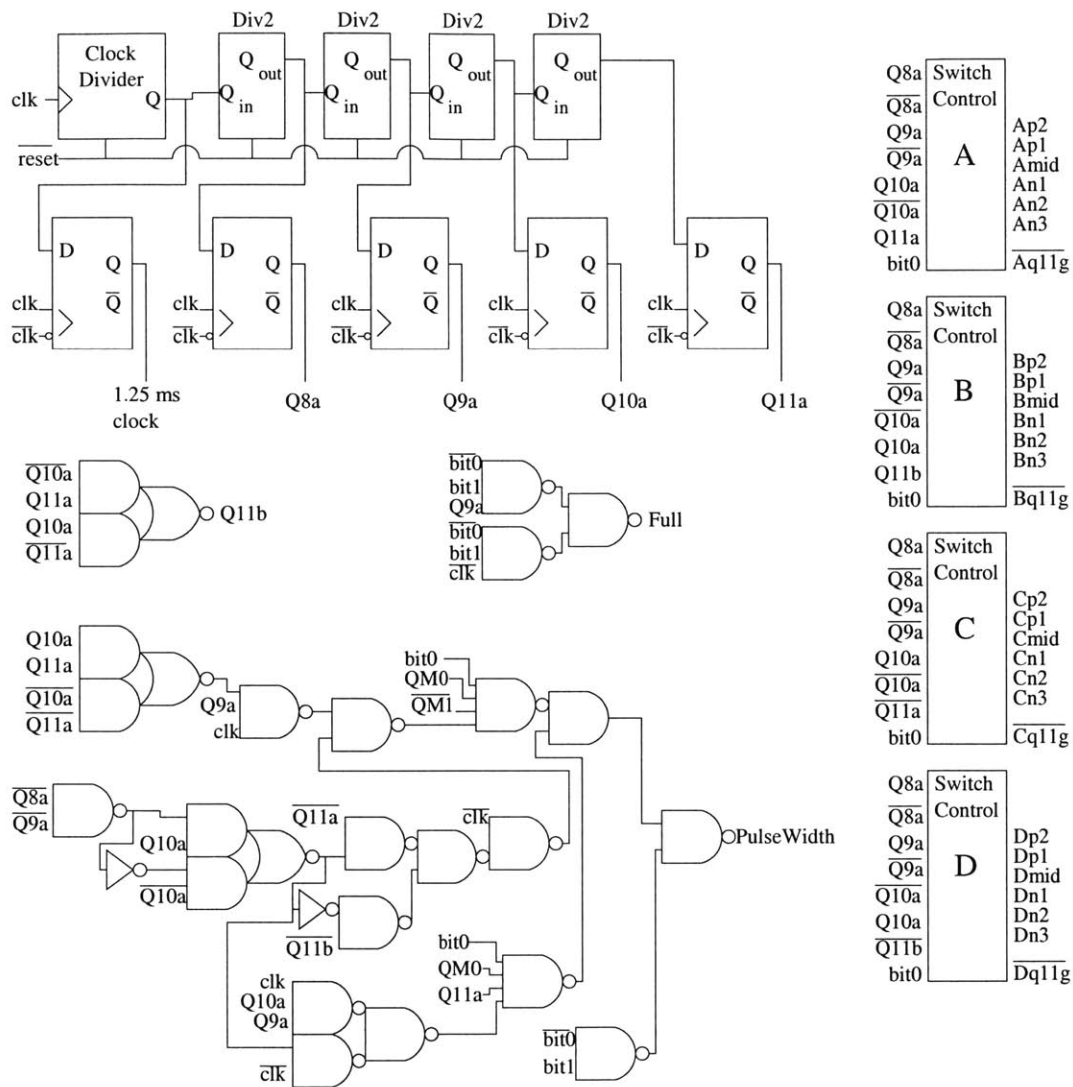


Figure 5-4: Main State Machine

50 and 100 Hz, so it is used by the switch controller to determine the frequency.

The final section of the main state machine contains the timing signals for the synchronous rectifier. In its final version, this system needs to count how many electrodes are loading each capacitor, and adjust the rectification time accordingly. Since we know in advance the stimulation patterns, we can simply calculate these timing signals with simple logic.

The rectifier by default charges for the short duration. If signal PulseWidth is high, it charges for the longer duration, and if signal Full is high, it charges while the AC voltage is higher than the capacitor voltage, except for a pause to turn on the Schottky rectifier. PulseWidth is set high if three or more electrodes are connected to a capacitor, and that capacitor is charging (if a positive capacitor, such as C_{p1} is loaded by three electrodes and the 125 KHz clock is high, then PulseWidth would be high). Full is set high if seven or more electrodes load a capacitor which needs to be charged.

Starting with the simpler case, signal Full is on only during pattern 3, all 15 electrodes stimulated at 100 Hz. In this pattern, there is always a negative capacitor which is heavily loaded, so Full is high for bit1 $\overline{\text{bit0}} \overline{\text{clk}}$. In addition, a positive capacitor is heavily loaded in states 3, 4, 6, and 7 (recall that states 8-15 are combined with states 0-7 at 100 Hz). So Full is also high for bit1 $\overline{\text{bit0}} \text{clk} Q_{9a}$. By Boolean logic, signal clk drops out when the two expressions are combined, so we have:

$$Full = bit1 \overline{\text{bit0}} (\overline{\text{clk}} + Q_{9a}) \quad (5.1)$$

The case for PulseWidth is a little more complicated. It should be on during pattern 4, when the M is displayed, for groups A and B, and when the I and T are displayed for group C. In addition, it was turned on during pattern 3 to ensure faster charging of p capacitors in between loads. An error resulted in PulseWidth being high in some cases during pattern 2, but this does not affect functionality, it simply makes stimulation using pattern 2 slightly less efficient. The full equation for PulseWidth

is:

$$\begin{aligned}
 PulseWidth = & bit1 \overline{bit0} \\
 & + bit0 \overline{Q_{M1}} Q_{M0} [Q_{9a} clk (Q_{10a} \oplus Q_{11a}) \\
 & \quad + \overline{clk} (\overline{Q_{11a}} (Q_{10a} \overline{Q_{9a}} \overline{Q_{8a}} + \overline{Q_{10a}} \overline{\overline{Q_{9a}} \overline{Q_{8a}}}) \\
 & \quad + \overline{Q_{11b}} (Q_{10b} \overline{Q_{9a}} \overline{Q_{8a}} + \overline{Q_{10b}} \overline{\overline{Q_{9a}} \overline{Q_{8a}}})] \\
 & + bit0 Q_{M1} Q_{11a} (clk Q_{10a} Q_{9a} \\
 & \quad + \overline{clk} (Q_{10a} \overline{Q_{9a}} \overline{Q_{8a}} + \overline{Q_{10a}} \overline{\overline{Q_{9a}} \overline{Q_{8a}}})) \quad (5.2)
 \end{aligned}$$

Clock Divider

The clock divider circuit is shown in Figure 5-5. It uses a series of the same divide by 2 flip-flops used as counters in the main state machine. To generate a 1.248 ms clock from an 8 μs clock, this circuit must divide the incoming clock by 156. It does this by first dividing it down by 4, then sending the resulting 32 μs clock into a resettable counter, which is reset every 39 cycles. When the counter output is 37 (on the 38th cycle), the reset signal is loaded into a D flip-flop register, and executes on the next cycle.

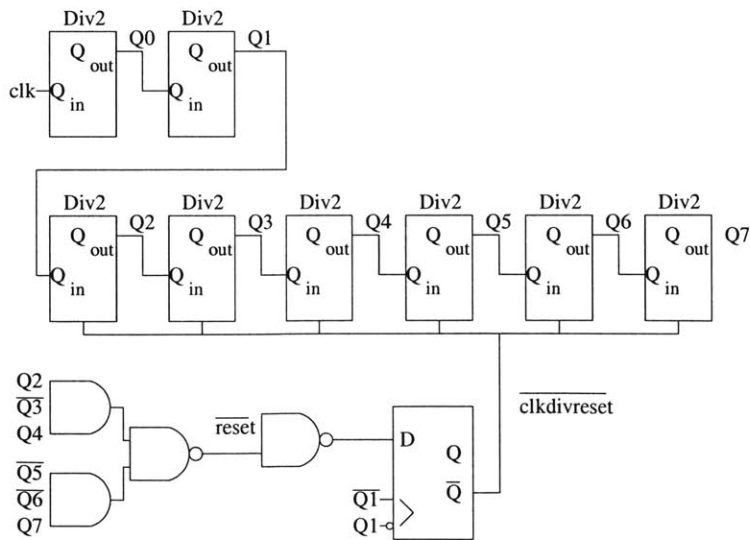


Figure 5-5: Clock Divider

The divide by 2 flip flop, introduced to me by Luke Theogarajan, is shown in

Figure 5-6. When the incoming clock, Q_{in} , goes from low to high, the amp on the left sets Q_n equal to $\overline{Q_{out}}$. When Q_{in} goes from high to low, the amp on the right sets Q_{out} equal to Q_n , or $\overline{Q_{out}}$ from the previous cycle. So the output toggles on a high to low transition of the input, the correct counter functionality.

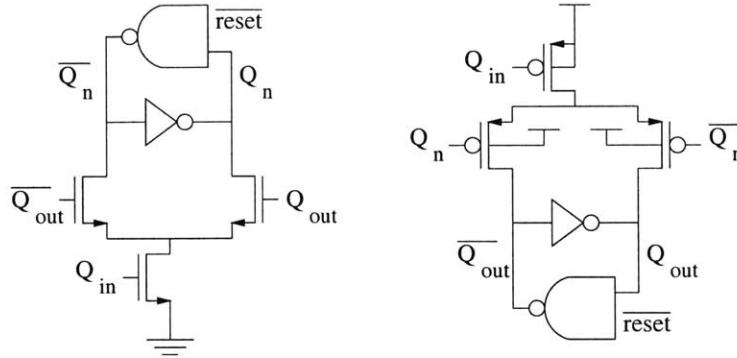


Figure 5-6: Divide By 2 Flip-Flop

Switch Control Logic

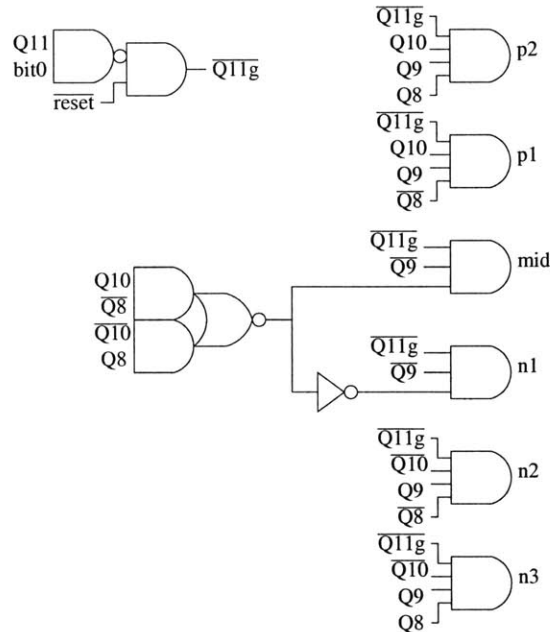


Figure 5-7: Switch Control Logic

The switches are controlled by the circuits in Figure 5-7. The leftmost circuit uses pattern bit 0 to determine the pattern frequency, and fixes Q_{11} low for 100

Hz patterns. The other six circuits are straightforward, turning on connections to certain voltages at certain states. For example, n2 turns on in state 2, which is $\overline{Q11g} \overline{Q10} \overline{Q9} \overline{Q8}$. The mid and n1 connections each happen twice, mid at states 0 and 5, n1 at states 1 and 4. So mid turns on at $\overline{Q11g} \overline{Q9} (\overline{Q10} \overline{Q8} + Q10 Q8)$ and n1 turns on at $\overline{Q11g} \overline{Q9} (\overline{Q10} Q8 + Q10 \overline{Q8})$, as calculated by the AOI circuit.

5.2.2 Electrode State Controls

The electrode state controller, shown in Figure 5-8, is quite simple, enabling certain electrodes during certain patterns. Each of the six circuits controls one of the six electrode pattern groups outlined above. The second signal, though abbreviated to E234, actually controls electrodes 2, 3, 4, 12, 13, and 14. A few other abbreviations should be explained. In the lower left circuit, $\overline{\text{all}}$ signifies pattern 2, all 15 electrodes on, which repeats in several other circuits. In the upper left circuit, $\overline{\text{Xall}}$ signifies pattern 1 or 2, the X pattern or all electrodes on. In the lower middle circuit, $\overline{\text{IT}}$ signifies pattern 3, states 3 or 4, when the pattern is I or T.

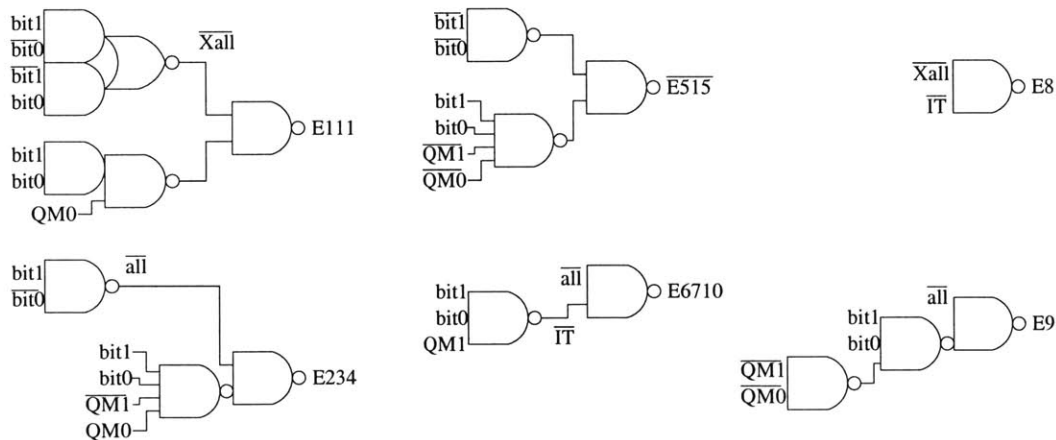


Figure 5-8: Electrode Control Logic

MIT Loop Control

The circuit which generates the MIT loop pattern, shown in Figure 5-9 is a grey code counter which is clocked at state 15 for group A. When clocked, the first D flip-flop

feeds its output to the second, which feeds the inverse of its output to the first. This yields a sequence of bits QM1 QM0 of 00, 01, 11, 10, 00, ...




Figure 5-9: MIT Pattern Loop Logic

5.2.3 Switch Driver Controls

The final segment of the digital control portion of the chip drives each of the 90 switches used to connect the 15 electrodes to one of 6 voltages, as shown in Figure 5-10. The signals to enable the 9 electrode pattern groups are registered with a local electrode group (A, B, C, D) clock. These signals are then ANDed with the switch control signals for that group (Ap2, etc.), to generate gate signals for each of the 90 switches. Since the switches are quite large (PMOS: $\frac{250\lambda}{2\lambda}$, NMOS: $\frac{100\lambda}{2\lambda}$), they require drivers. The PMOS driver consists of three inverters, sized $((\frac{W}{L})_P - (\frac{W}{L})_N): \frac{20\lambda}{2\lambda} - \frac{8\lambda}{2\lambda}; \frac{50\lambda}{2\lambda} - \frac{20\lambda}{2\lambda}; \frac{125\lambda}{2\lambda} - \frac{50\lambda}{2\lambda}$. The NMOS driver consists of two inverters, sized: $\frac{25\lambda}{2\lambda} - \frac{10\lambda}{2\lambda}; \frac{75\lambda}{2\lambda} - \frac{30\lambda}{2\lambda}$. Finally, on the right, this figure shows the electrode bias current source, leaking current to the electrode from capacitor C_{p1} .

Electrode Bias Current Sources

The current source which ensures that each electrode floats back to V_{Cp1} after stimulation is shown in Figure 5-11. It is a simple set of mirrors based on nbias, and will pull the electrode up to V_{Cp1} . The current source has a nominal current of 180 – 200 nA if the electrode voltage is a few hundred mV below V_{Cp1} , and drops off as the electrode voltage approaches V_{Cp1} . Should the electrode voltage exceed V_{Cp1} , the PMOS transistor will draw much more current from it, as much as 1 μA in the unlikely event that the electrode reaches a few hundred mV above V_{Cp1} .

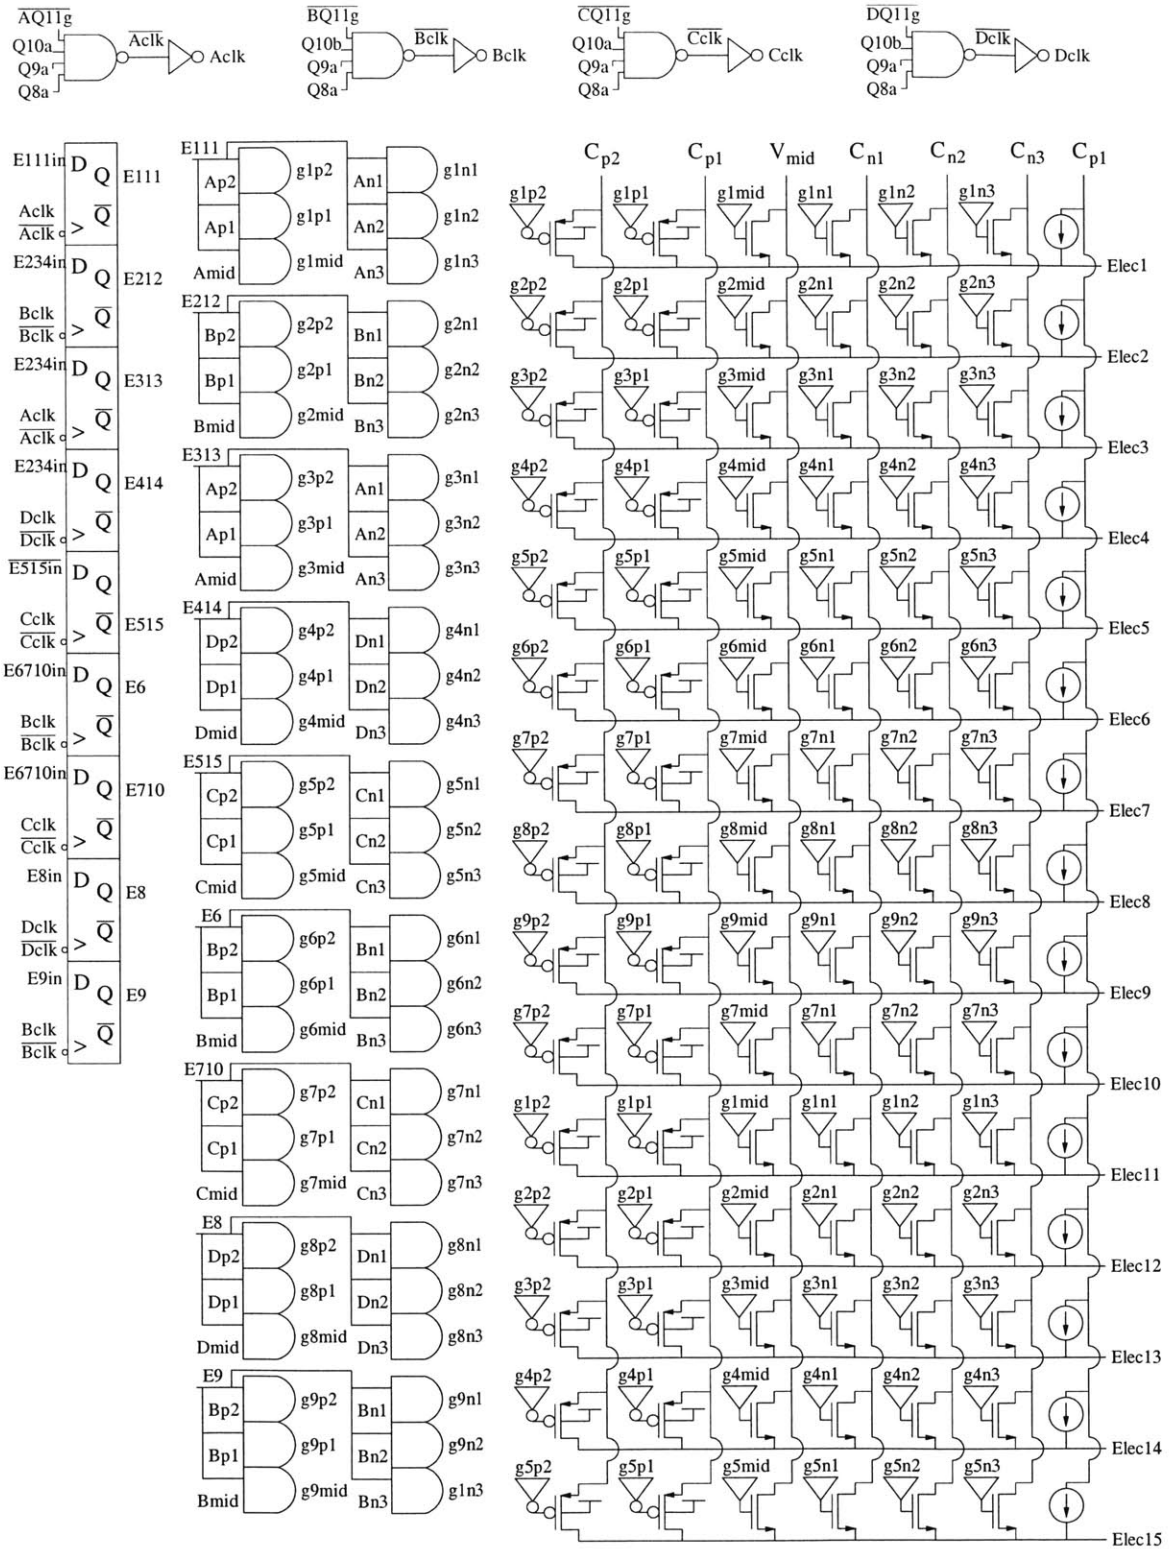


Figure 5-10: Switch Driver Logic and Switches

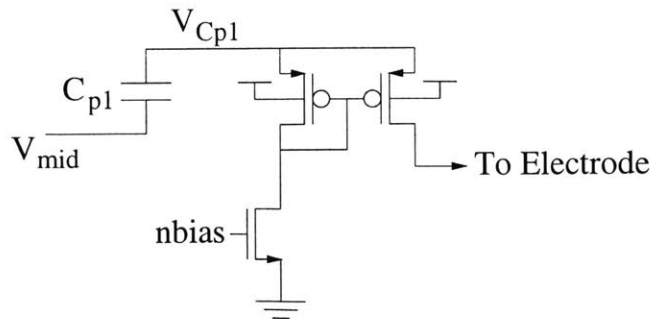


Figure 5-11: Bias Current Source

5.3 Design Problems

The ripple counter in the main state machine (Figure 5-4) will obviously give predictable state glitches. The state bits were registered in an attempt to better synchronize state changes and reduce glitching. The 1.25 ms clock was registered as well to delay it one clock cycle and synchronize it to the state bits. These registers reduced, but did not eliminate state glitching, so the switch controls were registered as shown in Figure 5-10.

5.4 Layout Issues

The switches and switch drivers take substantial chip area, and are repeated 15 times. Early attempts yielded layouts which were individually compact, but did not fit well with each other and with the pad frame. A different approach was then taken to ensure compact layout and easy connections to the pad frame and digital control circuitry.

The width of a single pad, 223λ , and the placement of pad, V_{dd} , and ground connections were considered, and the switches and drivers were fit into this box as compactly as possible. Figure 5-12 shows, at left, the empty box as defined by the dimensions of a single pad, and at right, the switches and drivers for one electrode within this box.




Figure 5-12: Switch and Driver Layout

At left, the box defined by a single pad. At right, the electrode switches, switch drivers, and logic, positioned to fit in the box and stack easily side by side.

5.5 Electrode Drive System Testing

Because of the limited number of pins on the chip, very few pins were available for monitoring the digital system. Most of the testing information instead came from observing the electrode pins to see that they were switched to the correct voltages in the correct sequence.

5.5.1 1.25 ms Clock

The one digital signal which was brought out to a pin was the 1.25 *ms* clock. As shown in Figure 5-5, this signal was designed to take a 32 μs clock, be normally high, and pulse low for one clock cycle out of every 39 (1.248 *ms*). Measurements indicate that this pulse width is 32.5 μs (due to a 2% error in the class E transmitter frequency) when no electrodes are being stimulated, but occasionally has a width of approximately 12, 16, or 24 μs when stimulation is turned on. The clock period is somewhat more stable, at about 1.26 *ms* (again, because of the transmitter frequency) with no stimulation, and ranging from 1.24 to 1.26 *ms* with stimulation. This is probably due to a problem with the divide by 2 flip-flop in Figure 5-6. It can easily be seen that a slowly rising or falling clock signal into the divide by 2 circuit can turn

on both the N and P differential pairs at the same time. This creates the potential for an oscillation, lasting until the clock signal turns off one of the differential pairs, and ending in an unknown state. Figure 5-13 shows this bug in timing diagrams. On the left is the ideal operation, on the right can be seen one example of the effect of a slow clock. In this example, Q_n begins to change on the rising clock edge, but as Q_n falls, Q_{out} begins to fall because the N differential pair is still on. This forces Q_n back up, and by this time the clock signal has reached its high value. This pattern can happen every cycle, only on rising or only on falling edges, or at completely random intervals. In this case, it is likely happening when there is sufficient substrate noise from the stimulation circuitry.




Figure 5-13: Divide by 2 Timing Diagram

At left, the ideal operation. At right, one possible effect of a slowly changing clock.

5.5.2 Electrode Patterns

The electrode output nodes were tested for each stimulation pattern to verify that the the appropriate electrodes were driven in the right sequence at the right frequency. All electrodes were confirmed to be off in pattern 00. The “X” pattern gave 50 Hz stimulation (sweeps through all capacitors in 10 *ms*, then sits for 10 *ms*) for electrodes 1, 5, 8, 11, and 15, as designed. The electrodes were appropriately separated and phase-shifted by their drive groups (A, B, C, D). The next pattern, also working as designed, gave 100 Hz (constant) stimulation to all 15 electrodes.

The final, “MIT” pattern, did not work as designed. Electrodes were separated by drive groups, and stimulation was at 50 Hz, but the “M I T -” loop did not work properly. It is likely that this is due to the state glitching discussed in Section 5.3.

Note that in Figure 5-9, the registers are clocked by the output of combinational logic which monitors the 4 electrode stimulation state bits (Q_8, Q_9, Q_{10}, Q_{11}). Glitching in these bits has the potential to generate state 1111 in a glitch, for example, when transitioning from 0111 to 1000. Sadly, this particular problem would not have occurred with the predictable glitching of the original ripple counter.

5.5.3 Driving a Series RC Load

The final test involving only the chip was to drive a series resistance (nominal, $1.5\text{ K}\Omega$; measured, $1.48\text{ K}\Omega$) and capacitance (nominal, $1\text{ }\mu\text{F}$; measured, 670 nF) from an electrode output pin. This arrangement allows monitoring of the capacitor voltage and easy calculation of the current, from $I = V_R/R = (V_{total} - V_C)/R$. The results are displayed in Figure 5-14, and are not substantially different from the calculated and measured voltage and current waveforms from Chapter 2 (Figures 2-13, 2-14, 2-20, 2-21).




Figure 5-14: Stimulation System Driving a Series Resistor and Capacitor

5.6 Conclusion

This portion of the system had more problems than more complicated portions. Two changes would have helped reduce errors, or at least make them easier to debug. A switch-mode simulator should have been used during the design phase. A full-model simulator was used for small subcircuits, but the entire digital circuit could not be simulated in this manner. The other change would have been to use 9 electrode output pins instead of 15, since there are only 9 unique signals. The switches are large enough to handle 2 electrodes in parallel, especially considering the large electrode resistance, and this change would have freed up 6 additional pins for debugging digital control signals.

This concludes the discussion of the three main portions of the stimulation system. We will now move on to discuss integration of these 3 parts and system testing.

Chapter 6

System Integration

6.1 Final VLSI Chip




Figure 6-1: Layout and Chip Micrograph of Full VLSI System

The analog and mixed-signal circuitry from Chapter 4 and the digital and switching circuitry from Chapter 5 were combined on the chip shown in Figure 6-1. At left is the layout, at right is a chip micrograph. The analog portion is on the top, the digital is below, and a horizontal row of substrate contacts (approximately through the middle of the chip) separate the two halves.

Figure 6-2 shows the micrograph in greater detail, with labels showing the function




Figure 6-2: Labeled Chip Micrograph of Full VLSI System

of each part of the chip. Particularly visible are the poly-poly2 voltage reference resistors in the upper left, the voltage reference MOS capacitors just to the right of those, the five identical clocked comparators in the middle left, the power supply decoupling capacitors across the middle right side, and the 15 identical switch / switch driver blocks running along the edge of the bottom half of the pad frame.

6.2 Full System

The coil jig from Chapter 3 serves as the mechanical foundation for the system, with the primary driver from Figure 3-9 attached to one of the plexiglass shelves. The other shelf holds the secondary PC board, shown in Figure 6-3.




Figure 6-3: Secondary Side Circuit Board

This board includes the chip from the above figures, the Schottky diode rectifier from Figure 4-1 (top, surface mount), a connector for the secondary coil (top, not visible, on the underside), a connector for the electrode cable (right), switches for digital control inputs (bottom), a potentiometer for setting the bias node for chip pad follower circuits (left), and of course, the 5 storage capacitors (4 at upper left, 1 at upper right). In addition, single row strip headers are used throughout the board for testing access.

The electrode array used contains 15 electrodes in a 3x5 grid, as shown in pattern Figures 5-1 and 5-2. This array, shown in Figure 6-4, is made of polyimide and contains 400 μm diameter iridium oxide electrodes, and thus is very similar to the electrode array shown in Figure 2-19 and used in bench tests described in Chapter 2. As stated previously, in Section 1.2.2, the final electrode resistance was approximately twice the target value, and the capacitance was roughly half the target value. In all of the measurements to follow, two electrode arrays were configured in parallel, and two electrodes in parallel are treated as a single electrode.




Figure 6-4: 15-Electrode Array

Finally, the entire system is shown in Figure 6-5. It includes the class E amplifier (far right), which drives the primary coil (just below the class E), coupling to the small secondary coil (barely visible to the left of the primary), which powers the secondary-side stimulator board (center), which drives the electrodes (attached to the green PC board in the background).

6.3 Full System Testing

The class E amplifier is powered by 2 separate power supplies, 10 V for the LM555 circuit and switch drivers, and 2-6 V for the choke. The system was turned on with a digital voltmeter monitoring the secondary system V_{dd} , and the choke supply voltage was slightly adjusted to ensure a secondary power supply of approximately 3.5 V.




Figure 6-5: Full Power and Stimulation System

6.3.1 Driving Electrodes

The electrodes were placed in the same saline irrigation solution used in Chapter 2, but with a separate, very large ($\sim 1 \text{ cm}^2$) platinum return electrode, connected to V_{mid} . Electrode current was measured by placing a resistor in series with one of the electrodes, and passing the voltage on either side of the resistor through a standard instrumentation amplifier, shown in Figure 6-6. Its component values are listed in Table 6.1.




Figure 6-6: Instrumentation Amplifier for Measuring Electrode Current

Component	Value/Part	Component	Value/Part
V_{dd}/V_{ss}	$\pm 10V$	A1, A2, A3	LT1355
C_{hp}	$1 \mu F$	R_{hp}	$100 K\Omega$
R_S	9.79Ω	R_1	$3.3 K\Omega$
R_2	$15 K\Omega, 1\%$	R_3	$13 K\Omega, 1\%$
Potentiometer	$10 K\Omega$ trim pot		

Table 6.1: Electrode Current Instrumentation Amplifier Components

Twisted pair wire was run from the electrode wires to the amplifier, where the sense resistor was physically located. The high-pass filter at the input gives AC coupling, ensuring that the inputs to op amps A1 and A2 are within the voltage rails. The gain of the instrumentation amplifier is $A_v = 1 + 2R_2/R_1$, or about 10. The gain of the non-inverting amplifier stage depends on the potentiometer setting, and was adjusted to get the desired overall voltage gain of 400. The amplifier output is:

$$V_{out} = A_v I_{elec} R_S = 3916 I_{elec} \quad (6.1)$$

Figure 6-7 shows the measured electrode current and voltage. The dashed line in the voltage plot, at 1.76 V, is V_{mid} . At the output of the amplifier, the current waveform has substantial pickup from the AC magnetic field, and some offset from the op amps. This was removed by measuring the no-current state and subtracting that output, with some phase alignment, from the output of the amplifiers. The waveform is still slightly fuzzy due to sampling rate limitations in the Tektronix TDS3014 oscilloscope. In order to capture over 10 *ms* of data, the 8 μ s sinusoid has a poor sampling rate.




Figure 6-7: Measured Electrode Current and Voltage Waveforms

The current shown in Figure 6-7 sums to $0.678 \mu C$ per phase, somewhat below the target value of $1 \mu C$, but still close to the human perceptual threshold.

6.3.2 System Power

Power consumed by the entire system was calculated by multiplying the secondary coil current by the secondary coil voltage. This gives the power consumption for everything downstream from the secondary coil. The I^2R power burned within the coil must then be added.

Measuring Secondary Coil Current and Voltage

Figure 6-8 shows the instrumentation amplifier used to measure the coil current, and the component values are listed in Table 6.2. It is very similar to the amplifier for measuring electrode current, but with 2 input stages, each giving gain of about 11. The non-inverting amplifier stage was adjusted to give overall gain of 500, yielding:

$$V_{out} = A_v I_{coil} R_S = 245 I_{coil} \quad (6.2)$$

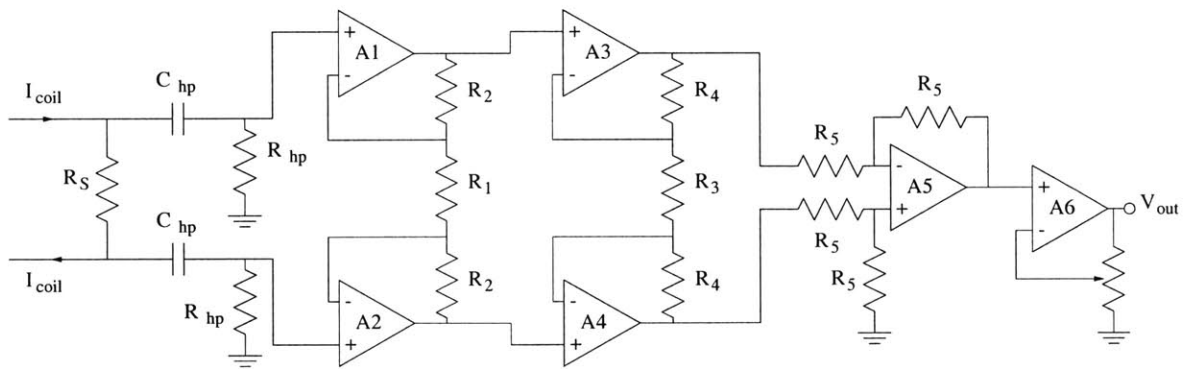


Figure 6-8: Instrumentation Amplifier for Measuring Secondary Coil Current

Figure 6-9 shows the instrumentation amplifier used to measure the coil voltage, with component values are listed in Table 6.3. This circuit has unity voltage gain, but provides the voltage difference across the secondary coil. This allows us to use the oscilloscope to monitor the coil voltage and current simultaneously on 2 channels using a common ground, which is necessary for performing power calculations.

Component	Value/Part	Component	Value/Part
V_{dd}/V_{ss}	$\pm 10V$	A1, A2, A3, A4	AD826
R_S	$0.49\ \Omega$	A5, A6	LT1355
C_{hp}	$0.22\ \mu F$	R_{hp}	$10\ K\Omega$
R_1	$2\ K\Omega$	R_2	$10\ K\Omega$
R_3	$2\ K\Omega$	R_4	$10\ K\Omega$
R_5	$1.15\ K\Omega, 1\%$	Potentiometer	$1\ K\Omega$ trim pot

Table 6.2: Coil Current Instrumentation Amplifier Components

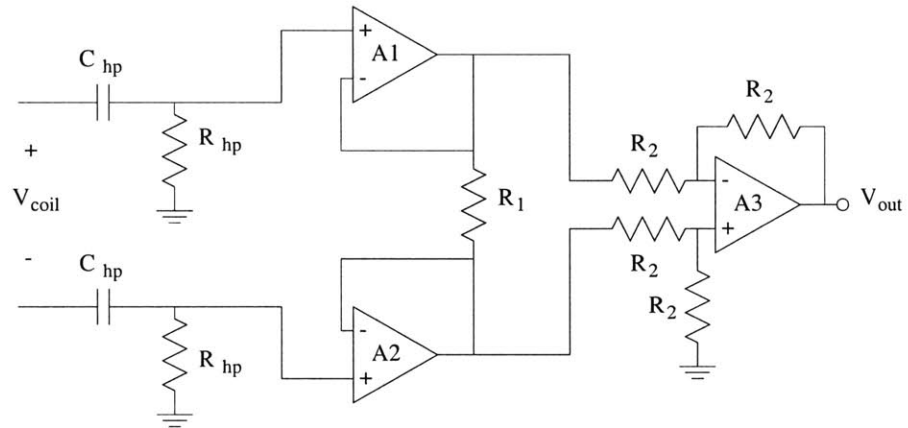


Figure 6-9: Instrumentation Amplifier for Measuring Secondary Coil Voltage

Component	Value/Part	Component	Value/Part
V_{dd}/V_{ss}	$\pm 10V$	A1, A2, A3	LT1355
C_{hp}	$0.1\ \mu F$	R_{hp}	$180\ K\Omega$
R_1	$10\ K\Omega$	R_2	$13\ K\Omega, 1\%$

Table 6.3: Coil Voltage Instrumentation Amplifier Components

Calculating Total System Power

Current and voltage waveforms were captured in Matlab and the waveforms were multiplied, yielding the instantaneous power for the entire implant system, with the exception of the losses in the secondary coil. However, the instantaneous power is very dependent upon whether a capacitor is being charged at that time, so some time-averaging is necessary. Since between 10 and 50 cycles may be recorded at one time without compromising sampling rate, average power was calculated for each such recording, and many such recordings were taken, and their powers averaged.




Figure 6-10: Coil Voltage and Current Waveforms, 15 Electrodes

Figure 6-10 shows small segments of typical voltage and current waveforms for the coil during heavy stimulation (15 electrodes, driven at 100 Hz). Though the waveforms look noisy, certain features are easy to pick out. Note the loading of the coil voltage on the rising part of the positive side of the waveform, and subsequent ringing. This is from the rectifier turning on for a short while, then turning off. A similar effect can be seen on the negative side, but the rectifier switch is only turned off for a brief time, then it is turned on again. You can see three of the negative




Figure 6-11: Coil Voltage and Current Waveforms, No Electrodes

waveforms which have no such distortion, and corresponding periods of practically zero current. At these times, no capacitor below V_{mid} needed to be charged. Note that the current peaks to nearly 10 mA . Figure 6-11 shows voltage and current waveforms for the secondary coil when no electrodes are being driven. Note that the current waveform is much lower in amplitude than in the previous figure. Even during brief charging events on the negative side, the current rarely exceeds 2 mA , and never reaches 4 mA .

Figure 6-12 shows a histogram of the power consumed while driving 15 electrodes at 100 Hz at the amplitude shown in Figure 6-7 ($0.678\text{ }\mu\text{C}$ per phase). The top histogram shows I^2R power consumed in the coil, the bottom shows power consumed in the rest of the system. Note that the coil power is far lower than the system power. Note also that the coil power looks bimodal, with the two peaks corresponding to recordings during periods of heavier and lighter charging loads. The same bimodality can be seen in the system power, but the peaks are far more spread out. The total power is the sum of the coil and system power, and the mean total power consumption




Figure 6-12: Power Consumption Histogram for 15 Electrodes
Coil power is shown in the top histogram, system power (not including the coil), in the bottom.

is 2.22 mW , with a standard deviation of 7.03×10^{-4} . Figure 6-13 shows the same type of histograms for the system with no stimulation. The mean total power when no electrodes are driven is $338 \mu\text{W}$, with a standard deviation of 2.02×10^{-5} . This data is summarized more completely in Table 6.4.

	15 Electrodes		No Electrodes	
	Mean	St.Dev.	Mean	St.Dev.
Coil Power	$44.4 \mu\text{W}$	1.41×10^{-5}	$2.79 \mu\text{W}$	1.10×10^{-7}
System Power	2.18 mW	6.91×10^{-4}	$335 \mu\text{W}$	2.01×10^{-5}
Total Power	2.22 mW	7.03×10^{-4}	$338 \mu\text{W}$	2.02×10^{-5}

Table 6.4: Power Measurement Data Summary

If the total power to drive 15 electrodes to $0.678 \mu\text{C}$ at 100 Hz is 2.22 mW , and the system consumes $338 \mu\text{W}$ with no electrodes driven, then electrode stimulation, including all of the system inefficiencies that arise because of it, consumes 1.88 mW , or $125 \mu\text{W}$ per electrode. For reference, a very aggressive current source design,




Figure 6-13: Power Consumption Histogram for No Electrodes
Coil power is shown in the top histogram, system power (not including the coil) in the bottom.

using the same $\pm 1.75\text{ V}$ voltage rails as this system and the same Schottky diode rectifiers, would use $271\ \mu\text{W}$ per electrode to deliver the same charge through the same electrodes at the same frequency:

$$\left(\frac{0.678\ \mu\text{C}}{5\ \text{ms}}\right) (1.75\text{ V} + 0.25\text{ V}) = 271\ \mu\text{W} \quad (6.3)$$

This equation calculates power during one phase, from one supply. The calculation from the opposite supply is identical, and since current is always drawn from one supply or the other, this represents the average power. The additional 0.25 V comes from the Schottky diode.

This represents a 53.7% savings over a very aggressive traditional design. A more typical current source design uses $\pm 2.5\text{ V}$ rails (though many use even higher rails). With all else being equal, this typical design uses $373\ \mu\text{W}$ per electrode, nearly three times the power consumed by our design.

Calculating Efficiency

We have determined that our system consumes $125 \mu W$ per electrode, but it is useful to know several more power numbers in order to get a better feel for the system's efficiency. From the charge per phase ($0.678 \mu C$) and the electrode resistance (two $2.31 K\Omega$ electrodes in parallel), we can determine the minimum possible I^2R power required to deliver this charge to an electrode, using the optimal constant current waveform:

$$\left(\frac{0.678 \mu C}{5 ms}\right)^2 \left(\frac{2.31 K\Omega}{2}\right) = 21.2 \mu W \quad (6.4)$$

Furthermore, from Figure 6-7, we can directly calculate the actual power into an electrode by multiplying the electrode current at every point by the voltage difference from V_{Cp1} , the electrode bias voltage. This gives the instantaneous power into the electrode, as shown in Figure 6-14. As predicted, the power in the first phase is positive, and the power at the beginning of the second phase is negative, showing energy being recovered from the electrode. The positive power at the end of the stimulation ensures that the electrode capacitance is discharged by the end of the cycle.

Integrating all of this power results in a net average power of $49.0 \mu W$, about 2.3 times the minimum possible power, leaving significant room for later improvement. If only the power *into* the electrode is considered, it amounts to $56.8 \mu W$, or 2.67 times the minimum possible.

With this information, we can calculate several efficiency values of interest. Efficiency is the fraction of power that is useful, and we have a choice of two numbers for total power and three numbers for useful power. For total power, we may use the total average system power, including overhead power, or we may use the average additional power required for the stimulation of 15 electrodes. For useful power, we may consider the net average power delivered to the electrodes, the average power delivered only *to* (not from) the electrodes, or the minimum power possible to deliver the charge. These six efficiencies are summarized in Table 6.3.2, along with efficien-




Figure 6-14: Power Into an Electrode During a Stimulation Cycle

gies for the aggressively designed current source and typically used current source (not including any overhead power).

If our system is to be compared with systems of unknown overhead power, the fourth, fifth, and sixth numbers in the table (without overhead) should be used. In this case, our system transports the required energy from the magnetic field to the electrode with over 39% efficiency. When considering the minimum possible required to charge the electrode, the efficiency drops to just under 17%, which, as expected from the earlier calculations, is more than twice the efficiency of the aggressive current source design under the same conditions, and nearly three times the efficiency of typically used current source designs.

These numbers show that there is still room for improvement. The overhead power is less of a concern, since electrode numbers for this application are likely to increase, but the 39.1% and 16.9% efficiency numbers are worth examining. This is explored further in Section 7.2.

Total System Power	2.22 mW	
Power <i>Into</i> Electrodes	15 x 56.8 μ W	38.4%
Total System Power	2.22 mW	
Net Power To Electrodes	15 x 49.0 μ W	33.1%
Total System Power	2.22 mW	
Minimum Possible Electrode Power	15 x 21.2 μ W	14.3%
Total - Overhead Power	1.88 mW	
Power <i>Into</i> Electrodes	15 x 56.8 μ W	45.3%
Total - Overhead Power	1.88 mW	
Net Power To Electrodes	15 x 49.0 μ W	39.1%
Total - Overhead Power	1.88 mW	
Minimum Possible Electrode Power	15 x 21.2 μ W	16.9%
Aggressive Current Source Design	15 x 271 μ W	
Minimum Possible Electrode Power	15 x 21.2 μ W	7.8%
Typical Current Source Design	15 x 373 μ W	
Minimum Possible Electrode Power	15 x 21.2 μ W	5.7%

Table 6.5: Power Efficiencies

Chapter 7

Conclusions and Future Work

7.1 Thesis Conclusions

This thesis presents a theoretical method of drastically reducing the power consumed in tissue stimulation, as well as a prototype implementation. Of the three main sections of the prototype system, the RF power transmission system of Chapter 3 worked as planned, the power handling circuitry of Chapter 4 worked sufficiently well, with a few design flaws and a few surprises due to layout flaws, and the electrode switching circuitry of Chapter 5 worked sufficiently well, with minor flaws due to lack of sufficient modeling and simulation, and substrate pickup from the 125 KHz magnetic fields.

Table 7.1 summarizes the power savings of the system presented in the preceding chapters. In this table, the term overhead refers to the power consumed by the system when it is not driving electrodes. This includes references, control circuitry, and any standby power consumption of the stimulation system.

It is worth exploring, to the extent possible, the power losses associated with this system. First, the measured power into the electrodes, $49 \mu W$, is 2.3 times the theoretical minimum for delivering the same charge. Theoretical calculations with ideal steps showed that a 4-step system could deliver the charge using only 15% more power than the minimum. However, the steps used here are equal steps, and do not perfectly account for the IR drop offset in the electrode voltage. Another possible

Theoretical	Minimum Possible Power	21.2 μW
This System	Measured Power Into Electrode	49 μW
This System	Net Power During Stimulation (Not Including Overhead)	125 μW
This System	Gross Power During Stimulation (Including Overhead)	148 μW
Aggressive Current Source System	Net Power During Stimulation (Not Including Overhead)	271 μW
Typical Current Source System	Net Power During Stimulation (Not Including Overhead)	373 μW

Table 7.1: Power per Electrode Comparison

reason for this extra power loss would be a change in electrode impedance over time. However, the electrode impedance was measured after several weeks of testing with this system, and the average resistance and capacitance were 2.28 $K\Omega$ and 450 μF , practically the same resistance as before, with slightly less capacitance.

The bulk of the rest of the losses occur between the coil and the electrode connection. Starting from the electrodes, this includes the electrode switches and their drivers, extra power consumed in the control circuitry due to electrode loading, extra power consumed in the rectifier switches and drivers, and leakage paths from the storage capacitors. The electrode switches were large enough (PMOS $250\lambda / 2\lambda$, NMOS $100\lambda / 2\lambda$, $\lambda = 0.8 \mu m$) that they added negligible resistance compared to the electrode resistance, yet they were not so large that their drivers consumed large switching power. The control circuitry consumes practically no more power during stimulation than during standby, and its total power is less than 200 μW . The likely source of extra power consumption is the rectifier switches, and the reason for this power loss is the simplistic algorithm used to charge the storage capacitors. The rectification durations were set conservatively to supply more than enough power for the electrodes, and the durations came out longer than expected due to process variations in the one-shot delay circuit. This long rectification time allows a large voltage to develop across the rectifier switch, burning unnecessary power. Finally, the leakage resistors (very long transistors operating in their linear regions) on the storage capacitors were designed to leak very little current compared with that consumed by the

electrodes, and judging from the standby (no electrode drive) power consumption, they function as designed.

Table 7.1 shows where power was burned in each portion of the system, as a measure of total power when driving 15 electrodes. The minimum possible power is far smaller than the total power, so there should be room for improvement. The entry labeled extra power into the electrodes is the difference between the calculated electrode power and the theoretical minimum. The standby or overhead chip power was taken directly from the measurements in Table 6.4. The power consumed in the chip due to stimulation is the total measured chip and electrode power from Table 6.4 (2.18 mW) minus the standby power and power into the electrodes. The coil power numbers are taken directly from Table 6.4.

Theoretical Minimum Possible Power	$15 \times 21.2 \mu\text{W} = 318 \mu\text{W}$
Extra Power Into Electrodes	$15 \times (49 - 21.2) \mu\text{W} = 417 \mu\text{W}$
Standby (Overhead) Chip Power	$335 \mu\text{W}$
Power Consumed in Chip Due to Stimulation	$(2.18\text{ mW} - \text{above}) = 1.11\text{ mW}$
Standby (Overhead) Coil Power	$2.79 \mu\text{W}$
Coil Power Due to Stimulation	$44.4 \mu\text{W} - 2.79 \mu\text{W} = 41.6 \mu\text{W}$
Total Consumed Power	2.22 mW

Table 7.2: System Power Consumption, 15 Electrodes

From the numbers in the table, one in particular is noteworthy. The power consumed in the chip due to stimulation is quite high, half of the total power for the entire system. This is likely due to the voltage developed across the rectifier transistors due to the long rectification times.

While the system as a whole has some room for efficiency improvement, it meets the goal of using far less power than the competing stimulation systems listed so far. However, it is possible that such systems could be designed more efficiently. A current source based system with lower voltage rails will use less power. If the neural tissue responds to total charge levels, and exact current control is not necessary, then the voltage supplies may be made quite low. This approach may preclude the use of cascode transistors, and may drastically reduce the output impedance of current

sources, but those are acceptable tradeoffs for the potentially vast reduction in power. Another means of reducing current source power is to simply switch the electrode directly to its return for a brief time at the beginning of the second current phase, and then drive positive current into the electrode at the end of the second current phase to ensure charge balancing. The one caveat for these methods is that if the voltage rails get too close to the electrode voltage, the supply voltage may be too low for operation of many of the control circuits. In addition, voltage rails which are too close to the voltage of the charged electrode may need to move if the electrode impedance shifts over time. Moving the supply rails can affect performance of control circuitry designed to run from a specific supply. The best compromise is to separate the control circuitry supply, using a moderate voltage to supply the relatively low-current controls, and a low voltage to supply the high-current electrodes. However, while the highest supply can easily be generated from a simple diode rectifier, the lower voltage supply (which must vary with electrode impedance) must be generated from some form of synchronous rectifier with controls to keep it at the correct voltage. So as a current source based stimulator moves toward lower supply voltages, it begins to resemble the system described in this thesis.

7.2 Future Work

While this stimulation system worked well, it had a few bugs and a few areas where the general design could be improved.

7.2.1 Changes to This Implementation

The first change to this version of the system would be the substrate and well isolation of the clocked comparators and their clocks, as described in Section 4.3.9. In general, much of the circuitry needed to have better substrate contacts. Specifically, the digital clock divider discussed in Section 5.2.1 was too sensitive to substrate noise. A more traditional D flip-flop clock divider, while taking more area, may have been more reliable for this application.

7.2.2 Broader Changes

The first area to target design changes is the section of the system which burns the most power: the synchronous rectifier. While the predictive comparator works perfectly and the rectifier switches are quite large (PMOS $1250\lambda/2\lambda$, NMOS $500\lambda/2\lambda$), the switches are simply turned on for too long a time. This allows a voltage to develop across the switches, costing power. In addition, significant current develops in the coil, which rings when the switch is turned off. If this is done, a timing method is needed which is more precise and more controllable than the one-shots used here. The one-shot delays may be made both more precise and controllable by using current-starved inverters, but some feedback method must be used to set their timing. The ideal method is to have a linear differential amplifier monitoring each capacitor and its associated reference voltage, with a low-pass filter to take a time average over the last few waveform cycles. If the capacitor voltage slips a few tens of mV below the reference voltage, the rectification time is increased. Also, adding a brief fixed duration rectification period onto the back side of the waveform would spread out the charging over time, reducing the required peak current. This can be implemented easily with an over-compensated predictive comparator, which switches before its inputs are equal.

Another method of improving control over the rectification timing is to use a clock which is much faster than the 125 KHz system clock. This requires a phase-locked loop, which may be more complexity and power than is necessary, but it would allow completely digital control of the rectifier timing. In order to get timing resolution that would compete with the analog system suggested above, the on-chip oscillator would have to operate at 20-50 MHz, consuming substantial power.

The method of leaking charge from the storage capacitors described in Section 4.3.7 is somewhat ad hoc, and should be replaced with some intelligent method of maintaining capacitor voltage. Using a charge pump for the rectifier switch gates will work, but is not worth the effort, especially the effort of building a below-ground charge pump, which requires floating diodes. Since the power lost through this path, even

with leakage transistors described in Section 4.3.7, is rather small, a simple means can be implemented of leaking off some charge when the capacitor voltage exceeds the reference voltage by too great a margin. The same linear differential amplifier described above could be the basis for this measurement.

Switching off the current through the coil during rectification caused a large voltage swing and resonance on the AC voltage, which contributed to substrate noise. This noise may be abated with a simple snubber circuit, like a series resistance and capacitance, across the secondary coil. The addition of a snubber may slightly increase power, so this should be done in conjunction with substantial redesign of the rectification timing controls.

One circuit change may make this design more tolerable to nerves by generating more nearly uniform stimulation current without substantially increasing the power consumption. Current source transistors may be added to each voltage level, limiting the output current immediately after the electrode is switched to it. As the electrode charges to nearly the voltage on that storage capacitor, the transistor begins to act like a switch instead of a current source. This would spread the current throughout the time during which the electrode is switched to that capacitor, while delivering only slightly less charge per level than the purely switch-based method.

Finally, as mentioned in Section 3.4, some measurement of the chip V_{dd} should be sent telemetrically to whatever circuit is driving the primary coil. This is potentially a difficult problem, since the data rate should be on the same order of magnitude as the power frequency, 125 KHz. Specifically, the V_{dd} data should be sent at approximately 125 KHz divided by the Q of the primary coil resonance system, or in the tens of KHz range. The back telemetry could be included as part of a higher frequency data communication system.

Appendix A

Mathematical Derivations and MATLAB Scripts

A.1 Proof of Optimal RC Charging Profile

Reprinted from internal technical paper M84 by John Wyatt, September 17, 2001.

More Efficient Charge Transfer Method II.

Problem - To charge a capacitor C through a resistor R from an initial charge Q_0 at $t=0$ to a final charge Q_1 at $t=T$ using the least total energy.




Figure A-1: Variables for Capacitor Charging Proof

Claim - The optimal solution is to drive with a constant current

$$\dot{q}(t) = i(t) = \frac{Q_1 - Q_0}{T}, \quad 0 \leq t \leq T \quad (\text{A.1})$$

Derivation - Let the optimal charging path be described by $q^*(t)$ (to be found), subject only to

$$q^*(0) = Q_1 \quad q^*(T) = Q_2 \quad (\text{A.2})$$

Let $\delta q(t)$ be any perturbation such that

$$q^*(t) + \lambda(\delta q(t)) \triangleq q_\lambda(t) \quad (\text{A.3})$$

meets the constraints A.2, i.e.,

$$\delta q(0) = \delta q(T) = 0 \quad (\text{A.4})$$

The total energy supplied is the sum of the capacitor energy supplied

$$E_C = \frac{Q_1^2 - Q_0^2}{2C} \quad (\text{A.5})$$

which is independent of λ , and the resistor energy

$$\begin{aligned} E_R(\lambda) &= \int_0^T [\dot{q}^*(t) + \lambda(\delta \dot{q}(t))]^2 R dt = \\ &R \int_0^T (\dot{q}^*(t))^2 + 2\lambda \dot{q}^*(t)(\delta \dot{q}(t)) + \lambda^2(\delta \dot{q}(t))^2 dt \end{aligned} \quad (\text{A.6})$$

If $q^*(t)$ is optimal, then

$$\left. \frac{dE_R(\lambda)}{d\lambda} \right|_{\lambda=0} = 0 \quad (\text{A.7})$$

for *any* choice of $\delta q(t)$ satisfying A.4. Differentiating A.6 and using A.7 gives

$$2R \int_0^T \dot{q}^*(t)(\delta \dot{q}(t))dt = 0 \quad (\text{A.8})$$

for *all* choices of $\delta q(\cdot)$ satisfying A.4, i.e., for *all* choices of $\delta \dot{q}(\cdot)$ such that

$$\int_0^T \delta \dot{q}(t)dt = 0 \quad (\text{A.9})$$

But the only waveforms that are orthogonal to all zero mean waveforms are constants. Thus $\dot{q}^*(t)$ is a constant.

A.2 Capacitor Bank Number and Voltage Calculations

A.2.1 3-Step Charging

```
%% 3 Step electrode drive system. What are the optimal sources,  
%% skipping a step on the return?
```

```
%% Shawn K. Kelly, last updated May 27, 2003
```

```
% Electrode parameters
```

```
R = 1.1e3;
```

```
C = 1000e-9;
```

```
tau=R*C;
```

```
% Target Electrode Charge
```

```
Qtarg = 1e-6;
```

```
Vtarg = Qtarg/C;
```

```
Steps = 3;
```

```
% Total time limit
```

```
Thalf = 5e-3;
```

```
T3 = Thalf;
```

```
T6 = 2*Thalf;
```

```
%timestep=5e-6;
```

```
I=Qtarg/Thalf;
```

```
% Other Times
```

```
T1 = Thalf/Steps;
```

```
T2 = 2*T1;
```

```
T4 = Thalf + T1;
```

```
T5 = Thalf + 2*T1;
```

```
delT=exp(-T1/tau);
```

```
X=1-delT;
```

```
Vstep = 0:.001*Vtarg/Steps:(I*R + Vtarg/Steps);
```

```
% Source Voltages
```

```
Vp1=Vstep;
```

```

Vmid=0;
Vn1=-Vstep;

Vtarget=Vp1-Vtarg;

% Electrode capacitance voltage at various times
% Energy lost from sources = Vs*Q
VCT1 = Vp1 + (Vmid-Vp1)*X;
Elost1 = Vmid.*VCT1*C;
VCT2 = VCT1 + (Vn1-VCT1)*X;
Elost2 = Vn1.*(VCT2-VCT1)*C;
Vn2 = (Vtarget - VCT2*delT)/X;
Elost3 = Vn2.*(Vtarget-VCT2)*C;
Elostdown = Elost1 + Elost2 + Elost3;

VCT4 = Vtarget+(Vmid-Vtarget)*X;
Elost4 = Vmid.*(VCT4-Vtarget)*C;
VCT5 = VCT4 + (Vp1-VCT4)*X;
Elost5 = Vp1.*(VCT5-VCT4)*C;
Vp2 = (Vp1 - VCT5*delT)/X;
Elost6 = Vp2.*(Vp1-VCT5)*C;
Elostup = Elost4+Elost5+Elost6;
Elosttot=Elostup+Elostdown;

[Emin,Imin]=min(Elosttot)
Vdn = [Vmid Vn1(Imin) Vn2(Imin)]
Vup = [Vmid Vp1(Imin) Vp2(Imin)]
delt = .001*T1;
V = [Vup(2) Vup(2) Vdn(1) Vdn(1) Vdn(2) Vdn(2) Vdn(3) Vdn(3) ...
      Vup(1) Vup(1) Vup(2) Vup(2) Vup(3) Vup(3) Vup(2) Vup(2)];

t = [-T1 0 delt T1 T1+delt T2 T2+delt T3 T3+delt T4 T4+delt T5 ...
      T5+delt T6 T6+delt 2*Thalf+T1];

Vid = Vup(2) + [0 0 -I*R -I*R-Vtarg I*R-Vtarg I*R 0 0];

tid = [-T1 0 delt Thalf Thalf+delt 2*Thalf 2*Thalf+delt 2*Thalf+T1];

vax = [-1 0.8];
tax = [-T1 2*Thalf+T1];
figure(1)
hold off
plot(t,V)
hold on

```

```

plot(tid,Vid,'r--')
plot([0 0],vax,'k:')
plot(tax,[0 0],'k:')

axis([tax(1) tax(2) vax(1) vax(2)])
ylabel('Drive Voltage')
xlabel('Time')
legend('Real Drive Waveform','Ideal Drive Waveform',2)

timestep=.01e-3;
time1 = 0:timestep:T1;
I1 = ((Vdn(1)-Vup(2))/R)*exp(-time1/tau);
time2 = (T1+timestep):timestep:T2;
I2 = ((Vdn(2)-VCT1(Imin))/R)*exp(-(time2-T1)/tau);
time3 = (T2+timestep):timestep:T3;
I3 = ((Vdn(3)-VCT2(Imin))/R)*exp(-(time3-T2)/tau);
time4 = (T3+timestep):timestep:T4;
I4 = ((Vup(1)-Vtarget(Imin))/R)*exp(-(time4-T3)/tau);
time5 = (T4+timestep):timestep:T5;
I5 = ((Vup(2)-VCT4(Imin))/R)*exp(-(time5-T4)/tau);
time6 = (T5+timestep):timestep:T6;
I6 = ((Vup(3)-VCT5(Imin))/R)*exp(-(time6-T5)/tau);

time = [-T1 -timestep time1 time2 time3 time4 time5 time6 ...
        2*Thalf+timestep 2*Thalf+T1];
I = [0 0 I1 I2 I3 I4 I5 I6 0 0];
Iid = [0 0 -Qtarget/Thalf -Qtarget/Thalf Qtarget/Thalf ...
       Qtarget/Thalf 0 0];
iax = [-2.5*Qtarget/Thalf 2.8*Qtarget/Thalf];
figure(2)
hold off
plot(time,I);
hold on
plot(tid,Iid,'r--');
plot([0 0],iax,'k:');
plot(tax,[0 0],'k:');
axis([-T1 (2*Thalf+T1) iax(1) iax(2)])
ylabel('Drive Current')
xlabel('Time')
legend('Real Drive Waveform','Ideal Drive Waveform',2)

%%%%% RESULTS

%E = .57974 uJ (.4033) 0 -.4033 -.6816 0 .4033 .4364

```

A.2.2 4-Step Charging

```
%% 4 Step electrode drive system. What are the optimal sources,  
%% skipping a step on the return?
```

```
%% Shawn K. Kelly, last updated May 27, 2003
```

```
% Electrode parameters
```

```
R = 1.1e3;
```

```
C = 1000e-9;
```

```
tau=R*C;
```

```
% Target Electrode Charge
```

```
Qtarg = 1e-6;
```

```
Vtarg = Qtarg/C;
```

```
Steps = 4;
```

```
% Total time limit
```

```
Thalf = 5e-3;
```

```
T4 = Thalf;
```

```
T8 = 2*Thalf;
```

```
I=Qtarg/Thalf;
```

```
% Other Times
```

```
T1 = Thalf/Steps;
```

```
T2 = 2*T1;
```

```
T3 = 3*T1;
```

```
T5 = Thalf + T1;
```

```
T6 = Thalf + 2*T1;
```

```
T7 = Thalf + 3*T1;
```

```
delt=exp(-T1/tau);
```

```
X=1-delt;
```

```
Vstep = .2:.001:.3;
```

```
% Source Voltages
```

```
Vp1=Vstep;
```

```
Vmid=0;
```



```

Vn1=-Vstep;
Vn2=-2*Vstep;

Vtarget=Vp1-Vtarg;

% Electrode capacitance voltage at various times
% Energy lost from sources = Vs*Q
foo = 1
VCT1 = Vp1 + (Vmid-Vp1)*X;
Elost1 = Vmid.*VCT1*C;
VCT2 = VCT1 + (Vn1-VCT1)*X;
Elost2 = Vn1.*(VCT2-VCT1)*C;
VCT3 = VCT2 + (Vn2-VCT2)*X;
Elost3 = Vn2.*(VCT3-VCT2)*C;
Vn3 = (Vtarget - VCT3*delT)/X;
Elost4 = Vn3.*(Vtarget-VCT3)*C;
Elostdown = Elost1 + Elost2 + Elost3 + Elost4;

VCT5 = Vtarget+(Vn1-Vtarget)*X;
Elost5 = Vn1.*(VCT5-Vtarget)*C;
VCT6 = VCT5 + (Vmid-VCT5)*X;
Elost6 = Vmid.*(VCT6-VCT5)*C;
VCT7 = VCT6 + (Vp1-VCT6)*X;
Elost7 = Vp1.*(VCT7-VCT6)*C;
Vp2 = (Vp1 - VCT7*delT)/X;
Elost8 = Vp2.*(Vp1-VCT7)*C;
Elostup = Elost5+Elost6+Elost7+Elost8;
Elosttot=Elostup+Elostdown;

[Emin,Imin]=min(Elosttot)
Vdn = [Vmid Vn1(Imin) Vn2(Imin) Vn3(Imin)]
Vup = [Vn1(Imin) Vmid Vp1(Imin) Vp2(Imin)]
delT = .001*T1;
V = [Vup(3) Vup(3) Vdn(1) Vdn(1) Vdn(2) Vdn(2) Vdn(3) Vdn(3)...
      Vdn(4) Vdn(4) Vup(1) Vup(1) Vup(2) Vup(2) Vup(3) Vup(3)...
      Vup(4) Vup(4) Vup(3) Vup(3)];

t = [-T1 0 delT T1 T1+delT T2 T2+delT T3 T3+delT T4 T4+delT T5...
      T5+delT T6 T6+delT T7 T7+delT T8 T8+delT T8+T1];

Vid = Vup(3) + [0 0 -I*R -I*R-Vtarg I*R-Vtarg I*R 0 0];

tid = [-T1 0 delT T4 T4+delT T8 T8+delT T8+T1];

vax = [-1.2 0.6];

```

```

tax = [-T1 T8+T1];
figure(1)
hold off
plot(t,V)
hold on
plot(tid,Vid,'r--')
plot([0 0],vax,'k:')
plot(tax,[0 0],'k:')

axis([tax(1) tax(2) vax(1) vax(2)])
ylabel('Drive Voltage')
xlabel('Time')
legend('Real Drive Waveform','Ideal Drive Waveform',2)

timestep=.01e-3;
time1 = 0:timestep:T1;
I1 = ((Vdn(1)-Vup(3))/R)*exp(-time1/tau);
time2 = (T1+timestep):timestep:T2;
I2 = ((Vdn(2)-VCT1(Imin))/R)*exp(-(time2-T1)/tau);
time3 = (T2+timestep):timestep:T3;
I3 = ((Vdn(3)-VCT2(Imin))/R)*exp(-(time3-T2)/tau);
time4 = (T3+timestep):timestep:T4;
I4 = ((Vdn(4)-VCT3(Imin))/R)*exp(-(time4-T3)/tau);
time5 = (T4+timestep):timestep:T5;
I5 = ((Vup(1)-Vtarget(Imin))/R)*exp(-(time5-T4)/tau);
time6 = (T5+timestep):timestep:T6;
I6 = ((Vup(2)-VCT5(Imin))/R)*exp(-(time6-T5)/tau);
time7 = (T6+timestep):timestep:T7;
I7 = ((Vup(3)-VCT6(Imin))/R)*exp(-(time7-T6)/tau);
time8 = (T7+timestep):timestep:T8;
I8 = ((Vup(4)-VCT7(Imin))/R)*exp(-(time8-T7)/tau);

time = [-T1 -timestep time1 time2 time3 time4 time5 time6...
        time7 time8 T8+timestep T8+T1];
I = [0 0 I1 I2 I3 I4 I5 I6 I7 I8 0 0];
Iid = [0 0 -Qtarget/T4 -Qtarget/T4 Qtarget/T4 Qtarget/T4 0 0];
iax = [-2*Qtarget/T4 2*Qtarget/T4];
figure(2)
hold off
plot(time,I);
hold on
plot(tid,Iid,'r--');
plot([0 0],iax,'k:');
plot(tax,[0 0],'k:');
axis([-T1 (T8+T1) iax(1) iax(2)])

```

```
ylabel('Drive Current')
xlabel('Time')
legend('Real Drive Waveform','Ideal Drive Waveform',2)
```

```
%%%% RESULTS
```

```
%E = .50634 uJ (.297) 0 -.297 -.594 -.8187 -.297 0 .297 .3629
```

A.2.3 5-Step Charging

```
%% 5 Step electrode drive system.  What are the optimal sources
%% skipping a step on the return?
```

```
%% Shawn K. Kelly, last updated May 27, 2003
```

```
% Electrode parameters
```

```
R = 1.1e3;
C = 1000e-9;
tau=R*C;
```

```
% Target Electrode Charge
```

```
Qtarg = 1e-6;
Vtarg = Qtarg/C;
```

```
Steps = 5;
```

```
% Total time limit
```

```
Thalf = 5e-3;
T5 = Thalf;
T10 = 2*Thalf;
I=Qtarg/Thalf;
```

```
% Other Times
```

```
T1 = Thalf/Steps;
T2 = 2*T1;
T3 = 3*T1;
T4 = 4*T1;
```

```
T6 = Thalf + T1;
T7 = Thalf + 2*T1;
T8 = Thalf + 3*T1;
T9 = Thalf + 4*T1;
```

```
delT=exp(-T1/tau);
X=1-delT;
```

```
Vstep = 0:.001*Vtarg/Steps:(I*R + Vtarg/Steps);
```

```
% Source Voltages
```

```
Vp1=Vstep;
```

```

Vmid=0;
Vn1=-Vstep;
Vn2=-2*Vstep;
Vn3=-3*Vstep;

Vtarget=Vp1-Vtarg;

% Electrode capacitance voltage at various times
% Energy lost from sources = Vs*Q
foo = 1
VCT1 = Vp1 + (Vmid-Vp1)*X;
Elost1 = Vmid.*VCT1*C;
VCT2 = VCT1 + (Vn1-VCT1)*X;
Elost2 = Vn1.*(VCT2-VCT1)*C;
VCT3 = VCT2 + (Vn2-VCT2)*X;
Elost3 = Vn2.*(VCT3-VCT2)*C;
VCT4 = VCT3 + (Vn3-VCT3)*X;
Elost4 = Vn3.*(VCT4-VCT3)*C;
Vn4 = (Vtarget - VCT4*delT)/X;
Elost5 = Vn4.*(Vtarget-VCT4)*C;
Elostdown = Elost1 + Elost2 + Elost3 + Elost4 + Elost5;

VCT6 = Vtarget+(Vn2-Vtarget)*X;
Elost6 = Vn2.*(VCT6-Vtarget)*C;
VCT7 = VCT6 + (Vn1-VCT6)*X;
Elost7 = Vn1.*(VCT7-VCT6)*C;
VCT8 = VCT7 + (Vmid-VCT7)*X;
Elost8 = Vmid.*(VCT8-VCT7)*C;
VCT9 = VCT8 + (Vp1-VCT8)*X;
Elost9 = Vp1.*(VCT9-VCT8)*C;
Vp2 = (Vp1 - VCT9*delT)/X;
Elost10 = Vp2.*(Vp1-VCT9)*C;
Elostup = Elost6+Elost7+Elost8+Elost9+Elost10;
Elosttot=Elostup+Elostdown;

[Emin,Imin]=min(Elosttot)
Vdn = [Vmid Vn1(Imin) Vn2(Imin) Vn3(Imin) Vn4(Imin)]
Vup = [Vn2(Imin) Vn1(Imin) Vmid Vp1(Imin) Vp2(Imin)]
delT = .001*T1;
V = [Vup(4) Vup(4) Vdn(1) Vdn(1) Vdn(2) Vdn(2) Vdn(3) Vdn(3)...
      Vdn(4) Vdn(4) Vdn(5) Vdn(5) Vup(1) Vup(1) Vup(2) Vup(2)...
      Vup(3) Vup(3) Vup(4) Vup(4) Vup(5) Vup(5) Vup(4) Vup(4)];

t = [-T1 0 delT T1 T1+delT T2 T2+delT T3 T3+delT T4 T4+delT T5...
      T5+delT T6 T6+delT T7 T7+delT T8 T8+delT T9 T9+delT T10...

```

```

    T10+delt T10+T1];

Vid = Vup(4) + [0 0 -I*R -I*R-Vtarg I*R-Vtarg I*R 0 0];

tid = [-T1 0 delt T5 T5+delt T10 T10+delt T10+T1];

vax = [-1.2 0.6];
tax = [-T1 T10+T1];
figure(1)
hold off
plot(t,V)
hold on
plot(tid,Vid,'r--')
plot([0 0],vax,'k:')
plot(tax,[0 0], 'k:')

axis([tax(1) tax(2) vax(1) vax(2)])
ylabel('Drive Voltage')
xlabel('Time')
legend('Real Drive Waveform','Ideal Drive Waveform',2)

timestep=.01e-3;
time1 = 0:timestep:T1;
I1 = ((Vdn(1)-Vup(4))/R)*exp(-time1/tau);
time2 = (T1+timestep):timestep:T2;
I2 = ((Vdn(2)-VCT1(Imin))/R)*exp(-(time2-T1)/tau);
time3 = (T2+timestep):timestep:T3;
I3 = ((Vdn(3)-VCT2(Imin))/R)*exp(-(time3-T2)/tau);
time4 = (T3+timestep):timestep:T4;
I4 = ((Vdn(4)-VCT3(Imin))/R)*exp(-(time4-T3)/tau);
time5 = (T4+timestep):timestep:T5;
I5 = ((Vdn(5)-VCT4(Imin))/R)*exp(-(time5-T4)/tau);
time6 = (T5+timestep):timestep:T6;
I6 = ((Vup(1)-Vtarget(Imin))/R)*exp(-(time6-T5)/tau);
time7 = (T6+timestep):timestep:T7;
I7 = ((Vup(2)-VCT6(Imin))/R)*exp(-(time7-T6)/tau);
time8 = (T7+timestep):timestep:T8;
I8 = ((Vup(3)-VCT7(Imin))/R)*exp(-(time8-T7)/tau);
time9 = (T8+timestep):timestep:T9;
I9 = ((Vup(4)-VCT8(Imin))/R)*exp(-(time9-T8)/tau);
time10 = (T9+timestep):timestep:T10;
I10 = ((Vup(5)-VCT9(Imin))/R)*exp(-(time10-T9)/tau);

time = [-T1 -timestep time1 time2 time3 time4 time5 time6...
        time7 time8 time9 time10 T10+timestep T10+T1];

```

```

I      = [0 0 I1 I2 I3 I4 I5 I6 I7 I8 I9 I10 0 0];
Iid    = [0 0 -Qtarg/Thalf -Qtarg/Thalf Qtarg/Thalf...
          Qtarg/Thalf 0 0];
iax    = [-2*Qtarg/Thalf 2*Qtarg/Thalf];
figure(2)
hold off
plot(time,I);
hold on
plot(tid,Iid,'r--');
plot([0 0],iax,'k:');
plot(tax,[0 0], 'k:');
axis([-T1 ((2*Thalf)+T1) iax(1) iax(2)])
ylabel('Drive Current')
xlabel('Time')
legend('Real Drive Waveform','Ideal Drive Waveform',2)

%%%%% RESULTS

%E = .48074 uJ (.2342) 0 -.2342 -.4684 -.7026 -.9123 0 .2342 .3391

```

A.3 Exact Magnetic Field Calculations

```
% Real magnetic field calculations!
% Shawn K. Kelly, last updated March 2003

% Magnetic permeability
mu = 4e-7*pi;
% Frequency
w = 2*pi*125000;

% Secondary and Primary Coil Radius, Number of turns
Rs= 5.5e-3;
Ns=60;

Rp= 18.5e-3;
Np = 45;
Ip = .8798;

z=15e-3;

dalpha = pi/180;
alpha = dalpha:dalpha:(2*pi);

H0 = (Np*Ip/2) * (Rp^2/((Rp^2 + z^2)^(1.5)))
da = Rs/200;
for i=1:200;
    a=i*da;
    aa(i)=a;
    x=a*sin(alpha);
    y=Rp - a*cos(alpha);
    r=sqrt(x.^2 + y.^2 + z^2);
    dH=(Np*Ip/(4*pi)) * (Rp./r.^2) * sin(acos(x/r)).* dalpha;
    %% find dH parallel to z, or dH cos beta
    dHnorm = dH.*y./sqrt(z^2 + y.^2);
    H(i) = sum(dHnorm);
    if (a==da)
        flux(i)=mu*mean([H(i) H0])*pi*a^2;
    else
        flux(i)=mu*mean([H(i) H(i-1)])*pi*(a^2 - (a-da)^2);
    end
end

end

plot(aa,H)
```



```
fluxtot=sum(flux);  
Hav = fluxtot/mu/pi/Rs^2  
Voc = fluxtot*w*Ns  
ylabel('Field Strength (A/m)')  
xlabel('Distance from Center of Secondary')
```

A.4 Geometric View of Coil Loading and Resonance

```
% Geometric view of secondary coil loading
% Shawn Kelly, 12/3/2002
% updated 4/7/03
% updated 6/5/03

%Constants
mu      = pi*4e-7;
rho     = 1.7e-8;
sigma  = 1.79;
reye   = 23.5e-3 / 2;

%Magnetic Field Parameters
rp      = 18.5e-3;
z       = 15e-3;
Hmult  = 3.46;
f       = 125000;
w       = 2*pi*f;
Hnom    = sqrt(2)*16.3e6/f;
H       = Hmult*Hnom*rp^3/((rp^2 + z^2)^1.5);

%Coil Parameters
rmin    = 4.3e-3;
rmax    = 6.5e-3;
t       = 700e-6;
alph    = pi/4;
N       = 60;

%Calculated Coil Parameters
rav     = mean([rmin rmax]);
A       = t*(rmax-rmin);
Vol     = t*pi*(rmax^2-rmin^2);
lR      = 2*pi*rav*N;
AR      = alph*A/N;

%Coil Model
Lambda  = mu*H*(N*pi*rav^2);
Voc     = w*Lambda
Rs      = rho*lR/AR
%L      = pi*mu*rav*N^2/2;
```

```

L      = 54e-6;
Pav    = Voc^2/(8*Rs); %Available, not average

%Load
Pmin   = 10e-3;
Pmin1  = 5e-3;
Pmin2  = 15e-3;
Vmin   = 2.2;
Vmin1  = 1.5;
Vmin2  = 3;

%Begin
step   = .001;
num    = 2/step + 1;

xZs    = Voc*Rs/(Rs^2 + w^2*L^2);
yZs    = -Voc*w*L/(Rs^2 + w^2*L^2);
rI     = Voc/(2*Rs);
cxI    = rI;
cyI    = 0;
yI     = (cyI-rI):(step*rI):(cyI+rI - step*rI);
xI     = -sqrt(rI^2 - (yI-cyI).^2) + cxI;
xI(num)= cxI;
yI(num)= cyI+rI;

rP     = sqrt((2/Rs)*(Pav - Pmin));
cxP    = rI;
cyP    = 0;
yP     = (cyP-rP):(step*rP):(cyP+rP - step*rP);
xP     = -sqrt(rP^2 - (yP-cyP).^2) + cxP;
xP(num)= cxP;
yP(num)= cyP+rP;

rP1    = sqrt((2/Rs)*(Pav - Pmin1));
cxP1   = rI;
cyP1   = 0;
yP1    = (cyP1-rP1):(step*rP1):(cyP1+rP1 - step*rP1);
xP1    = -sqrt(rP1^2 - (yP1-cyP1).^2) + cxP1;
xP1(num)= cxP1;
yP1(num)= cyP1+rP1;

rP2    = sqrt((2/Rs)*(Pav - Pmin2));
cxP2   = rI;
cyP2   = 0;
yP2    = (cyP2-rP2):(step*rP2):(cyP2+rP2 - step*rP2);

```

```

xP2      = -sqrt(rP2^2 - (yP2-cyP2).^2) + cxP2;
xP2(num)= cxP2;
yP2(num)= cyP2+rP2;

rV       = sqrt(Vmin^2/(Rs^2 + w^2*L^2));
cxV      = Voc*Rs/(Rs^2 + w^2*L^2);
cyV      = -Voc*w*L/(Rs^2 + w^2*L^2);
yV       = (cyV-rV):(step*rV):(cyV+rV - step*rV);
%xV      = -sqrt(rV^2 - (yV-cyV).^2) + cxV;
xV       = sqrt(rV^2 - (yV-cyV).^2) + cxV;
xV(num)= cxV;
yV(num)= cyV+rV;

rV1      = sqrt(Vmin1^2/(Rs^2 + w^2*L^2));
cxV1     = Voc*Rs/(Rs^2 + w^2*L^2);
cyV1     = -Voc*w*L/(Rs^2 + w^2*L^2);
yV1      = (cyV1-rV1):(step*rV1):(cyV1+rV1 - step*rV1);
%xV1     = -sqrt(rV1^2 - (yV1-cyV1).^2) + cxV1;
xV1      = sqrt(rV1^2 - (yV1-cyV1).^2) + cxV1;
xV1(num)= cxV1;
yV1(num)= cyV1+rV1;

rV2      = sqrt(Vmin2^2/(Rs^2 + w^2*L^2));
cxV2     = Voc*Rs/(Rs^2 + w^2*L^2);
cyV2     = -Voc*w*L/(Rs^2 + w^2*L^2);
yV2      = (cyV2-rV2):(step*rV2):(cyV2+rV2 - step*rV2);
%xV2     = -sqrt(rV2^2 - (yV2-cyV2).^2) + cxV2;
xV2      = sqrt(rV2^2 - (yV2-cyV2).^2) + cxV2;
xV2(num)= cxV2;
yV2(num)= cyV2+rV2;

y1 = min(1.1*min(yV2),-1.1*rI);
y2 = max(1.1*max(yV2),1.1*rI);
x1 = min(1.1*min(xV2),-.1*rI);
x2 = max(1.1*max(xV2),2.2*rI);
zz = [0 0];
xz = [x1 x2];
yz = [y1 y2];
rz = [rI rI];
vzx = [cxV cxV];
vzy = [cyV cyV];
P1 = (2/Voc)*[.01 .01];
P2 = (2/Voc)*[.015 .015];
P3 = (2/Voc)*[.02 .02];
P4 = (2/Voc)*[.025 .025];

```

```

hold off
plot(xI,yI)
hold on
plot(xP,yP,'r')
plot(xP1,yP1,'r:')
plot(xP2,yP2,'r:')
plot(xV,yV,'g')
plot(xV1,yV1,'g-.')
plot(xV2,yV2,'g-.')
plot(zz,yz,'k')
plot(xz,zz,'k')
plot(rz,yz,':')
plot(vzx,yz,':')
plot(xz,vzy,':')
plot(P1,yz,'r--')
plot(P2,yz,'r--')
plot(P3,yz,'r--')
title('Geometric View')
axis([-0.05 .1 -.14 .07]);
ylabel('Imaginary');
xlabel('Real');

%Pout = Voc*real(Iopt)
mindifl = 1e-3;
minjl = 1;
minkl = 1;

for j=1:num
    [dif,k] = min(abs(xP-xV(j)) + abs(yP - yV(j)));
    if ((dif < mindifl) && (yV(j) > cyV))
        dif;
        mindifl = dif;
        minjl = j;
        minkl = k;
    end
end

plot(xV(minjl),yV(minjl),'ko')
plot(xZs,yZs,'kx')

Pout = .5*Voc*xV(minjl)
PH = sigma*pi/15 * reye^5 * (w*mu*H)^2

Rtot = Voc/xV(minjl)

```

```

Rlm = Rtot - Rs
Xlm = -w*L
Clr = L/(w^2*L^2 + Rlm^2)
Rlr = L/(Rlm*Clr)

Iload = Vmin/(sqrt(Rlm^2 + Xlm^2))

time=0:.001*8e-6:8e-6;
Voutmag = Voc*sqrt(Rlm^2+Xlm^2)/Rtot;
Voutplot = Voutmag*sin(w*time);
Iloadplot = Voutmag*sin(w*time - atan(Xlm/Rlm));
%figure(2)
%hold off
%plot(time,Voutplot)
%hold on
%plot(time,Iloadplot,'r')
%axis([-0.001 .01 -0.01 .01])

```

Bibliography

- [1] W. F. Agnew, D. B. McCreery, T. G. H. Yuen, and L. A. Bullara. Local anaesthetic block protects against electrically-induced damage in peripheral nerve. *Journal of Biomedical Engineering*, 12(4):301–308, July 1990.
- [2] William F. Agnew and Douglas B. McCreery, editors. *Neural Prostheses: Fundamental Studies*, chapter 2. Biophysics and Bioengineering Series. Prentice Hall, Englewood Cliffs, New Jersey, 1990.
- [3] Mel Bazes. Two novel fully complementary self-biased cmos differential amplifiers. *IEEE Journal of Solid-State Circuits*, 26(2):165–168, February 1991.
- [4] Xenia Beebe and T. L. Rose. Charge injection limits of activated iridium oxide electrodes with 0.2 ms pulses in bicarbonate buffered saline. *The IEEE Transactions on Biomedical Engineering*, 35(6):494–495, June 1988.
- [5] R. C. Black and P. Hannaker. Dissolution of smooth platinum electrodes in biological fluids. *Applied Neurophysiology*, 42(6):366–374, 1979.
- [6] S. B. Brummer and M. J. Turner. Electrical stimulation with Pt electrodes. II. Estimation of maximum surface redox (theoretical non-gassing) limits. *IEEE Transactions on Biomedical Engineering*, 24(5):440–443, September 1977.
- [7] Russell Erich Caulfield. Power limits influencing retinal prosthesis design. S.M. thesis, Massachusetts Institute of Technology, 2001.

- [8] N. Donaldson and T. Perkins. Analysis of resonant coupled coils in the design of radio frequency transcutaneous links. *Medical & Biological Engineering & Computing*, 21:612–627, 1983.
- [9] 2002. From approximate calculations of field-frequency product based on publicly available knowledge of power transmission and implant size.
- [10] Frederick W. Grover. *Inductance Calculations: Working Formulas and Tables*. Dover Publications, Inc., New York, 1962.
- [11] David Halliday, Robert Resnick, and Kenneth Krane. *Physics*, volume 2. John Wiley and Sons, Inc., New York, 4th edition, 1992.
- [12] Mark S. Humayan et al. Visual perception elicited by electrical stimulation of retina in blind humans. *Archives of Ophthalmology*, 114(1):40–46, January 1996.
- [13] Mark S. Humayan et al. Pattern electrical stimulation of the human retina. *Vision Research*, 39(15):2569–2576, July 1999.
- [14] P. F. Johnson and L. L. Hench. An *in vitro* analysis of metal electrodes for use in the neural environment. *Brain Behavior and Evolution*, 14(1):23–45, 1977.
- [15] Shawn K. Kelly. A system for electrical retinal stimulation for human trials. M.Eng. thesis, Massachusetts Institute of Technology, 1998.
- [16] Shawn K. Kelly. Low-power techniques for a retinal prosthesis. Poster at *The Association for Research in Vision and Ophthalmology* conference, May 2003.
- [17] R.J. Kemp, P.N. Murgatroyd, and N.J. Walker. Self resonance in foil inductors. *Electronics Letters*, 11(15):337 – 338, July 1975.
- [18] P.M. Lin and Leon O. Chua. Topological generation and analysis of voltage multiplier circuits. *IEEE Transactions on Circuits and Systems*, CAS-24(10):517–530, October 1977.
- [19] Wentai Liu et al. A neuro-stimulus chip with telemetry unit for retinal prosthetic device. *IEEE Journal of Solid-State Circuits*, 35(10):1487–1497, 2000.

- [20] D. B. McCreery, W. F. Agnew, T. G. H. Yuen, and L. A. Bullara. Comparison of neural damage induced by electrical stimulation with faradaic and capacitor electrodes. *Annals of Biomedical Engineering*, 16(5):463–481, 1988.
- [21] Douglas B. McCreery, William F. Agnew, Ted G. H. Yuen, and Leo Bullara. Charge density and charge per phase as cofactors in neural injury induced by electrical stimulation. *IEEE Transactions on Biomedical Engineering*, 37(10):996–1000, October 1990.
- [22] Alex C. H. MeVay and Rahul Sarpeshkar. Predictive comparators with adaptive control. *IEEE Transactions on Circuits and Systems II, Analog and Digital Signal Processing*, 50(9):579–588, September 2003.
- [23] Urs Oesch and Jiří Janata. Electrochemical study of gold electrodes with anodic oxide films – II. Inhibition of electrochemical redox reactions by monolayers of surface oxides. *Electrochimica Acta*, 28(9):1247–1253, September 1983.
- [24] Kotaro Ogura, Shiro Haruyama, and Kyuya Nagasaki. The electrochemical oxidation and reduction of gold. *Journal of the Electrochemical Society*, 118(4):531–535, April 1971.
- [25] Joseph F. Rizzo III, John Wyatt, John Loewenstein, Shawn Kelly, and Doug Shire. Methods for acute electrical stimulation of retina with microelectrode arrays and measurement of perceptual thresholds in humans. *Investigative Ophthalmology and Visual Science*. Accepted for publication, expected publication date December 2003 or January 2004.
- [26] Joseph F. Rizzo III, John Wyatt, John Loewenstein, Shawn Kelly, and Doug Shire. Perceptual efficacy of electrical stimulation of human retina with a microelectrode array during acute surgical trials. *Investigative Ophthalmology and Visual Science*. Accepted for publication, expected publication date December 2003 or January 2004.

- [27] Kenneth Roach. Irox electrodes... M.Eng. thesis, Massachusetts Institute of Technology, 2003.
- [28] L. S. Robblee, J. L. Lefko, and S. B. Brummer. Activated Ir: An electrode suitable for reversible charge injection in saline solution. *Journal of the Electrochemical Society*, pages 731–733, March 1983.
- [29] L. S. Robblee, J. McHardy, W. F. Agnew, and L. A. Bullara. Electrical stimulation with Pt electrodes. VII. Dissolution of Pt electrodes during electrical stimulation of the cat cerebral cortex. *Journal of Neuroscience Methods*, 9(4):301–308, December 1983.
- [30] T. L. Rose and L. S. Robblee. Electrical stimulation with Pt electrodes. VIII. Electrochemically safe charge injection limits with 0.2 ms pulses. *IEEE Transactions on Biomedical Engineering*, 37(11):1118–1120, November 1990.
- [31] Robert V. Shannon. A model of safe levels for electrical stimulation. *IEEE Transactions on Biomedical Engineering*, 39(4):424–426, April 1992.
- [32] Nathan O. Sokal and Alan D. Sokal. Class e – a new class of high-efficiency tuned single-ended switching power amplifiers. *IEEE Journal of Solid-State Circuits*, SC-10:168–176, June 1975.
- [33] L. "J." Svensson and J. G. Koller. Driving a capacitive load without dissipating fcv^2 . *IEEE Symposium on Low-Power Electronics*, pages 100–101, October 1994.
- [34] P. Troyk and M. Schwan. Closed-loop class e transcutaneous power and data link for microimplants. *IEEE Transactions on Biomedical Engineering*, 39(6):588–599, 1992.
- [35] James Weiland. *Electrochemical Properties of Iridium Oxide Stimulating Electrodes*. PhD thesis, The University of Michigan, 1997.
- [36] James D. Weiland and David J. Anderson. Chronic neural stimulation with thin-film, iridium oxide electrodes. *IEEE Transactions on Biomedical Engineering*, 47(7):911–918, July 2000.

- [37] Robert L. White and Thomas J. Gross. An evaluation of the resistance to electrolysis of metals for use in biostimulation microprobes. *The IEEE Transactions on Biomedical Engineering*, 21(6):487–490, November 1974.
- [38] John Wyatt. 3-d coil calculations. Unpublished notes, October 2000.
- [39] John Wyatt. More efficient charge transfer method ii. Unpublished notes, September 2001.
- [40] John Wyatt. Geometric picture of optimal secondary loading in sinusoidal steady state. Unpublished notes, 2002.
- [41] John Wyatt. Optimal secondary loading in sinusoidal steady state. Unpublished notes, 2002.
- [42] John Wyatt and Joseph Rizzo. Ocular implants for the blind. *IEEE Spectrum*, May 1996.
- [43] Babak Ziaie, Steven C. Rose, Mark D. Nardin, and Khalil Najafi. A self-oscillating detuning-insensitive class-e transmitter for implantable microsystems. *IEEE Transactions of Biomedical Engineering*, 48(3):397–400, March 2001.