Design of a Miniature High-Speed Carbon-Nanotube-Enhanced Ultracapacitor for Electronics Applications

by

Matthew E. D'Asaro

B.S. Electrical Engineering
University of Washington, 2010

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

MASTER OF SCIENCE IN ELECTRICAL ENGINEERING
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2012

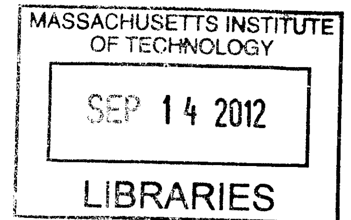

ARCHIVES

© Massachusetts Institute of Technology 2012. All rights reserved.

The author hereby grants to M.I.T. permission to reproduce
and to distribute publicly paper and electronic
copies of this thesis document in whole and in part
in any medium now known or hereafter created.

Signature of Author: ______________________________
Department of Electrical Engineering and Computer Science
May 23, 2012

Certified by: ______________________________
Joel E. Schindall
Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by: ______________________________
Leslie A. Kolodziejski
Chair of the Committee on Graduate Students

Design of a Miniature High-Speed Carbon-Nanotube-Enhanced Ultracapacitor for Electronics Applications

by

Matthew E. D'Asaro

Submitted to the Department of Electrical Engineering and Computer Science
on May 23, 2012, in Partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Electrolytic capacitors, the current standard for high-value capacitors, are one of the most challenging components to miniaturize, accounting for up to 1/3 of the volume in some power devices, and are the weak link with regard to reliability, accounting for the majority of failures in consumer electronics. As a potential alternative vertically aligned carbon nanotubes are utilized to create miniature high-speed ultracapacitors. Because the nanotubes are grown on silicon using low pressure chemical vapor deposition, this technique also opens the possibility of high-value integrated (on-die) capacitors. Using this technique a capacitance density of 52 $\mu F/mm^2$ was achieved. Separately, through careful design of the electrode geometry it is demonstrated that the ionic resistance, the primary factor responsible for the long time constant of ultracapacitors, scales approximately linearly with electrode finger width, thereby demonstrating a workable method for making miniature high-speed ultracapacitors. This work represents the first known example of controlling an ultracapacitor time constant purely though modification of the mechanical structure of the electrodes. It is further projected that using advanced lithography and growth techniques this speed could be increased to 120 Hz. Finally, a variety of packaging techniques are examined for both integrated and discrete applications of this technology.

Thesis Supervisor: Joel E. Schindall, PhD
Title: Bernard M. Gordon Professor of the Practice

Acknowledgments

The completion of this thesis would not have been possible without the assistance, guidance, support, and advice from fellow members of the MIT community, friends, and family. I would like to acknowledge all of those who have helped me including special thanks to:

- My advisor, Professor Joel Schindall, for his continuous support and encouragement.
- Professor John Kassakian for his advice, support and encouragement.
- My friend and colleague, David Jenicek, who worked with and assisted me on a day-to-day basis throughout the project. Without his contributions this project would not have made nearly the progress that it did.
- Sarah Mountjoy, for her long hours helping grow and test the numerous samples required to collect the data in this thesis.
- Dave Otten, for bringing his remarkable engineering skills to bear on so many of the technical problems we encountered during the process of upgrading and using the CNT growth equipment. His temperature controller is a remarkable piece of equipment and has been a tremendous asset to the project.
- Kurt Broderick, for teaching me how to safely and correctly use the facilities at MTL and for providing practical ideas, suggestions, and advice about the deposition process as the project moved along.
- Dr. John Miller for taking the time to visit our laboratory and provide us with crucial knowledge which can otherwise only be obtained though years of experience.
- My grandfather, Arthur D'Asaro, without whose inspiration I would not have made the choices that have led me to where I am today.
- My parents, Danna and Eric D'Asaro and my sister Laura for their continuous moral support, ideas, and occasional assistance with editing my publications.

Contents

Abstract 3

Acknowledgments 5

Prologue 13

1 Background 15
- 1.1 Introduction to the Electrolytic Capacitor . . . 15
- 1.2 Problems with the Electrolytic Capacitor . . . 15
- 1.3 Introduction to the Ultracapacitor . . . 17
- 1.4 Previous Work on High-Speed Ultracapacitors . . . 18
- 1.5 Previous Work on Miniature Ultracapacitors . . . 19
- 1.6 Introduction to the Interdigitated Electrode Miniature Ultracapacitor (IDEMU) Concept . . 19
- 1.7 Thesis Outline . . . 20

2 Catalyst Deposition and Nanotube Growth 23
- 2.1 Overview of the Optimized Deposition Process . . . 23
- 2.2 Challenges to the Deposition Process and Potential Future Work . . . 24
  - 2.2.1 Reducing the Resistance of the Current Collector . . . 24
  - 2.2.2 Reducing Contact Resistance Between the Nanotubes and Current Collector . . . 25
  - 2.2.3 Improving Electrical Connections to the Current Collector . . . 25
- 2.3 Overview of Growth System and Growth Process . . . 26
- 2.4 Challenges to the Growth Process and Potential Future Work . . . 28
  - 2.4.1 Temperature During Reduction . . . 29
  - 2.4.2 Gas Preheating . . . 29
  - 2.4.3 Chamber Pressure . . . 29
  - 2.4.4 Tube Furnace . . . 30
  - 2.4.5 Gas Flow Rates and Growth Temperature . . . 30

3 Electrode Structure Optimization 31
- 3.1 Basic Electrode Structure . . . 31
- 3.2 Calculating electrical resistance . . . 32
- 3.3 Minimizing Finger Spacing . . . 32
- 3.4 Variable Finger Widths . . . 33

4 Electrolyte Selection and Handling 37
- 4.1 Electrolyte Requirements . . . 37
- 4.2 Testing with Aqueous Electrolytes . . . 38
  - 4.2.1 Wetting CNTs in Aqueous Electrolytes . . . 39
  - 4.2.2 Sulfuric Acid . . . 39
  - 4.2.3 Sodium Sulfate . . . 39

4.2.4 Delamination of CNTs in Aqueous Electrolytes . . . 40
4.3 Testing in Non-Aqueous Electrolytes . . . 41
4.3.1 Acetonitrile . . . 41
4.3.2 Propylene Carbonate . . . 42

5 Packaging **45**
5.1 Wire Attachment . . . 45
5.2 Coverslip Encapsulation . . . 45
5.3 Epoxy Encapsulation . . . 47

6 Results and Analysis **49**
6.1 Analysis of Cyclic Voltammograms . . . 49
6.2 Results and Analysis for Wide Finger Devices Optimized for Maximum Capacitance Density 49
6.3 Analysis of Impedance Spectrograms . . . 52
6.3.1 Resistor . . . 52
6.3.2 Capacitor . . . 53
6.3.3 Series Resistor and Capacitor . . . 53
6.3.4 Parallel Resistor and Capacitor . . . 54
6.3.5 Ionic Resistances . . . 54
6.3.6 Equivalent Circuit Model of an Ultracapacitor . . . 54
6.4 Results for High-Speed Narrow Finger Devices . . . 57
6.5 Analysis of Results for High-Speed Narrow Finger Devices . . . 58
6.5.1 Analysis of Relationship Between Finger Width and Ionic Resistance . . . 58
6.5.2 Background on Time Constant . . . 59
6.5.3 Analysis of Time Constant vs. Finger Width Data . . . 60
6.6 Ion Migration Distance . . . 61

7 Conclusion **63**

Bibliography **67**

A Raw Impedance Spectroscopy Data **69**

B Growth Procedure **75**

C Mass Flow Controller System **79**
C.1 Mass Flow Controller Calibration . . . 79
C.2 Documentation for MIT / LEES Mass Flow Controller Interface Box . . . 80

D Heater System **85**
D.1 Growth Chamber . . . 85
D.2 Heater Support and Electrical Interface . . . 85
D.3 Replacing the Heater . . . 86
D.4 Heater Control . . . 87

List of Figures

1.1 Diagram of an electrolytic capacitor . . . 16
1.2 Mainboard of a 2009 Apple laptop charger with electrolytic capacitors highlighted in red . . . 17
1.3 Schematic of an opposing-electrode ultracapacitor using carbon nanotubes as the electrode material. Note that the electrode material is actually much thicker than shown here. . . . 18
1.4 Structure of IDEMU. The nanotubes grow vertically out of the silicon substrate while the ions move sideways, parallel to the silicon substrate. . . . 19

2.1 Partial cross-section of sample after deposition but before growth (not to scale) . . . 24
2.2 Pictures of equipment used in nanotube growth . . . 27
2.3 Diagram of growth chamber (not to scale) . . . 28
2.4 A sample on the heater . . . 28
2.5 Scanning electron microscope image of a shorted sample. Adjacent fingers are from different electrodes and so their bridging causes an electrical short-circuit. . . . 29

3.1 Basic structure of all devices investigated. Units are in microns (µm). . . . 31
3.2 Structure of variable finger spacing devices. This example has ten fingers and 50 µm spacing. Units are in microns (µm) . . . 33
3.3 Structure of the first attempt at variable finger width devices. This example has eighteen fingers per electrode, 50 µm spacing, and 200 µm finger width. Units are in microns (µm) . . . 34
3.4 Structure of the second attempt at variable finger width devices. This example has twenty fingers, 150 µm spacing, and 80 µm finger spacing. Units are in microns (µm) . . . 35

4.1 Solvation shell of a positive ion in water. . . . 38
4.2 Damage to sample tested in sulfuric acid . . . 40
4.3 Delaminated nanotube fingers suspended in solution . . . 40
4.4 Hypothesized mechanism of nanotube delamination . . . 41
4.5 Device wetted in propylene carbonate showing the individual fingers folded sideways and shorting out to each other. The effect looks the same in acetonitrile. . . . 43

5.1 Cross-section of proposed cover slip package. Shown here with through silicon vias (not to scale). . . . 46
5.2 Cross-section of proposed potted package. Shown here without through silicon vias (not to scale). . . . 47
5.3 Experiment in encapsulating a droplet of water with Krylon Appliance Enamel. Note that due to the permeability of the paint, the water has evaporated despite being covered. . . . 48

6.1 Voltage waveform applied to the capacitor by the cyclic voltammeter . . . 50
6.2 Idealized cyclic voltammogram for a capacitor . . . 50
6.3 Cyclic voltammograms for wide finger devices . . . 51
6.4 Idealized impedance spectrogram for a resistor . . . 52
6.5 Idealized impedance spectrogram for a capacitor . . . 53
6.6 Idealized impedance spectrogram for a capacitor in series with a resistor . . . 53

6.7 Idealized impedance spectrogram for a capacitor in parallel with a resistor 55
6.8 Idealized impedance spectrogram for an ionic resistor 55
6.9 Equivalent circuit for an ultracapacitor 55
6.10 Idealized impedance spectrogram for an ultracapacitor. R_{TOT} is the series resistance, the quantity (R_{TOT} - R_L) is the ionic resistance, the quantity (R_L - R_S) is the contact resistance, and R_S is the electrical series resistance. 56
6.11 Example of resistances being extracted from real EIS data. R_{TOT} is marked with an 'X' and R_L is marked with a circle. 57
6.12 Ionic resistance vs. finger width for sample set 1 and sample set 2 59
6.13 Time constant vs. finger width for sample set 1 and sample set 2 60
6.14 A simple model of a parallel-plate carbon nanotube ultracapacitor with 1 cm^2 plate area in the uncharged state (spacing between plates is drawn larger than scale) 61
6.15 A simple model of a parallel-plate carbon nanotube ultracapacitor with 1 cm^2 plate area showing effective ion movement when the device charges (spacing between plates is drawn larger than scale) 62

A.1 Unscaled impedance spectrograms for first set of samples 70
A.2 Impedance spectrograms scaled to show ionic resistance for first set of samples 71
A.3 Unscaled impedance spectrograms for second set of samples 72
A.4 Impedance spectrograms scaled to show ionic resistance for second set of samples 73

C.1 Diagram of the setup required for a bubble test. 80
C.2 Flow rates of the new mass flow controllers plotted against the flow rates registered by the old mass flow controllers. 81
C.3 Reproduction of the pinout label on the back of the MIT / LEES mass flow controller interface box. 82
C.4 Schematic of the MIT / LEES mass flow controller interface box. Note that in the present setup, gas 1 is argon, gas 2 is acetylene, gas 3 is hydrogen, and gas 4 is not connected. 83

D.1 Cross section of the growth chamber without heater support installed. Note that the vacuum fittings that this assembly mounts to are not shown (not to scale). 86
D.2 Cross section of the heater assembly. Note that connecting wires are not shown (not to scale). 87

List of Tables

2.1 Table of gas flow rates and pressures 26

3.1 Table of parameters for variable finger width samples 35

4.1 Radii of various ions in water (from Reed [21]) 38

6.1 Data collected from first set of variable finger width samples 58
6.2 Data collected from second set of variable finger width samples 58

B.1 Gas flow rates for reduction and growth 77

Prologue

Ever since I was a small child I have been fascinated with electronics and electronics repair. One of my earliest memories is sitting transfixed by the vacuum tubes glowing in a radio at my Grandmother's house. When I got a few of years older, I wanted to actually make the radio work, not just watch the tubes glow. With my father's help we quickly found the problem - a dried up and leaking electrolytic capacitor! I fixed the radio and my interest in electronics grew.

Over the years since then I have fixed literally hundreds of pieces of electronics ranging from simple circuits such as 1920's radio sets and fluorescent light ballasts to complex and delicate instruments including microwave signal generators and oscilloscopes. Each project has provided a unique challenge and has helped to stretch my knowledge of electronics, past and present. However, there has been one component that has played a role in more than half of all failures I have diagnosed. It is the electrolytic capacitor.

Because of their high failure rate and their tendency to spray corrosive electrolyte on surrounding parts when they do fail, I have developed a strong dislike for electrolytic capacitors, or 'electrolytics' as they are colloquially known. So when, after having been accepted to MIT as a graduate student, I heard a lecture by Professor Joel Schindall on his carbon nanotube ultracapacitor technology I immediately wanted to try to apply the technology to producing a miniature ultracapacitor capable of competing with the traditional electrolytic capacitor. This thesis is the culmination of that effort.

Chapter 1

Background

This chapter explores the history and current state-of-the-art for electrolytic capacitors and ultracapacitors. It then introduces the concept of a miniature integrated ultracapacitor and outlines the specific challenges that have to be overcome in order to create such a device. The details of these challenges and the work done so far to overcome them is explored in the remaining chapters.

1.1 Introduction to the Electrolytic Capacitor

The present standard for high-value capacitors is the aluminum electrolytic capacitor. Patented in 1901 by Charles Pollak [1], the electrolytic capacitor was the first capacitor with sufficient capacitance density to make a practical off-line power converter (a plug-in power supply) and thus allowed radios to run off of the AC power line instead of batteries. "The first really practical AC set—and sold in quantity, was RCA's Radiola 30, from September, 1925" [2]. Eighty-seven years later the electrolytic capacitor is still the preferred device for filtering the ripple from rectified 60 Hz line voltage to make clean DC, offering sufficient speed and energy density for this application as well as low cost.

The device consists of two aluminum plates separated by an electrolyte. When a potential is first applied to the plates after fabrication a thin layer of alumina (aluminum oxide) forms on the positive anode plate. This oxide is insulating and so electrically isolates the plates, resulting in a capacitor with the anode plate forming one side, the alumina grown on this plate forming the dielectric, and the electrolyte, which is conducting and electrically connected to the other plate, forming the other side. This oxide layer can be extremely thin, and thus the device can have high capacitance density, because it is self-healing - if a hole forms in it the exposed aluminum will quickly react to re-form the insulating alumina layer. See Figure 1.1.

1.2 Problems with the Electrolytic Capacitor

Despite all of its advantages, the aluminum electrolytic capacitor also has numerous drawbacks. The two most prevalent are its low volumetric energy density and its high failure rate. As an example of their low energy density, see Figure 1.2. This figure shows the mainboard of a recent (2009) laptop charger made by the Apple computer company. The electrolytic capacitors, highlighted in red, take up approximately 1.2 in^2

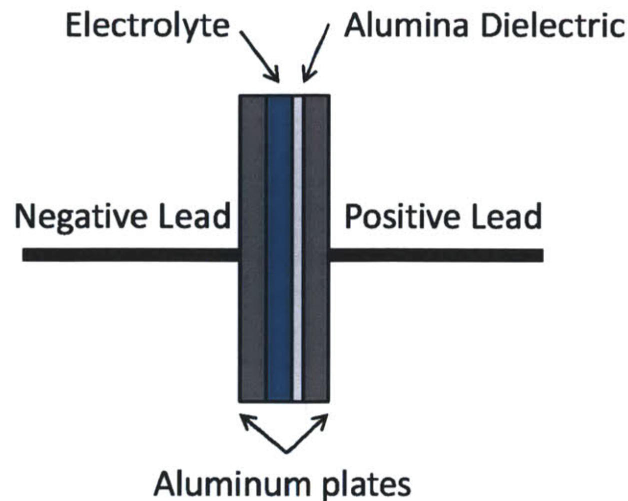

Figure 1.1: Diagram of an electrolytic capacitor

which is 22% of the 5.4 in^2 of the total available board area, second only to the magnetic components. This same percentage applies to device volume, since all components in this design are the same height. In personal experience, this is typical, if not slightly lower than average, with electrolytic capacitors using up to 1/3 of total board area in off-line switching power supplies.

There are two major uses of electrolytic capacitors in off-line switching power supplies. The first is as a primary side filter for smoothing out the rectified line frequency, and the second is as a secondary side filter for smoothing out ripple at the switching frequency. As switching frequencies increase there is reason to believe that the secondary filters can be replaced with smaller value higher speed capacitors such as ceramic multi-layer capacitors. However, since the line frequency is fixed, the capacitance of the primary side filter is dependent only on the current drawn from the converter and the amount of ripple that is tolerable [11]. Thus, this capacitor becomes one of the largest components in the supply.

Furthermore, in extensive personal experience repairing consumer and scientific electronics, about half of all failures have electrolytic capacitors as their root cause. Electrolytic capacitor failure is widely recognized among people working on electronics, to the point that an entire hobbyist website, `http://www.badcaps.net/`, is devoted to their replacement. The mechanism of failure is simple and fundamental to the design. When the capacitor self-heals, aluminum is converted to alumina by the following reaction: $2Al(s) + 3H_2O \rightarrow Al_2O_3 + 3H_2$, thus releasing hydrogen gas that vents through the seal on the capacitor can, thereby using up electrolyte [12]. As the electrolyte is consumed the internal resistance of the capacitor increases, thereby causing heating and additional gas generation. Eventually the device fails open-circuit.

Figure 1.2: Mainboard of a 2009 Apple laptop charger with electrolytic capacitors highlighted in red

1.3 Introduction to the Ultracapacitor

An ultracapacitor, also known as a super capacitor or electrolytic double layer capacitor, offers the highest energy density of any existing capacitor technology. Invented at Standard Oil and patented in 1966 [3], the device operates on an entirely different principle from a traditional electrolytic capacitor. The device consists of two porous electrodes, traditionally made of activated carbon, immersed in an electrolyte and separated by a piece of ion-permeable insulating material. This separator, which is commonly made of cellulose, serves to electrically insulate the electrodes from each other while allowing ions to easily pass through. When a potential is applied between the plates, the ions in the electrolyte are attracted to and migrate into the electrodes with the opposite charge. These ions cancel out the fields created by the electric charges on the plates, allowing additional electric charge to build up, which in turn attracts more ions. The process continues until the all the surface area of the electrodes is covered with ions, implying that the capacitance of the device increases with the porosity of the plates. Figure 1.3 shows such a device; however, in a traditional ultracapacitor the electrodes would be made of a porous sponge-like material rather than the highly ordered carbon-nanotubes as shown in the figure. With highly porous plates, farad-range capacitances are easily obtained. However, because ions have to move inside the device and because ionic resistances are typically much higher than electrical resistances, ultracapacitors have tended to be far too slow for electronic filtering applications [4].

The speed of a traditional ultracapacitor is limited by the ionic resistance of the porous electrodes. That is, as ions migrate into and out of the electrode material under the influence of the electric field created in the electrolyte by the potential across the device, collisions with the electrode material slow their progress. In this sense, the electrode material can be modeled as an ionic resistor with an ionic resistance, R_{ionic} which is a function of the ionic resistivity, ρ_{ionic}, the length (distance the ions have to migrate), L, and a cross-sectional area A per Equation 1.1.

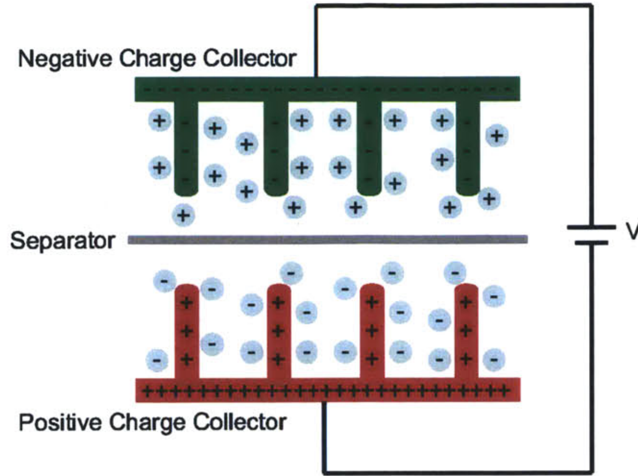

Figure 1.3: Schematic of an opposing-electrode ultracapacitor using carbon nanotubes as the electrode material. Note that the electrode material is actually much thicker than shown here.

$$R_{ionic} = \frac{\rho_{ionic} L}{A} \tag{1.1}$$

1.4 Previous Work on High-Speed Ultracapacitors

Attempts have been made to create ultracapacitors that operate at higher speed than traditional activated carbon devices. One such example is the work done at MIT by Riccardo Signorelli (PhD 2009), Prof. Joel Schindall and Prof. John Kassakian, which forms the basis for the technique explored in this thesis. The idea behind this work was to reduce the ionic resistivity of the electrode material by replacing the traditional activated carbon with carbon nanotubes. Carbon nanotube forests provide less ionic resistance than activated carbon because the paths along which the ions must migrate into the electrode material are less convoluted [5]. See Figure 1.3 for a diagram of this device. Although this technique improves speed by up to 5x as compared with activated carbon based ultracapacitors, by itself this technique does not yield speeds sufficient for filtering at 120 Hz as required to make a primary-side filter for an off-line power converter. For a more detailed description of this device and of ultracapacitor technology in general, see references [17] and [18].

Sufficient speed for that application was achieved by Dr. John Miller by using a forest of approximately 0.6 μm graphene "potato chips" grown out of a nickel substrate and wetted with a phosphonium ionic liquid [4]. While this device provides sufficient speed, the extremely thin layer of electrode material severely limits the capacitance density of the device. Also, since the graphene must be grown at 800 °C it is not possible to integrate this technology into a standard IC.

1.5 Previous Work on Miniature Ultracapacitors

Lower energy commercial ultracapacitors for electronics applications are also available and are sold by companies such as Panasonic, Nichicon, Cornell Dubilier, and AVX. However, these devices exhibit high equivalent series resistance (ESR) and low speed in comparison with traditional electrolytic capacitors, and are thus still primarily sold as burst power or "battery replacement" devices for use in applications like complementary metal oxide semiconductor (CMOS) memory backup.

There are two published microelectromechanical systems (MEMS) ultracapacitors at the time of this writing. The first, which uses the same interdigitated topology that is proposed in this thesis, was done by Y. Q. Jiang [7]. However, no attempt is made in that paper to optimize the speed of the device for filtering applications, nor do they achieve the capacitance density demonstrated in this thesis.

The second, by H. J. In, uses an opposing electrode topology and must be produced using a far more sophisticated manufacturing process where the electrodes are grown in a planar topology and then folded into an opposing topology on micro-fabricated gold hinges [8].

1.6 Introduction to the Interdigitated Electrode Miniature Ultracapacitor (IDEMU) Concept

The idea behind the IDEMU is to utilize the carbon nanotube (CNT) enhanced ultracapacitor technology developed at MIT by Riccardo Signorelli, Prof. Joel Schindall and Prof. John Kassakian, to create miniature planar high-speed ultracapacitors of the structure shown in Figure 1.4.

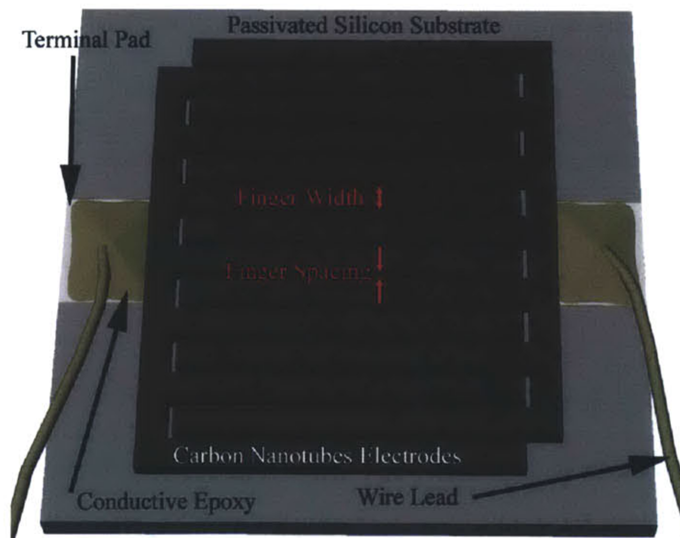

Figure 1.4: Structure of IDEMU. The nanotubes grow vertically out of the silicon substrate while the ions move sideways, parallel to the silicon substrate.

There are several advantages to this structure. First, because the electrodes are fabricated on the same substrate at the same time, no mechanical assembly of the device is required. This means that it can be

miniaturized to scales that would be prohibitive to the assembly of individual parts. Second, because the nanotubes are held rigidly in place by the substrate, and the separation between the electrodes is defined by the lithographic technique used to create the interdigitated pattern, the need for an ion-permeable separator to prevent the electrodes from making electrical contact is completely eliminated, preventing the wasted space and increased ionic resistance that separators introduce. Third, and most importantly, it has the potential to allow a device to be constructed that is simultaneously fast and high capacitance. As the width of each interdigitated finger of the electrodes is decreased, the effective pore depth seen by the ions should also decrease, and with it the ionic resistance of the electrodes. In a device with conventional geometry, the pore depth seen by the ions also decreases with the height of the nanotubes, but here it is linked to the width of the electrode fingers, since the ions are migrating parallel to the silicon substrate, so the RC time constant remains the same. However, in the IDEMU, as the width of the fingers is decreased, the total surface area available to store ions can be maintained by increasing the number of fingers, thus allowing for a high-value, high speed device. In other words, per Equation 1.1 the ionic resistance is dependent on the distance into the electrode that the ions have to migrate, which, in this structure, is set by the finger width, not the nanotube height. Therefore, this is a way to control the time constant of an ultracapacitor by only changing the mechanical structure of the electrodes, something that is not known to have been proposed or done before.

All of the steps required to produce the CNT electrodes can be performed using only standard semiconductor processing equipment, giving the technology the potential to be manufactured inexpensively using existing equipment and to be easily integrated into existing technologies. Furthermore, it has been demonstrated that nanotubes can be grown at CMOS compatible temperatures [6], which means that this technology could potentially be integrated into ICs, allowing previously unobtainable on-die capacitances to be realized. (Trench capacitors are the current state of the art for high-value on-die capacitors but they only demonstrate a capacitance of about 450 nF/mm^2 [10].)

1.7 Thesis Outline

The remainder of this document attempts to answer the question: "Can an ultracapacitor be constructed that competes with electrolytic capacitors and if so, can it be integrated onto an IC?"

Specifically:

- What should the electrodes be made of?

  Chapter 2 explains catalyst deposition and nanotube growth.

- What mechanical structure should be used?

  Chapter 3 describes lithography and device structure optimization.

- What electrolyte should be used?

  Chapter 4 discusses electrolyte selection and wetting.

- How should the device be packaged?

  Chapter 5 proposes methods of packaging.

- What has been achieved?

  Chapter 6 outlines the results obtained.

Chapter 2

Catalyst Deposition and Nanotube Growth

This chapter details the procedure for creating high-speed planar ultracapacitor electrodes on silicon. The challenges associated with this procedure, and with the technique in general, are also explored.

2.1 Overview of the Optimized Deposition Process

Device construction begins with a polished single-crystal silicon wafer that has been thermally oxidized to give it an approximately 50 µm silicon dioxide layer. This wafer represents a passivated IC on top of which a nanotube capacitor is to be grown. Onto this wafer the interdigitated electrode pattern is formed using lithography. The idea is to deposit the current-collector metal and nanotube catalyst only where the electrodes are to be formed, thereby controlling the position and shape of the resulting nanotube forest electrodes with lithographic precision.

The lithographic process uses positive masking. Specifically, photoresist is applied to the wafer, exposed, and developed such that bare silicon is exposed only where nanotubes are to be grown. Positive masking is used, as opposed to negative masking (where the metal would be applied to the whole wafer and then etched off where it is not needed) so as to avoid islands of photoresist, which were determined to be easily damaged when the unwanted metal was removed (liftoff).

The lithography process begins when a quarter of a four inch wafer is cleaned by spinning it at approximately 600 RPM while spraying it first with acetone and then isopropyl alcohol. The wafer is then dried by spinning it at 3000 RPM for 10 to 15 seconds. After cleaning, a coating of OCG 825 photoresist is applied by pouring a one inch diameter puddle of photoresist onto the wafer and then spinning it at 3000 RPM for 30 seconds. The resultant coating is visually inspected for fine particulates, and if there are acceptably few, the photoresist is cured by placing the quarter wafer on a 130 °C hotplate for two minutes.

After the photoresist is applied, the 1/4 wafer is placed on the chuck of a Karl Suss MA4 mask aligner and the mask (a plastic transparency) is placed on top of it and aligned to maximize the number of usable samples. The resist is then exposed for 30 seconds, after which it is developed in OCG 934 for 30 seconds under agitation. After rinsing and drying with high-pressure nitrogen, the 1/4 wafer is ready for the deposition

process.

Next, three layers are deposited on the wafer using an AJA 3-target sputtering machine. These layers consist of: molybdenum (100 nm), alumina (10 nm), and iron (1 nm). The molybdenum underlayer serves as an electrical conductor connecting all of the nanotubes in a given electrode together. Iron is the catalyst for nanotube growth, and the other two layers, alumina and aluminum, have been determined experimentally to promote good nanotube growth [1]. After sputtering, the remaining photoresist is removed with Microstrip® in an ultrasonic cleaner using the following process. A dish of water is heated to just under the boiling point on a hotplate. Into this dish a beaker containing the sample and Microstrip® is placed for roughly one to two minutes. The beaker is then transferred to an ultrasonic cleaner for one to two minutes before being returned to the heated water bath on the hotplate. This process is repeated for approximately 20 minutes until all of the excess metal is removed. The finished wafer is diced into individual 1 cm x 1 cm square devices by one of two processes. Initially the wafers were diced using a die-saw. However, it was determined that the debris produced by this process damaged the catalyst layer, so later samples were diced by manually scribing and cleaving them. These squares (called samples) are then ready for nanotube growth. Figure 2.1 shows a partial cross section of the completed sample. Finger A_N and Finger B_N represent two fingers of opposite electrodes.

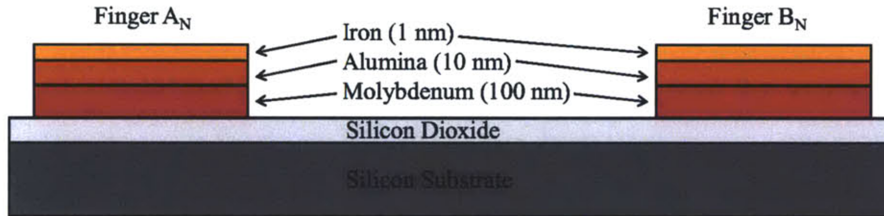

Figure 2.1: Partial cross-section of sample after deposition but before growth (not to scale)

2.2 Challenges to the Deposition Process and Potential Future Work

Since the deposition process relies exclusively on standard semiconductor processing it has the fewest challenges associated with it of all the steps in the construction process. However, the following issues still need to be addressed.

2.2.1 Reducing the Resistance of the Current Collector

The resistance of the molybdenum layer is of concern. As shown in Chapter 3, the resistance of the current collector, even of the widest finger device (400 µm wide fingers) is 1.9 ohms, which is significant in the overall performance of the device. This gets dramatically worse in the narrowest finger (10 µm wide fingers) device at 11 ohms.

The obvious ways to improve this performance are to either use a metal with a higher conductivity for the current collector and/or use a thicker layer of it. First, to address the use of other metals, the conductivity

of Molybdenum is 4.85×10^{-8} Ωm, which is high compared to that of aluminum, the standard material used for metalization in semiconductor processing, at 2.42×10^{-8} Ωm [13]. However, experiments on aluminum have not been successful, presumably because the nanotube growth is currently done at a temperature which exceeds that of the melting point of aluminum. Tungsten, at 4.82×10^{-8} Ωm was also tried earlier in the development process, not because of resistivity issues but because tungsten was the metal used for the current collector in earlier CNT ultracapacitor research [5]. However, tungsten was determined to be completely unsuitable because it does not adhere to the oxidized silicon substrate, resulting in peeling when the photoresist is lifted off. The second choice is a thicker layer of molybdenum. However, increasing the metal thickness greatly increases the liftoff time. In corroboration, increasing the thickness of molybdenum from 50 nm to 100 nm approximately doubled the time required for liftoff to occur completely.

An alternative (but unexplored) deposition process would be to apply the current collector, alumina and iron catalyst to the entire wafer, then apply photoresist, expose the photoresist, develop the photoresist and then etch, removing the metal where it is not needed. This has the advantage of potentially being able to handle thicker current collectors with no liftoff problems but due to the wide array of variables in this project, this approach has been left unexplored. As a final point, the current collector resistivity issue will become much less of a problem as the devices are scaled down, since the width to length ratio of the fingers can be reduced (and with it the electrical resistance of the current collector) without adversely affecting the ionic resistance of the device (i.e. the fingers can still be narrow on an absolute scale).

2.2.2 Reducing Contact Resistance Between the Nanotubes and Current Collector

As will be shown in Chapter 6, the electrical contact resistance between the nanotubes and the current collector dwarfs the more fundamental source of series resistance: the ionic resistance of the electrodes. Thus, it is imperative for a working device that this contact resistance be lowered. The assumed source of this resistance is the alumina layer between the current collector and the catalyst from which the nanotubes grow. This layer appears to be essential for nanotube growth, but alumina is an insulator. Thus, another material that is more conductive than alumina but which still promotes nanotube growth needs to be found, or the thickness of the alumina layer needs to be reduced to the point that its resistance is negligible. To this end, titanium nitride was deposited in place of alumina for one set of samples, but, unfortunately, this did not allow for nanotube growth.

2.2.3 Improving Electrical Connections to the Current Collector

Another, separate issue is that of connecting to the current collector. To ultimately create a useful device, a solderable or wire-bondable conductor will have to be applied to the contact pads. This will require multiple deposition steps and mask alignments that were not required for the kind of experimental fabrication done for this project. However, aside from requiring more expensive chrome masks and a more complex deposition procedure, the steps required are perfectly standard in the semiconductor industry and should not provide significant technical challenge.

As a final note, early in the development of the process an extremely thin (0.3 nm) aluminum layer was placed over the catalyst layer in an attempt to protect the samples from oxidation in air prior to growth.

Table 2.1: Table of gas flow rates and pressures

Gas	Reduction Flow Rate (sccm)	Growth Flow Rate (sccm)
Argon	642	642
Hydrogen	88	66
Acetylene	0	28
Chamber Pressure (Approximate)	10 Torr	20 Torr
Heater Temperature (°C)	100 °C (heater current is off)	780 °C

When no difference was observed in the growth of the samples this was abandoned as it required that the sputtering machine be brought up to air to have one of the targets changed partway through the deposition process, greatly increasing the time required for the deposition process.

2.3 Overview of Growth System and Growth Process

The growth equipment used to produce the vertically aligned carbon nanotubes consists of a custom-built low pressure chemical vapor deposition (CVD) system. This system consists of three parts: a resistive heater to heat the sample, shown in Figure 2.2(a), a vacuum pump to remove the air and gases from the growth chamber, shown in Figure 2.2(b), and a gas distribution system to introduce high purity gas into the growth chamber, shown in Figure 2.2(c). The heater is simply a strip of heavily doped silicon wafer, which is heated by means of a high current power supply. The temperature of this heater is monitored by an infrared thermometer positioned below the growth chamber such that it 'looks' up at the bottom of the heater. A control system monitors the heater temperature and adjusts the current flowing through the heater accordingly, to keep the temperature constant. This heater is contained inside a glass tube made of high purity quartz so that it can withstand the pressure of the atmosphere when under vacuum and the high temperatures generated by the heater. This tube is sealed to the rest of the vacuum system such that the tube can slide to one side, exposing the heater support so that samples can easily be inserted and removed from the chamber. A simplified diagram of the growth chamber is shown in Figure 2.3. See Appendix D for a detailed description of the construction of the growth chamber, heater and heater support.

The growth process is summarized as follows. For the first run of the session, the system is heated by turning on the heater (under evacuation) for about 15 minutes until the tubing feels warm, and all of the gases are purged by running them simultaneously for three minutes. This serves to put the system into the same state that it is in after the completion of a growth run. When preheating was not performed, it was observed that the first growth of the session was different from later growths. After prepping the system, the sample is placed on the heater and the vacuum pump is turned on. Next the argon mass flow controller (MFC) is turned on and set to 642 sccm for three minutes. This is to flush out excess air in the growth chamber. Then the argon is turned off and the pressure is allowed to drop to under 10 mTorr. After this, a valve is closed, constricting the connection between the chamber and the vacuum pump to a piece of 1/8 in. stainless tubing (from a 2 in. connection) before the MFCs are set for the reduction flow rates shown in Table 2.1. The purpose of this step is to reduce unwanted oxides on the surface of the sample, which have been shown to diminish CNT growth [9]. The sample is allowed to reduce for five minutes and then the heater is turned on to the temperature shown in Table 2.1 and the gas flow rates are set to the growth flow rates. Growth continues for 15 minutes after which all the gases are shut off as well as the heater. The

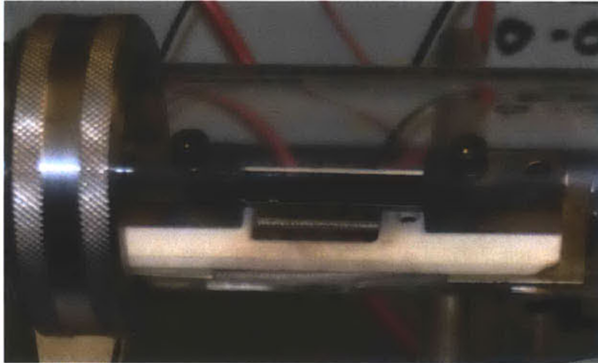

(a) Heater used to heat the sample during growth

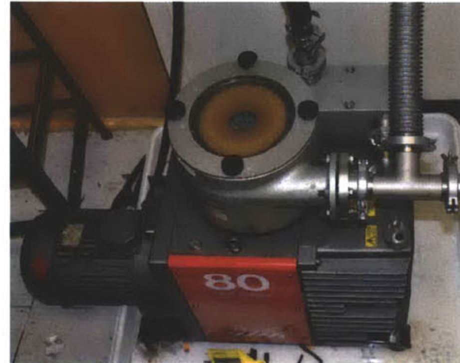

(b) Vacuum pump used to remove air and gasses from the chamber

(c) Gas distribution system used to inject high purity gas into the growth chamber during reduction and growth

Figure 2.2: Pictures of equipment used in nanotube growth

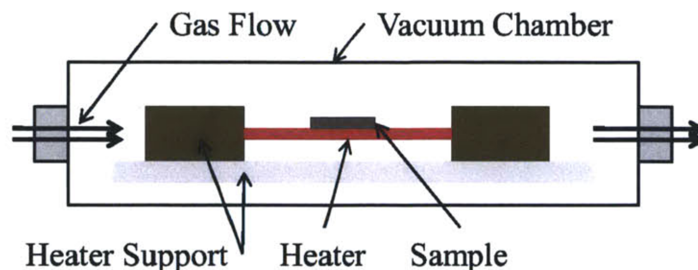

Figure 2.3: Diagram of growth chamber (not to scale)

restriction valve is again opened and the system is allowed to evacuate again for five minutes. This serves the safety functions of ridding the chamber of residual gases and allowing the sample to cool before it is exposed to air. The height of the resulting nanotubes is on the order of 100 µm but varies with the specific sample. Figure 2.4 shows what the sample looks like when it is sitting on the heater. Note the dull red color of the heater.

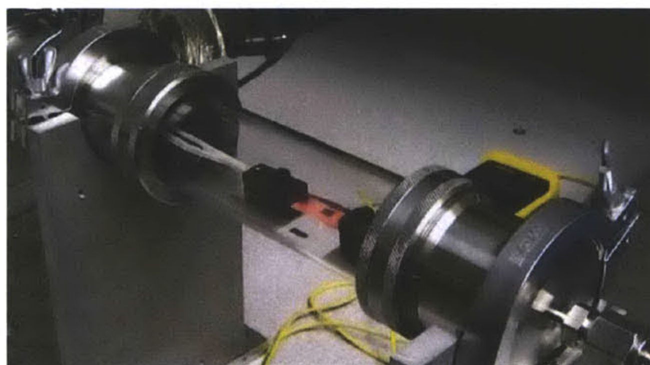

Figure 2.4: A sample on the heater

2.4 Challenges to the Growth Process and Potential Future Work

The growth process is by far the most temperamental part of the construction process. Frequently the samples do not grow in spots or, worse, the nanotubes grow sideways and short out the sample as in Figure 2.5. This section outlines in detail these problems, the work done so far to address them, and what can be done in the future.

Not all of the factors influencing nanotube growth are understood as yet. However, it is clear that aside from the deposition of the catalyst, the main factors affecting the results are the gas ratios during the reduction and growth, and the heater temperature during growth. Extensive experiments have been performed to attempt to optimize these parameters and eliminate other variables. A few examples are outlined below.

2.4.1 Temperature During Reduction

Heater temperature during reduction was explored in two different manners. First, the heater was heated to a known temperature and then turned off, at which point a timer was started. After fixed amounts of time, varying between a few seconds and several minutes, which permitted the heater to cool to varying degrees, the sample was placed onto the heater and the growth process was initiated. No association between wait time and the resulting growth quality was observed. Second, the heater was explicitly turned on during the reduction phase. This resulted in contradictory results as the first sample grown in this manner seemed to grow better but later ones seemed worse, and so this approach was abandoned.

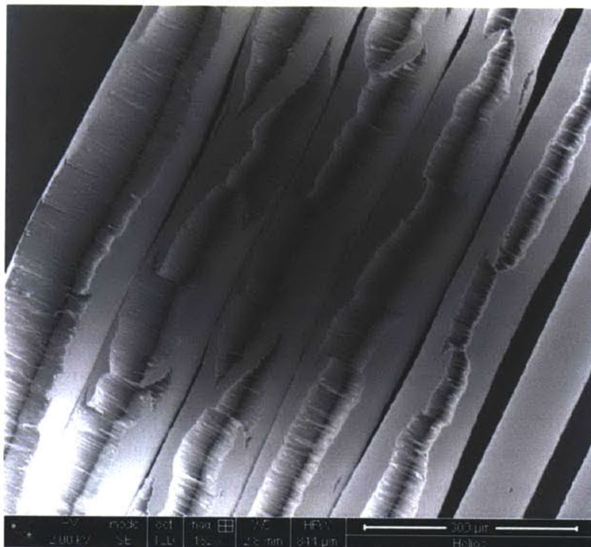

Figure 2.5: Scanning electron microscope image of a shorted sample. Adjacent fingers are from different electrodes and so their bridging causes an electrical short-circuit.

2.4.2 Gas Preheating

Another factor that was explored was the temperature of the gas entering the chamber. This was controlled by placing a tube furnace around the line entering the growth chamber and setting it to various temperatures around those used for growth (700 °C to 800 °C). It has been reported [9] (although with ethylene as opposed to acetylene) that heating the gas before allowing it to enter the chamber changes its chemical composition in a manner that favors nanotube growth. However, no significant improvement was observed. According to Zhong, [19], the precursor gas needed for CNT growth is acetylene, which is used in this project. Others have used ethylene, which forms acetylene at high temperatures, explaining their success with preheating and the results obtained here which indicate that it is not useful.

2.4.3 Chamber Pressure

Chamber pressure is not actively controlled in our system because the pressure is set by the constriction size of the tubing, the characteristics of the vacuum pump, and the rate at which gases are introduced into

the chamber. Thus, to determine the effect of pressure on growth rate, the constriction during growth was increased by partially closing a valve, but no noticeable increase in growth was observed. In a future system it would be highly desirable to have a gate valve in series with the vacuum pump, which would be controlled electronically to maintain a constant pressure. This would eliminate pressure as a variable and would allow independent control of gas pressure and gas flow rate through the chamber.

2.4.4 Tube Furnace

Other nanotube groups [9] have experimented with a tube furnace instead of a heater to provide the high temperatures needed for nanotube growth. The advantage is that this guarantees even heating of the whole sample. However, it was found to be impracticable for the reason that the tube furnace has such large thermal mass that if a sample is placed into a hot furnace it immediately and completely oxidizes, rendering it useless. Also, once the furnace is hot it takes hours to cool, which prevents the rapid cycling of samples needed for efficient experimentation. That said, it might be feasible to use a tube furnace if a sample insertion system were built that allowed the sample to be inserted into the furnace while the furnace is under vacuum.

2.4.5 Gas Flow Rates and Growth Temperature

Finally, and most importantly, the optimal ratio of gases and the optimal growth temperature were determined through repeated experiments at different flow rates and temperatures. The flow rates appear to be much less critical than the temperature, and the values in Table 2.1 were readily determined and have been maintained throughout most of the experimentation. However, growth appears to be very sensitive to the heater temperature, and the optimal temperature seems to change somewhat from batch to batch. For that reason, temperatures ranging from 675 °C to 780 °C have been used, with the highest temperatures yielding the best results in the latest experiments. If the temperature is set too low the growth will be short and tends to have an excess of amorphous carbon in it. On the other hand, a temperature that is too high tends to produce uneven growth or even the complete absence of growth if the temperature is excessive. This is due to the trade-off between activation energy needed for growth, favoring a high temperature, and the self-pyrolysis (chemical breakdown) of acetylene if the temperature is excessive.

These are all of the major challenges to growth that have been recognized so far. Although they are significant, using the existing system produces good quality growth a high percentage of the time and is certainly sufficient to produce growths for experimental and even prototyping purposes. The next chapter discusses the design of the mechanical structure of the electrodes and its effect of the capacitance and speed of the resulting devices.

Chapter 3

Electrode Structure Optimization

This chapter addresses the geometric configuration of the interdigitated electrode structure and its effect on time constant and capacitance. The two basic parameters are the width of the fingers and the width of the space between the fingers. Thinner fingers allow for faster speed, but also increase the electrical resistance of the current collector and decrease the total capacitance. Wider space between fingers increases reliability but at the expense of lower capacitance.

3.1 Basic Electrode Structure

The basic structure, common to all devices investigated, is shown in the scaled drawing shown in Figure 3.1.

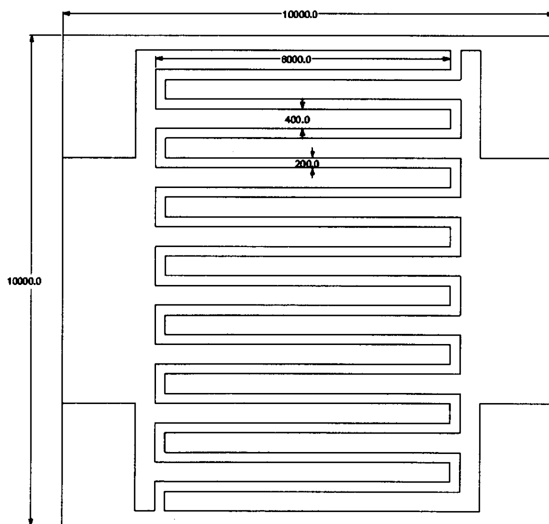

Figure 3.1: Basic structure of all devices investigated. Units are in microns (μm).

As shown, the structure consists of interdigitated fingers of uniform width separated by uniform spacing. On either side, all of the fingers are attached together and to a pad for making electrical connections. In

this original design the spacing between the fingers is 200 µm while the finger width is 400 µm and there are eight fingers per electrode. The entire device is 1 cm x 1 cm as this is the maximum size that will fit on the heater in the growth chamber (see Chapter 2).

3.2 Calculating electrical resistance

One of the important parameters for a given mechanical structure is its effective electrical series resistance. This is the effective electrical resistance seen in series with the capacitor considering only the effect of the resistivity of the current collector material deposited in the interdigitated pattern. Calculating this resistance is straightforward as each finger can be modeled as a resistor made of a thin strip of resistive material. Equation 3.1 calculates the resistance of such a strip of material with thickness t, width W, and resistivity ρ.

$$R = \rho \frac{L}{Wt} \tag{3.1}$$

Since the capacitance is distributed evenly along each finger the average resistance seen in series with this capacitance is half the resistance of a finger, but since charge must travel through two fingers (one in each electrode) to charge or discharge the capacitor the total resistance seen is equal to the resistance of a finger. Finally, since there are n fingers per side, all in parallel with each other, the total resistance is divided by the number of fingers per electrode. Therefore, Equation 3.2 represents the total calculated electrical resistance of the fingered structure. Note that this calculation neglects resistance in the bonding pad and interconnect wires, but this should not be significant, especially in later designs when the width of the bonding pad was dramatically increased.

$$R = \rho \frac{L}{nWt} \tag{3.2}$$

For example, if the structure shown in Figure 3.1 is deposited with 20 µm of molybdenum, as was the case for the original samples, then the electrical ESR is calculated to be 4.5 Ω.

3.3 Minimizing Finger Spacing

After verifying the basic functionality of the structure shown in Figure 3.1, the next step was to determine the minimum spacing between the fingers. Minimizing the space between fingers maximizes the capacitance, as it increases the device area used for electrode material. However, it has two other advantages as well. First, from a packaging perspective, if the finger spacing is sufficiently narrow, surface tension alone will hold electrolyte between the fingers, meaning that the excess electrolyte can be removed from the sample, leaving a well defined structure of wetted nanotubes. This has the potential to simplify packaging as it would eliminate loose liquid without having to resort to gelling the electrolyte (see Chapter 5). Second, by minimizing finger spacing the ionic resistance of the electrolyte between the fingers is minimized, potentially increasing device speed.

To find a minimum finger spacing, a set of four different types of samples were made. Each had a constant finger width of 400 µm, just like the original samples, but the finger spacings were 200 µm (the

original spacing), 100 µm, 50 µm, and 20 µm. To approximately maintain a 1 cm^2 sample size, the number of fingers per electrode was increased to nine, ten, and eleven respectively for the three new spacing sizes. For example, Figure 3.2 is a scale drawing of the ten finger 50 µm spacing sample.

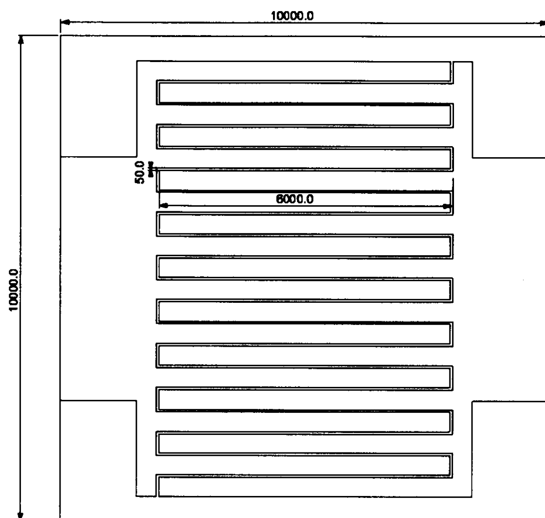

Figure 3.2: Structure of variable finger spacing devices. This example has ten fingers and 50 µm spacing. Units are in microns (µm)

These new samples were then grown and the yield rate was noted, thus determining that the yield dropped below 50% for the finest (20 µm) spacing and thereby setting the minimum spacing, at least for the equipment used, to be approximately 50 µm. Yield here refers to the percentage of samples that are not shorted after being grown.

3.4 Variable Finger Widths

The final, most critical, and ultimately most challenging part of electrode structure optimization was determining the effect of finger width. As explained in Chapter 1 the ionic resistance should decrease linearly with decreasing finger width, and confirming this was one of the major goals of this project.

First attempt at variable finger widths

Since the minimum finger spacing had already been determined to be 50 µm, the first set of variable finger width samples were made using that same spacing. The finger widths were 400 µm, 200 µm, 80 µm, and 40 µm. Again, the number of fingers was increased to maintain approximately 1 cm^2 total sample area. Figure 3.3 is a scaled drawing of the 200 µm wide samples.

Unfortunately, these samples did not work correctly for a variety of reasons. First and foremost, although the 50 µm spacing had been tried for a sample with relatively few fingers and produced acceptable yield, when the number of fingers was dramatically increased (the 40 µm wide finger samples had 90 fingers, for example) the yield decreased to practically zero. The majority of devices failed because one or more bunches

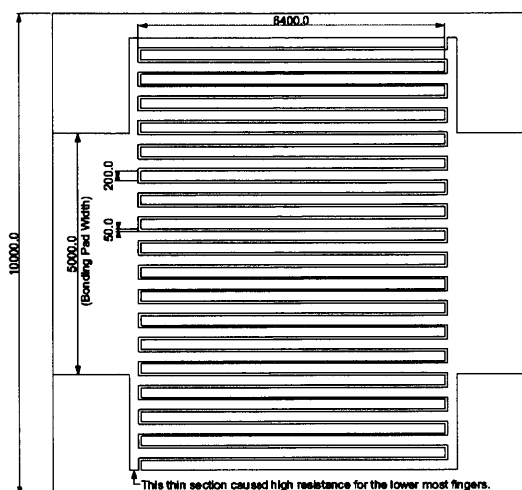

Figure 3.3: Structure of the first attempt at variable finger width devices. This example has eighteen fingers per electrode, 50 µm spacing, and 200 µm finger width. Units are in microns (µm)

of CNTs grew laterally, shorting adjacent fingers. The probability of this occurring scales with the length of the boundary between the sets of fingers and thus with the number of fingers. The second problem was that in this first design, the width of the material connecting the uppermost and lowermost fingers to the bonding pad scaled with the finger width, causing greatly increased (and therefore non-negligible) electrical resistance in these fingers. See Figure 3.2. Finally, the widest finger sizes turned out to be so wide that the nanotube fingers were wider than the nanotubes were long. Such wide fingers are undesirable because in that case the shortest path for the ion flow to take would be to come vertically out of one finger and vertically down into the adjacent finger in an arching path. This would invalidate the linear ionic resistance scaling with finger width theory because if the nanotubes are shorter than the fingers, it would be the nanotube height, not the finger width that would set the length of the dominant ionic resistor and thus the ionic resistance. For all of these reasons, this first design of variable finger width samples was abandoned.

Second attempt at variable finger widths

To solve the problems with the first generation of variable finger width samples, a second set was designed. This set used 150 µm finger spacing with finger widths of 10 µm, 20 µm, 40 µm, and 80 µm and with redesigned bonding pads. The 80 µm sample is shown in Figure 3.4. This redesign solved all of the problems previously discussed and successfully demonstrated a roughly linear relationship between finger width and ionic resistance (see Chapter 6). The electrical resistance of each of these samples, as computed by Equation 3.2, is summarized in Table 3.1.

As can be seen, the electrical resistances are still relatively large, but this could be completely mitigated by using a thicker metalization or more fingers (assuming that the shorting problem could be resolved through better growth equipment).

This concludes the analysis of the mechanical structure of the electrodes. See Chapter 6 for the excellent

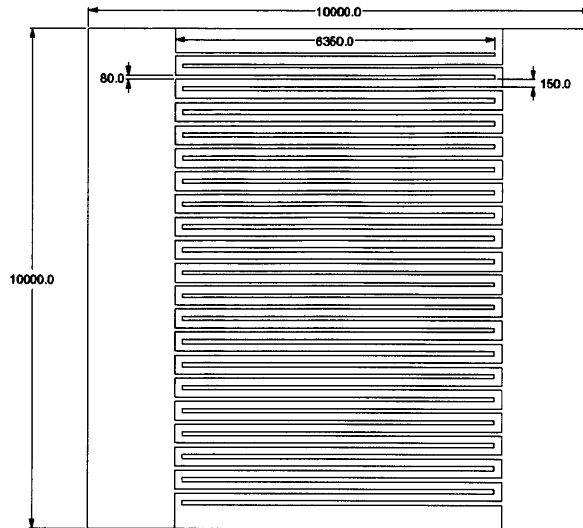

Figure 3.4: Structure of the second attempt at variable finger width devices. This example has twenty fingers, 150 µm spacing, and 80 µm finger spacing. Units are in microns (µm)

Table 3.1: Table of parameters for variable finger width samples

Finger Width (µm)	Finger Spacing (µm)	Finger Length (mm)	Number of Fingers	ESR (Ω)
10	150	6.35	28	11
20	150	6.35	27	5.7
40	150	6.35	24	3.2
80	150	6.36	20	1.9

results obtained from this last design. The next chapter discusses electrolyte selection and handling, and its effect on device performance.

Chapter 4

Electrolyte Selection and Handling

In order for an ultracapacitor to function, its electrodes must be submerged in an electrolyte (see Section 1.3). However, numerous mechanical and chemical issues occur when the electrolyte is added to the system. This section discusses those challenges and the various electrolytes that have been tested in an attempt to overcome them.

4.1 Electrolyte Requirements

The most fundamental requirement of the electrolyte is that it contain mobile ions. This means that (aside from exotic materials such as ionic liquids) the electrolyte must consist of a solvent with an ionic compound (salt) dissolved in it. At this level, common saltwater would work as a (crude) electrolyte. However, there are other requirements as well. It is desirable that the solution have a low ionic resistance. That is, the ions should move easily through the solution when a potential (electric field) is applied across it. This is necessary for a high speed device. Furthermore, to decrease the ionic resistance of the device and to increase the available surface area, and therefore the capacitance, it is desirable that the chosen compound produce ions with a small diameter. Since the ions are dissolved in the solvent, the radius that matters is that of the solvation shell, and this is affected both by the selection of the salt and the solvent. See Figure 4.1 and Table 4.1. (There is some controversial evidence that under certain circumstances an ion can be separated from its solvation shell in an ultracapacitor. This effect, if it is proven to exist, is beyond the scope of this document. [20]) To complicate matters, the voltage limit of an ultracapacitor is set by the voltage at which the electrolyte breaks down by electrolysis. For example, at STP (standard temperature and pressure) this is about 1.2 V for water [13]. Per equation 4.1, the energy stored in a capacitor increases as the square of the voltage across it, so it is highly desirable to increase the operating voltage of the device, indicating that a solvent other than water should be used. On top of these basic requirements, the electrolyte must not chemically interact with the other materials used in the cell (capacitor) and needs to be stable both over time and over the operating and storage temperature ranges of the device. Each of these electrolyte parameters introduces various trade-offs, with the added complexity that the higher-performance options are more difficult to handle. Put differently, there is no one perfect electrolyte, so the various options must be considered carefully. The rest of this chapter is devoted to the choices of electrolytes thus far explored.

Table 4.1: Radii of various ions in water (from Reed [21])

Ion	Dry Radius Å	Hydration (Mol H_2O)	Hydrated Radius (Å)
Li^+	0.78	14	7.3
Na^+	0.98	10	5.6
K^+	1.33	6	3.8
Rb^+	1.49	0.5	3.6
NH_4^+	1.43	3	-
Mg^{2+}	0.78	22	10.8
Ca^{2+}	1.06	20	9.6
Ba^{2+}	1.43	19	8.8
Al^{3+}	0.57	57	-

$$E_{capacitor} = \frac{1}{2}CV^2 \tag{4.1}$$

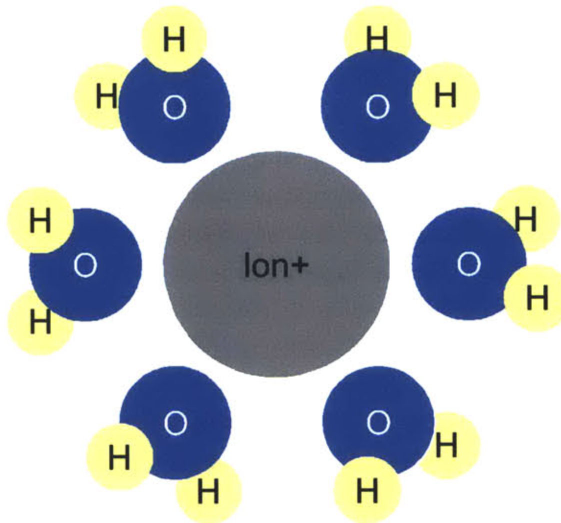

Figure 4.1: Solvation shell of a positive ion in water.

4.2 Testing with Aqueous Electrolytes

The majority of experiments performed were done on aqueous (water based) electrolytes because of their relative ease of handling compared to organic electrolytes. However, this limits the usable maximum voltage for the cells to 1 V and eventually causes the nanotubes to delaminate from the current collector.

4.2.1 Wetting CNTs in Aqueous Electrolytes

Carbon nanotubes are highly hydrophobic (contact angle of approximately 145°, per Murakami [14]), meaning that water (and water based solutions) will not penetrate down into a forest of them but rather will bead up on top in almost spherical droplets. This is highly undesirable as the electrolyte must penetrate into the electrodes in order for the device to function. This problem was overcome by the following procedure, kindly suggested by John Miller of JME Consulting.

- Using a pipette apply one droplet of water to the surface of the CNT electrode.

  The droplet beads up and stays on top.

- On top of this drop apply one drop of concentrated isopropyl alcohol (IPA).

  Once the IPA is introduced, both droplets combine and penetrate into the nanotube forest, releasing the trapped air in it in the form of bubbles.

- Wash out the water and IPA by liberally applying electrolyte.

  It is undesirable that the electrolyte be diluted with or contaminated with the alcohol.

In a revised version of this procedure, the alcohol is applied directly to the dry sample and then washed out with the electrolyte, thereby simplifying the procedure. The advantage of the original procedure, however, is that it clearly demonstrates that the IPA allows the water to penetrate the normally hydrophobic electrode.

4.2.2 Sulfuric Acid

Sulfuric acid was the first electrolyte experimented with. It is a good choice because of its low ionic resistance and small ion size, and is widely used in the ultracapacitor industry. However, it was quickly determined to be unsuitable for experimentation due to its highly corrosive nature. The acid rapidly dissolved the silver plating from the wires connected to the ultracapacitor and, due to bubbling, it was also apparent that it was dissolving the silver in the epoxy used to hold the wires onto the bonding pads on the edges of the sample (see Chapter 5). Figure 4.2 shows the sample in question. Note the exposed copper on the right-hand interconnect wire and the discoloration of the nanotubes (assumed to be due to dissolved silver deposition) and the darkening of the silver epoxy. Further experiments with sulfuric acid could proceed, but doing so would require that all metals used in the device be compatible. This could be achieved by using gold-wire bonding to attach to the device, but that would add considerably to the expense and complexity of testing, and so was not pursued.

4.2.3 Sodium Sulfate

To allow testing with an easy-to-handle aqueous electrolyte, but not have the sample damaged by a strong acid, it was proposed by David Jenicek (PhD candidate, MIT) that a neutral salt be used, thereby providing ions without lowering the pH of the system. A review of literature showed that sodium sulfate was sometimes used for this purpose, due to its relatively small ion size, low ionic resistance, safety, availability, and compatibility with other materials used in the cell [15]. This proved a good choice and all remaining testing in aqueous electrolyte was done with sodium sulfate. The solution proved easy to work with, and using the

Figure 4.2: Damage to sample tested in sulfuric acid

process described previously, was easily wetted to the samples. However, since it is aqueous, it is only useful for testing below 1 V (1.2 V is the absolute limit, but an 0.2 V safety window is usually needed), and as described in the next section, it caused the devices to disintegrate within several hours of wetting.

4.2.4 Delamination of CNTs in Aqueous Electrolytes

From the first time sodium sulfate was used, it was observed that the carbon nanotubes delaminate from the current collector after the sample has been wetted for several hours. This first manifests itself as a drastic decrease in capacitance, followed by the nanotube fingers separating completely from the substrate. Interestingly, the nanotubes stick together, such that the individual fingers remain intact, although separated from the substrate and suspended in solution as shown in Figure 4.3. It was further observed that this process is accelerated by the presence of an electric potential across the device. If a bias is applied, the device can start to disintegrate in under an hour. To see if the problem is due to sodium sulfate specifically, or to all aqueous electrolytes, a sample was placed in pure deionized water for approximately 48 hours and was found to delaminate just like those placed in sodium sulfate.

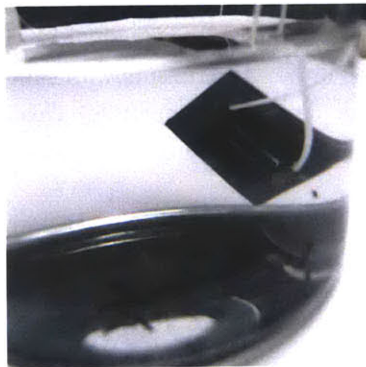

Figure 4.3: Delaminated nanotube fingers suspended in solution

The delamination process, therefore, occurs due to the presence of water and although the exact mechanism has not been determined with certainty, it is hypothesized that, as described by Murakami and

Maruyama, it is due to the highly hydrophobic nature of carbon nanotubes [14]. When water is forced into the hydrophobic nanotube forest, it creates microscopic forces on the nanotubes as the water tries to wet to the substrate but is repelled from the nanotube, as shown in Figure 4.4. This force separates the nanotube from the substrate, eventually delaminating the entire electrode.

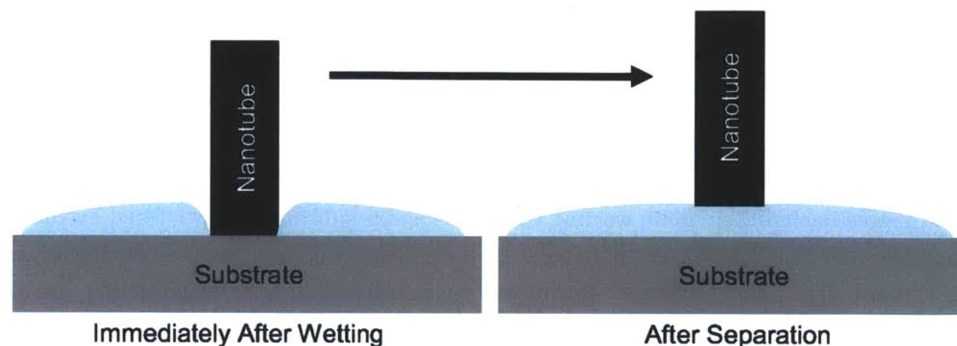

Figure 4.4: Hypothesized mechanism of nanotube delamination

4.3 Testing in Non-Aqueous Electrolytes

Due to the problems associated with aqueous electrolytes, non-aqueous electrolytes are frequently used in commercial and experimental ultracapacitors. The two commonly used solvents for non-aqueous electrolytes are acetonitrile and propylene carbonate. Both of these solvents allow the cells to be used up to 2.7 V. Acetonitrile has a higher ionic conductivity, but also has a much higher vapor pressure (boiling point 82 °C), evaporating quickly if not kept in a sealed container. It also has the disadvantage of being moderately toxic if inhaled or ingested and decomposing into cyanide gas if overheated such as in an over-voltage condition in an ultracapacitor. It is widely used in the US, but due to safety concerns it is banned for use in ultracapacitors in Japan. Propylene carbonate, on the other hand, is very safe (so much so that it is frequently used in cosmetics) and has a much lower vapor pressure (boiling point 240 °C). However, the non-aqueous electrolytes present another challenge in that they must be handled in an environment free from moisture. They are hydrophilic and if water is introduced will decompose by electrolysis when the cell is brought above 1 V, resulting in highly reactive hydrogen and oxygen being produced that will then damage the cell. In both cases, the salt commonly used in non-aqueous electrolytes is TEMA/BF4 (triethylmethylammonium tetrafluoroborate), desirable because of its high solubility and small ion size.

4.3.1 Acetonitrile

A 1 M solution of TEA/BF4 (tetraethylammonium tetrafluoroborate) in acetonitrile was kindly provided by Riccardo Signorelli (PhD, 2009) and thus testing in acetonitrile was done with TEA/BR4 rather than TEMA/BF4. These salts are extremely similar and the high cost of purchasing a small quantity of TEMA/BF4 meant that the TEA/BF4 solution was used instead.

Due to the nonpolar nature of the acetonitrile, there is no problem wetting it to the carbon nanotubes -

applying a drop causes it to immediately sink in rather than bead up as is the case with aqueous electrolytes. Furthermore, after leaving a sample sitting in a solution of acetonitrile for over two years, no evidence of nanotube separation could be seen. In fact, in the course of all the experiments done with acetonitrile throughout the duration of this project and all those done previously by Signorelli, no issues with nanotube separation were observed.

Samples Shorting in Acetonitrile

Although no delamination was observed, another equally serious problem emerged. An unknown mechanism causes the very thin fingers of nanotubes to fall sideways against each other, shorting out the device. This occurs to some degree immediately upon wetting the device and is substantially worsened by the application of an electric potential. For this reason, no usable data was obtained from samples wetted in acetonitrile. Investigating this problem proved difficult. No matter which electrolyte is used, after they are wetted it is impossible to dry the samples without destroying them, as the evaporating liquid causes the nanotubes to bunch up and fall over, completely destroying the fingered structure. Since acetonitrile evaporates extremely quickly, it was not easily possible to photograph the damaged samples under optical or electron microscopes.

Thus, there are several unconfirmed hypotheses as to why the samples short in the acetonitrile. The first is that the nanotubes do not fully wet in the aqueous electrolyte, and the resulting surface tension has the effect of drawing the individual tubes together, forming a relatively rigid mechanical structure resistant to being pulled over. This is evidenced by the fact that the nanotube fingers, when wetted in water, appear to be thinner at the top than at the bottom. The second hypothesis is that the acetonitrile is acting as a lubricant, allowing the nanotubes to slide relative to one another, thereby allowing the entire finger structure to fold over sideways. These two hypotheses are not mutually exclusive and some combination of them may be occurring.

4.3.2 Propylene Carbonate

Due to the high cost of dry TEMA/BF4 salt, it was not possible for us to prepare a solution of TEMA/BF4 in propylene carbonate. However, a sample of pure propylene carbonate was acquired and used to attempt to create a solution with sodium sulfate. This proved unsuccessful as sodium sulfate is apparently insoluble in propylene carbonate. That said, it was determined that the carbon nanotubes easily wet in the propylene carbonate, just as they do in the acetonitrile. Unfortunately, the propylene carbonate is no better than the acetonitrile with regard to samples shorting - the nanotube fingers fold sideways and short in propylene carbonate just as they do in acetonitrile. However, since propylene carbonate is not nearly as volatile as acetonitrile, it was possible to get pictures of the damage under an optical microscope. Figure 4.5 shows the effect.

Due to the problem with the non-aqueous electrolytes causing the samples to short, non-aqueous electrolytes were abandoned for the remainder of the project, and all useful data (see Chapter 6) was taken using sodium sulfate. The data was taken immediately after wetting the samples to avoid delamination. That said, the non-aqueous electrolytes, because they don't cause said delamination and allow for operation at a higher voltage, are probably a better choice for a commercial device. The use of a gelling agent to thicken the electrolyte and thereby hold the fingers in place without adversely effecting ion mobility is a

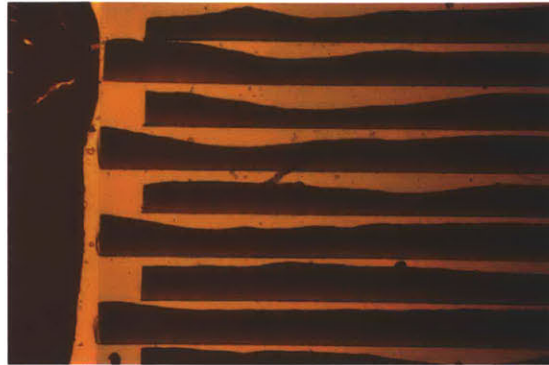

Figure 4.5: Device wetted in propylene carbonate showing the individual fingers folded sideways and shorting out to each other. The effect looks the same in acetonitrile.

promising option for controlling the shorting issue. This concludes discussion of electrolytes. In the next chapter, various methods of packaging the device are explored.

Chapter 5

Packaging

One of the major remaining hurdles to a working miniature ultracapacitor is a reliable, small, inexpensive, and easy-to-manufacture package. Unlike traditional semiconductor devices, ultracapacitors require a wet electrolyte to operate, thus necessitating a hermetically sealed package capable of preventing the electrolyte from evaporating or becoming contaminated with environmental moisture. The fact that the electrolytes of choice cannot be heated too much above the boiling point of water and are easily contaminated by common metal ions such as iron and silver eliminates common hermetic techniques based on soldering, brazing, or high-temperature fusing of glasses and ceramics. This leaves adhesive-based or mechanical bonding as the only reasonable choices. Due to lack of time and resources no serious experiments on packaging were performed; however, two potential packaging options, both derived from existing semiconductor packaging techniques, were considered and are described in detail below.

5.1 Wire Attachment

Before beginning to describe future systems, it is worth discussing the 'packaging' performed on the experimental devices created so far. Because the devices are simply placed in a beaker of electrolyte for testing, there is no need to have a proper 'package' per se. However, it is still necessary to attach wires to the contact pads on the samples so that electrical testing can be performed. Wire bonding and soldering were investigated, but both would require an additional deposition layer with a separate, aligned, mask. Also, soldering has the potential to introduce unwanted ions while the diameter of wire bonding wire is so small that it makes handling difficult. For these reasons silver-plated Teflon coated wire was chosen to be attached to the samples with silver epoxy. The use of silver epoxy was taken from the work done by Jiang et al. in [7].

5.2 Coverslip Encapsulation

One method considered was to encapsulate the device by placing a piece of impermeable material on top of the nanotube electrodes and sealing it around the edges with adhesive, presumably an epoxy resin. This design is shown in schematic form in Figure 5.1.

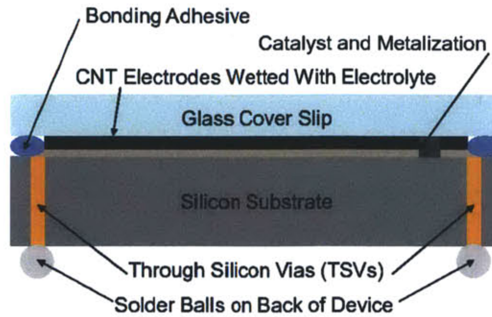

Figure 5.1: Cross-section of proposed cover slip package. Shown here with through silicon vias (not to scale).

This design derives from a procedure already used for encapsulating image sensors used in digital cameras with the exception that in that application the space between the cover slip and the device is not filled with a liquid electrolyte [16]. One advantage of this option is that the technology to create this kind of package already exists, and therefore should be less expensive and easier to manufacture than a packaging technology invented for ultracapacitors. It also has the advantage that glass is one of the few truly hermetic materials and is completely compatible with all potential electrolytes. This leaves the sealing material as the only potentially permeable part of the package, and even this can be mitigated by making the seal very wide and short, thereby greatly increasing the distance over which electrolyte and/or contaminants would have to diffuse. The package could be designed with through silicon vias (TSVs) such that the finished device could be bumped on the back and directly soldered to the board, as shown in Figure 5.1, or, as a less expensive alternative, it could be configured with bonding pads as used with the alternative packaging scheme shown in Figure 5.2. That said, the TSVs offer the advantage of allowing the sealing compound to mate directly with the silicon substrate, simplifying the adhesion requirements of the adhesive. TSVs would also allow for a smaller total package size.

The main challenge of this scheme is the difficulty of finding and applying the adhesive such that even in the presence of the electrolyte it still bonds to the cover slip and silicon die. The design is potentially expensive, especially with the TSVs, and requires mechanical assembly of separate parts, something which, from a manufacturing standpoint, is undesirable. On this last point however, it is conceivable that an entire wafer of devices could be covered in one step and then diced afterwards, although this further complicates the matter of dispensing the adhesive and electrolyte.

A slight variant of the design calls for the silicon substrate to be etched away, allowing the top of the CNTs to be level with the top of the substrate, thereby containing the electrolyte and greatly reducing the amount of adhesive that must be applied. However, such deep etching would be expensive to do on a large scale and would complicate the process of depositing the catalyst material and metalization as it would now have to be deposited in a recess. Other variants, which all include the same drawbacks illustrated above, involve etching the glass or making a three-layer sandwich structure.

5.3 Epoxy Encapsulation

The second proposed encapsulation method, which is shown in Figure 5.2, involves simply potting the entire device in an appropriate epoxy resin. This has the distinct advantage that it would be very inexpensive to perform once perfected and is already the dominant method of encapsulating semiconductor devices. It also has the advantage that the tops of the nanotubes would end up embedded in the encapsulant, protecting the fingers from moving inside the device after construction and shorting out. However, finding an appropriate encapsulant that will not itself contaminate the electrolyte either when liquid or when cured and which will form a sufficiently hermetic seal when dry that it will not allow contamination or evaporation of the electrolyte is a tall order. Acetonitrile, for example, is highly miscible with epoxy resin. Should such an encapsulant be found, this technique could be used with either TSVs or with bond pads and would allow for relatively easy integration of ultracapacitors with existing semiconductor packaging lines.

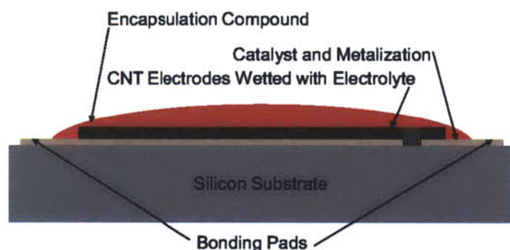

Figure 5.2: Cross-section of proposed potted package. Shown here without through silicon vias (not to scale).

To help with the problem of finding a compatible sealant, two modifications to the basic design described above are proposed. The first is to gel the electrolyte, holding it in place while the encapsulant is applied and allowed to cure. This was previously discussed in Chapter 4 in the context of providing mechanical support to the electrode fingers to prevent shorting. A second modification, and one which was experimentally demonstrated to a limited extent, is to spray a thin layer of material over the wetted electrodes before the encapsulent is applied to prevent the electrolyte from mixing with the uncured encapsulant. This inner material does not need to provide long term permeability resistance nor does it need to be mechanically strong. It only needs to separate the electrolyte from the encapsulant for long enough to allow full curing to occur.

To test this idea, a small droplet of water, representing a drop of aqueous electrolyte, was placed on a 1 cm square of passivated silicon, representing the substrate of a device. Onto this water was sprayed a thick coating of Krylon Appliance Enamel, chosen for its immiscibility with water and its ability to form a continuous layer on top of the droplet. The results were promising in that the paint dried and provided a remarkably tough skin over the water droplet. However, since the paint is thin and permeable the water evaporated out from under it in the course of a week. Figure 5.3 shows the experiment after the evaporation of the water. VacSeal, made by Space Environment Laboratories, was also tried as it has very low permeability. Unfortunately, it would not form a coherent film over the water droplet.

Ultimately, for all of these schemes, an advanced sealant needs to be found, ideally one that is immiscible with non-aqueous electrolytes. However, that goal proved to be beyond the scope of this project. The next chapter describes the results obtained so far.

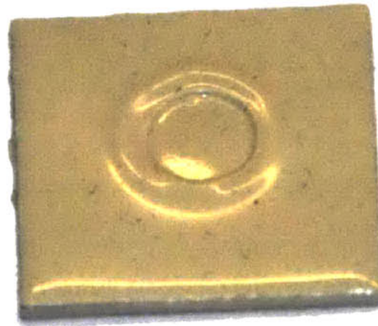

Figure 5.3: Experiment in encapsulating a droplet of water with Krylon Appliance Enamel. Note that due to the permeability of the paint, the water has evaporated despite being covered.

Chapter 6

Results and Analysis

This chapter describes the analysis methods used to evaluate finished devices, and details the results obtained from the samples produced over the course of this project. The results are divided into two sections. The first section describes the results of the initial samples that used wide fingers for maximum low-speed capacitance. The second section describes the results of the devices optimized for higher speeds.

6.1 Analysis of Cyclic Voltammograms

One of the primary measurement instruments used is a cyclic voltammeter. Made by Arbin, this instrument measures the current through the capacitor while sweeping the voltage linearly between bounds (usually -1 and 1 volt) in a triangle-wave pattern as in Figure 6.1. The current is then plotted as a function of voltage, to form a plot called a cyclic voltammogram (CV), an idealized version of which is shown in Figure 6.2. Analysis of a cyclic voltammogram is performed as follows. Per Equation 6.1 the voltammogram for an ideal capacitor should be a rectangle with the capacitance equal to the current divided by the voltage sweep rate. Thus, the capacitance for a device is determined by dividing the current recorded on the voltammogram at zero volts by the sweep rate. The current around zero voltage is used because the apparent capacitance increases with applied voltage due to parasitic effects.

$$i_c = C\frac{dV}{dt} \tag{6.1}$$

6.2 Results and Analysis for Wide Finger Devices Optimized for Maximum Capacitance Density

The initial devices created used very wide fingers (500 µm wide with 50 µm spacing) to maximize the capacitance per unit area that could be obtained, and thereby maximize the amount of energy that could be stored. This approach also simplified fabrication and minimized the potential for shorting between adjacent fingers (due to the wide spacing). Specifically, only approximately 8% of the active device area did not have nanotubes on it and thus very large specific capacitances, the highest known to be reported for capacitors

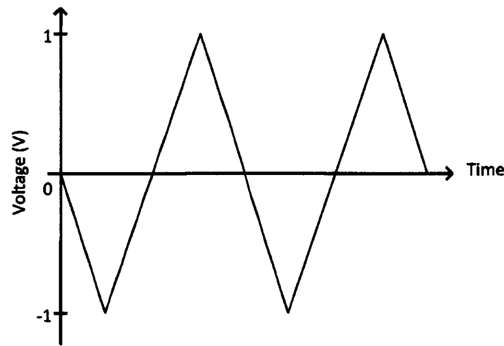

Figure 6.1: Voltage waveform applied to the capacitor by the cyclic voltammeter

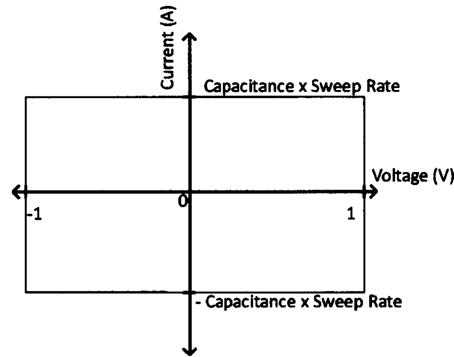

Figure 6.2: Idealized cyclic voltammogram for a capacitor

on silicon, were obtained [22].

For repeatability, three devices of this nature were constructed and characterized. Performing CVs on these devices yielded a wide range of capacitances, presumably due to variations in the consistency, height, and density of the carbon nanotube growth. Specifically, capacitance values of 0.67 mF, 1.6 mF, and 3.5 mF were obtained. Figures 6.3(a), 6.3(b), and 6.3(c) show the voltammograms corresponding to these three samples. The sweep rate for the graphs shown is 50 mV/sec, chosen because it is slow enough that the ESR and the observed decrease in capacitance with frequency become negligible. Since this device is planar and intended for eventual use in compact electronics, the specific capacitance (capacitance per unit area) is important for comparison with other similar technologies. The active area of the interdigitated nanotube structure, without the connecting terminals, has dimensions of 7 mm x 9.5 mm, which corresponds to 66.5 mm^2, yielding a maximum capacitance density of 52.6 $\mu F/mm^2$, which is over an order of magnitude higher than the 4.28 $\mu F/mm^2$ achieved by Y. Q. Jiang with a structurally similar planar ultracapacitor [7]. For comparison, this is over two orders of magnitude higher than for typical trench capacitors, which at about 450 nF/mm^2, is the current state-of-the art for high-value on-die capacitors [10].

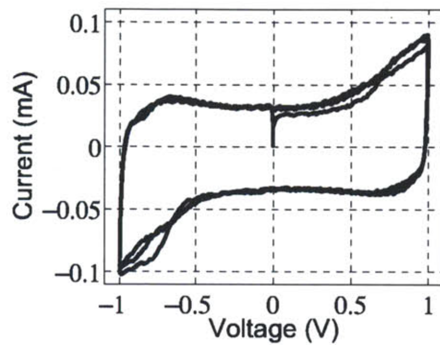

(a) 0.67mF sample

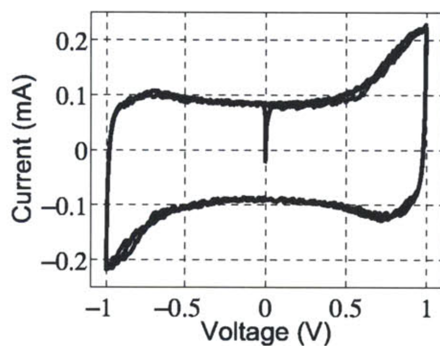

(b) 1.6mF sample

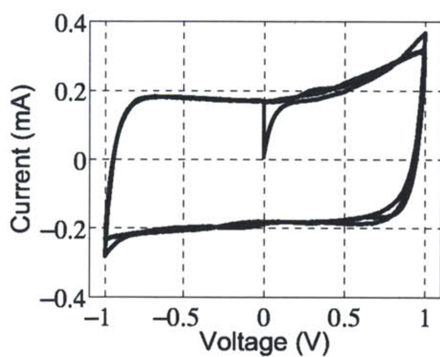

(c) 3.5mF sample

Figure 6.3: Cyclic voltammograms for wide finger devices

6.3 Analysis of Impedance Spectrograms

After confirming the concept of an interdigitated electrode capacitor and proving that high-value devices could be created, the next step was to optimize the speed of the device. As explained in Chapters 1 and 3, the speed of an ultracapacitor is limited primarily by the ionic resistance of the electrodes, which in theory should scale linearly with the distance ions have to migrate into the electrode. Thus, the ionic resistance of the device should scale linearly with finger width and verifying this was one of the major goals of this project as it demonstrates the interdigitated electrode miniature ultracapacitor (IDEMU) concept - the possibility of creating high specific capacitance, high speed ultracapacitors by altering the mechanical structure of the electrode. However, to do that a new measurement tool was needed. While the CV does a good job of determining capacitance, and the speed of the device can be crudely estimated by increasing the sweep rate of the CV and noting the change in measured capacitance, it does not allow the individual sources of internal resistance to be separated from each other. Instead, this property was measured using electrochemical impedance spectroscopy, a technique frequently used to characterize the performance of batteries and fuel cells.

The data for an impedance spectrogram is taken by measuring the complex impedance (both real and imaginary components) of the device over a very wide range of frequencies, in the case of this work, 0.1 Hz to 100 kHz. This data is then plotted with the real component on the x-axis, the negative of the imaginary component on the y-axis, and the frequency as a parametric variable. The following subsections describe the impedance spectrograms for specific components.

6.3.1 Resistor

A circuit consisting of a single resistor will have a single point on the x-axis at the value of the resistor because a resistor has a constant real impedance independent of frequency. The impedance spectrogram for this case is shown in Figure 6.4

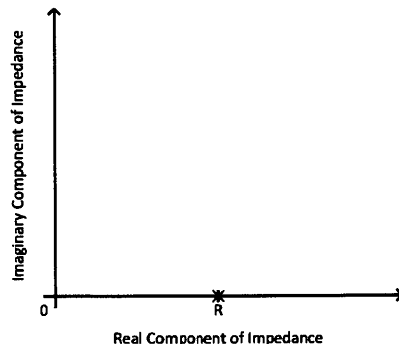

Figure 6.4: Idealized impedance spectrogram for a resistor

6.3.2 Capacitor

Similarly, the impedance of a capacitor, as shown in Equation 6.2, is negative and purely imaginary. Thus, since the negative impedance is plotted along the y-axis of an impedance spectrogram, an ideal capacitor will form a vertical line going down from infinity with increasing frequency along the y-axis. The impedance spectrogram for this case is shown in Figure 6.5

$$Z = \frac{-j}{\omega C} \tag{6.2}$$

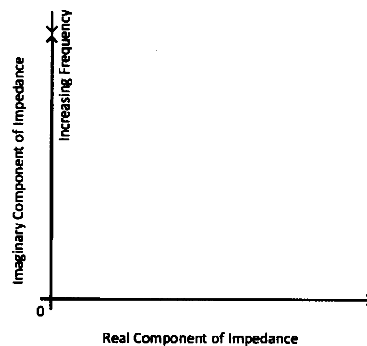

Figure 6.5: Idealized impedance spectrogram for a capacitor

6.3.3 Series Resistor and Capacitor

Combining the two cases above, a circuit consisting of a resistor in series with a capacitor will form a vertical line coming up from the x-axis at the value of the series resistance. This is because the real component of the impedance will always remain constant at value R, and the imaginary component will decrease from infinity with frequency. The impedance spectrogram for this case is shown in Figure 6.6.

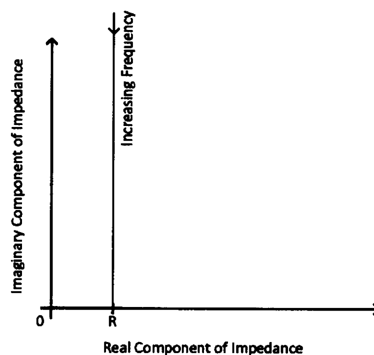

Figure 6.6: Idealized impedance spectrogram for a capacitor in series with a resistor

6.3.4 Parallel Resistor and Capacitor

A more complex case consists of a resistor in parallel with a capacitor. This will form a semi-circle with diameter equal to the value of the parallel resistance as derived below. First, the impedance of the parallel RC circuit is found by taking the inverse of the sum of the inverses of the impedances of the resistor and the capacitor, resulting in Equation 6.3.

$$Z = \frac{1}{\frac{1}{R} + j\omega C} \tag{6.3}$$

Then, Equation 6.3 can be simplified as shown in 6.4 such that the real and imaginary components can be separated.

$$Z = \frac{R}{1 + Rj\omega C} = \frac{R(1 - Rj\omega C)}{(1 + Rj\omega C)(1 - Rj\omega C)} = \frac{R - R^2 j\omega C}{1 + R^2\omega^2 C^2} \tag{6.4}$$

Next, the real and imaginary components of the impedance are separated resulting in equations 6.5 and 6.6.

$$Z_{real} = \frac{R}{1 + R^2\omega^2 C^2} \tag{6.5}$$

$$Z_{imag} = \frac{-R^2\omega C}{1 + R^2\omega^2 C^2} \tag{6.6}$$

Finally, since an impedance spectrogram plots the negative imaginary component of the impedance versus the real component, y is taken to be negative Z_{imag}. To further simplify, a parametric variable k is set equal to the quantity $R\omega C$. This results in the parametric equations 6.7 and 6.8, which define x and y in terms of the parametric variable k. These equations are easily recognized as those for a semicircle of radius R extending from $x = y = 0$ to $x = R, y = 0$.

$$x = \frac{R}{1 + k^2} \rightarrow \frac{x}{R} = \frac{1}{1 + k^2} \tag{6.7}$$

$$y = \frac{Rk}{1 + k^2} \rightarrow \frac{y}{R} = \frac{k}{1 + k^2} \tag{6.8}$$

The impedance spectrogram for this case is shown in Figure 6.7.

6.3.5 Ionic Resistances

For reasons beyond the scope of this document, a pure ionic resistor will form a 45° line starting from the origin as shown in Figure 6.8 [5].

6.3.6 Equivalent Circuit Model of an Ultracapacitor

Combining the cases described in the sections above, the equivalent circuit shown in Figure 6.9 can be obtained from the idealized impedance spectrogram shown in 6.10. This is one model of the equivalent circuit for an ultracapacitor. The electrical series resistance arises from the resistance of the current collector (as described in Chapter 3), the resistance of the connection between the current collector and the connecting

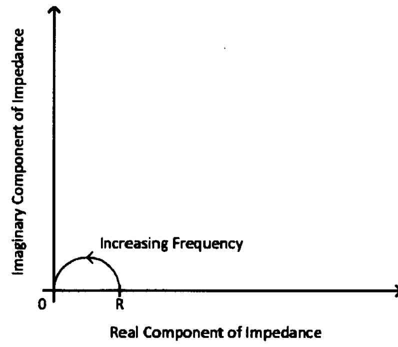

Figure 6.7: Idealized impedance spectrogram for a capacitor in parallel with a resistor

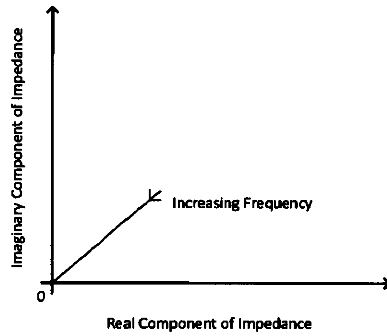

Figure 6.8: Idealized impedance spectrogram for an ionic resistor

wires, and the resistance of the connecting wires themselves. The contact resistance and contact capacitance arise from the fact that there is resistance between the nanotubes and the current collector and that the bottoms of the nanotubes and the top of the current collector form a small parallel plate capacitor. Finally, the ionic resistance arises from the time it takes ions to migrate into and out of the electrodes as discussed in detail in Chapter 1.

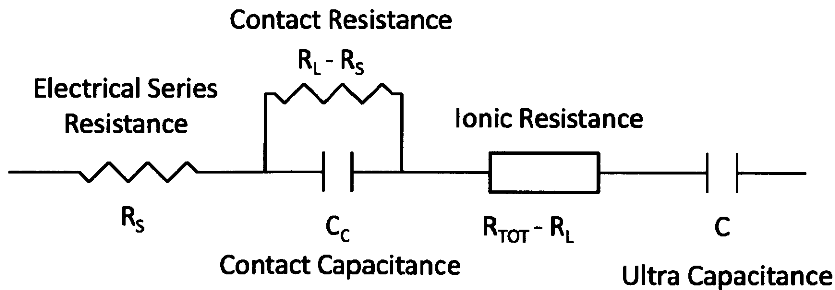

Figure 6.9: Equivalent circuit for an ultracapacitor

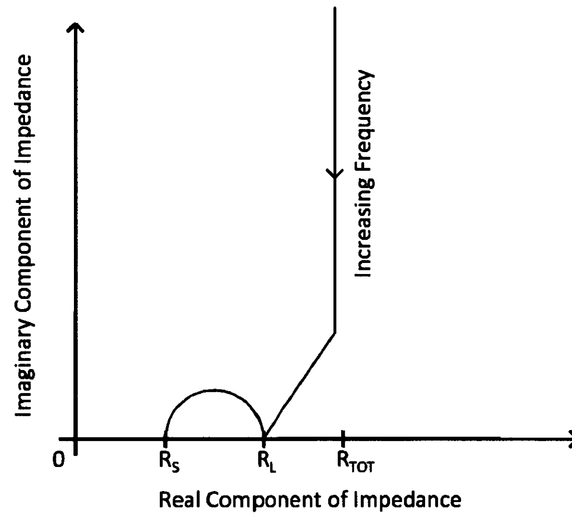

Figure 6.10: Idealized impedance spectrogram for an ultracapacitor. R_{TOT} is the series resistance, the quantity (R_{TOT} - R_L) is the ionic resistance, the quantity (R_L - R_S) is the contact resistance, and R_S is the electrical series resistance.

6.4 Results for High-Speed Narrow Finger Devices

Samples were created with narrow fingers to test the hypothesis that the ionic resistance of the device can be decreased, and therefore the speed increased, by decreasing the width of the fingers. Two such sets were created to validate the results. Each set of samples had four different finger widths, 10 µm, 20 µm, 40 µm, and 80 µm, all spaced 150 µm apart. These two sets will be referred to as sample set one and sample set two throughout this document. After growing these two sets of samples, they were characterized with the impedance spectrometer to find their ionic resistances and with the cyclic voltammeter, to find their low frequency capacitances. From the resulting impedance spectrograms (shown in Appendix A) the total electrical series resistance (consisting of the electrical and contact resistances) and the ionic resistance were extracted as explained in the previous section. Figure 6.11 is an example of a real EIS plot with the point indicating R_{TOT} marked with a circle and the point representing R_L marked with an 'X'. Because of the frequency limitation of the impedance spectrometer, the semi-ellipse representing the contact resistance and contact capacitance is incomplete.

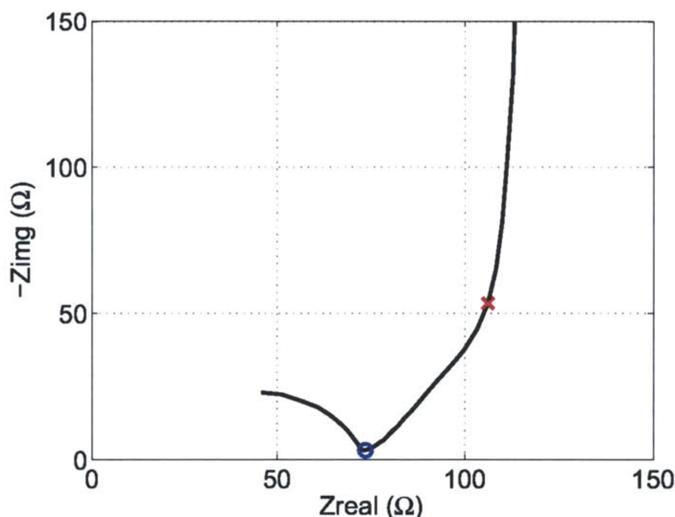

Figure 6.11: Example of resistances being extracted from real EIS data. R_{TOT} is marked with an 'X' and R_L is marked with a circle.

Note, however, that due to non-idealities in the real impedance spectrograms there is significant ambiguity in describing the location of the inflection point that determines the transition between the regions of the plot representing the ionic resistance and the ultra capacitance and thus in determining the value of the ionic resistance. The largest source of this error is that due to leakage resistance across the ultracapacitor, the 'vertical' line representing the ultracapacitance is not actually vertical but starts to curve down again in another semi-ellipse. This curving makes it hard to distinguish the end of the 45° line representing the ionic resistance from the start of the large semi-ellipse representing the ultracapacitance.

Table 6.1 and Table 6.2 list the ionic resistance, total electrical series resistance, and capacitance for each of the samples in the two sample sets. Note that unrounded data is NOT meant to imply any specific degree of precision.

Table 6.1: Data collected from first set of variable finger width samples

Finger Width (μm)	Total ESR (R_{TOT}) (Ω)	Ionic Resistance (R_{TOT} - R_L) (Ω)	Capacitance (mF)
10	45.64	4.01	1.39
20	39.34	8.11	2.33
40	59.28	16.34	2.51
80	108.8	15.1	0.85

Table 6.2: Data collected from second set of variable finger width samples

Finger Width (μm)	Total ESR (R_{TOT}) (Ω)	Ionic Resistance (R_{TOT} - R_L) (Ω)	Capacitance (mF)
10	60.95	6.23	1.27
20	40.29	23.37	1.14
40	73.56	32.64	1.54
80	78.43	49.47	1.68

6.5 Analysis of Results for High-Speed Narrow Finger Devices

Here the results for the variable finger width devices are analyzed.

6.5.1 Analysis of Relationship Between Finger Width and Ionic Resistance

The first, and arguably most important question that can be answered from this data is whether or not there is indeed a linear relationship between finger width and ionic resistance. To that end, the ionic resistance for both sets of data is plotted against the finger width in Figure 6.12. From this figure, it can be seen that ionic resistance increases roughly linearly between 10 μm and 40 μm, but then starts to drop off at 80 μm. However, this drop-off is expected because the nanotubes are on the order of 100 μm high, so when the finger width reaches 80 μm the nanotubes fingers are as wide as they are tall, meaning that the ions have three paths into and out of them - two through the sides and one through the top, thus decreasing the ionic resistance. For comparison, the thinner fingers are so much taller than they are wide that the ion conductivity into the fingers from the top is negligible. Furthermore, the 80 μm point for sample set one is suspect as the capacitance measured for that sample is much lower than the other even though, being the widest finger sample (and thus having the least wasted space between fingers), it should have the largest capacitance of the three. For these reasons, the best-fit lines only include the first three points in each data set.

The second observation is that the slopes of the two sample sets are not the same. This too, is to be expected. Since the samples were grown at different times on different substrates, the characteristics, such as growth density and height of the nanotubes is different. Per Equation 6.9, the ionic resistance depends on the finger width, L, the finger area, A, which is the finger width times the nanotube height, and the ionic conductivity of the material, ρ_{ionic}. Thus, a difference in nanotube height, or a difference in ionic conductivity could easily cause a different slope in the linear relationship between finger width and ionic resistance.

Finally, it is observed that the intercepts for both lines are near zero. This is also expected as a zero length ionic resistor should have zero ionic resistance!

$$R_{ionic} = \frac{\rho_{ionic} L}{A} \tag{6.9}$$

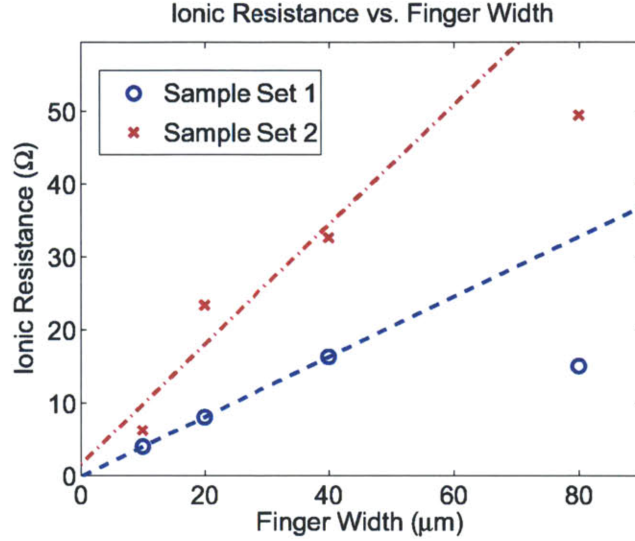

Figure 6.12: Ionic resistance vs. finger width for sample set 1 and sample set 2

6.5.2 Background on Time Constant

Determining that ionic resistance is linear with finger width is good, but ultimately what matters for filtering is the time constant. An RC circuit, in this case formed by the series electrical and ionic resistance of the ultracapacitor in series with the capacitance of the ultracapacitor, forms a first order low-pass filter circuit with cutoff frequency f_c, which is related to the time constant τ as shown in Equation 6.10. The time constant, in turn, is related to R and C per Equation 6.11.

$$f_c = \frac{1}{2\pi\tau} \tag{6.10}$$

$$\tau = RC \tag{6.11}$$

For frequencies below the cutoff frequency, the filter does not reduce their amplitude. However, for frequencies above the cutoff frequency, the amplitude is decreased by 20 dB (ten times) for every factor of ten that the input is higher than the cutoff frequency. Thus, to effectively filter a 120 Hz signal, the cutoff frequency must at minimum be about ten times higher, or 1.2 kHz. Solving for the time constant for a 1.2 kHz cutoff frequency yields 0.13 ms. Similarly, a cutoff frequency of 120 Hz corresponds to a time constant of 1.3 ms.

6.5.3 Analysis of Time Constant vs. Finger Width Data

To compare the time constants needed for 120 Hz filtering with those obtained so far, the series electrical resistance is neglected as it can presumably be reduced through the methods discussed in Chapter 3 and is not a fundamental limit to the speed of an ultracapacitor. With this assumption, the time constant is computed as the product of the ionic resistance and the low frequency capacitance of the device. This time constant is plotted against finger width in Figure 6.13

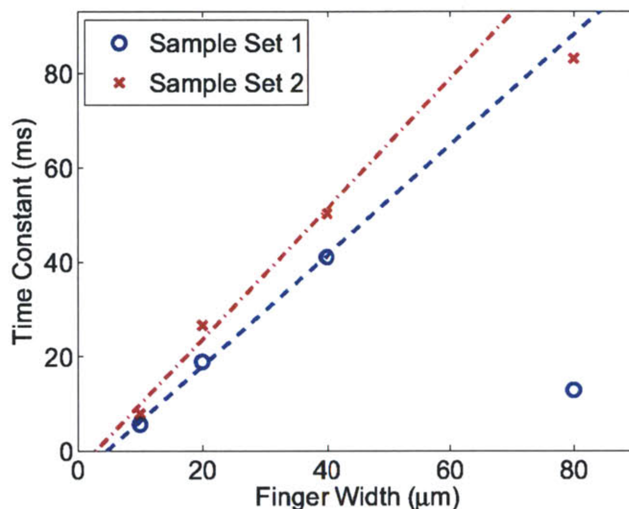

Figure 6.13: Time constant vs. finger width for sample set 1 and sample set 2

From this figure it is obvious that, neglecting the final point in each data set for the reasons discussed previously, the time constant is very linear with finger width and intercepts near zero, both properties that are expected. This provides strong evidence that the time constant of an ultracapacitor can be controlled purely by adjusting the geometry of the electrodes, a result that is not known to have been shown before. The slopes of the correlation are similar for both sets of samples, being 1.17 ms per μm for sample set one and 1.38 ms per μm for sample set two. Because these are roughly similar, it is acceptable to average them to about 1.3 such that a rough estimate of the finger width needed for 120 Hz filtering can be obtained. This admittedly crude model predicts that to form a filter with a cutoff frequency of 120 Hz (a time constant of 1.3 ms), the absolute minimum for any 120 Hz filtering to occur, the finger width would have to be about 1.7 μm wide. For effective filtering, a cutoff frequency of 1.2 kHz corresponding to a time constant of 0.13 ms, the width would need to be about 0.17 μm. Although the narrowest finger widths created in this project were 10 μm, there is no reason why, using modern lithography techniques, the width of the fingers could not be reduced to 0.17 μm or thinner as the nanotubes themselves are on the order of 5 to 10 nm in diameter [5].

6.6 Ion Migration Distance

The basis for the ionic resistance model for porous electrode material is that the ions must physically move into and out of the electrode structure, and the results of the preceding chapter are based and interpreted upon this premise. To verify this assumption, in this section the model is explored analytically.

The diagram shown in Figure 6.14 shows a simple model of a parallel-plate carbon nanotube ultracapacitor with 1 cm^2 plate area in the uncharged state. Based on previous results, such a device has a capacitance of approximately 15 mF at 1 V in aqueous sodium sulfate electrolyte. Equation 6.12 calculates the charge on a single plate of a capacitor based on the capacitance, C, and the voltage across the device, v. Using this formula, when fully charged, there will be a charge of 0.015 coulombs stored on each plate.

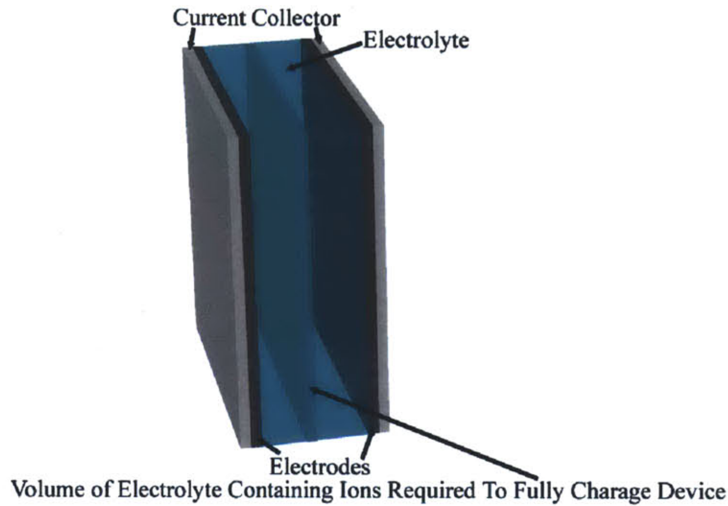

Figure 6.14: A simple model of a parallel-plate carbon nanotube ultracapacitor with 1 cm^2 plate area in the uncharged state (spacing between plates is drawn larger than scale)

$$q = Cv = (15 \times 10^{-3}\ \text{F})(1\ \text{V}) = 0.015\ \text{C} \tag{6.12}$$

Sodium ions have a charge of +2 and sulfate ions have a charge of -2 and the sodium sulfate electrolyte has a concentration of 1 mol / L. Equation 6.13 calculates the volume, V, of electrolyte needed to hold a charge, q, worth of positive or negative ions, based on a molarity, m, of the electrolyte, the charge of an electron, 1.602 x 10^{-19}, and Avogadro's constant, 6.022 x 10^{23}. Therefore, by equation 6.13 a volume of 0.00518 cm^3 will hold all of the ions needed to fully charge one plate of the device.

$$V = 0.015\ \text{Coulombs} \frac{1\ \text{Carrier}}{1.602 \times 10^{-19}\ \text{Coulombs}} \frac{1\ \text{Ions}}{2\ \text{Carriers}} \frac{1\ \text{L}}{6.022 \times 10^{23}\ \text{ions}} \frac{1000\ \text{mL}}{1\ \text{L}} = 7.8 \times 10^{-5}\ \text{mL} \tag{6.13}$$

Thus, according to this model, within the electrolyte, a 1 cm x 1 cm section of positive ions will move approximately 78 µm toward the negatively charged electrode, while a 1 cm x 1 cm section of negative ions

will move approximately 78 µm toward the positively charged electrode in order to fully charge the device. The diagram in Figure 6.15 shows the effective movement of the ions when the device is fully charged. The red layer represents the positive ions, while the green layer represents the negative ions. Of course, in the real device the ion movement is more complex, but this model does confirm that there is significant movement of ions into and out of the electrode material (which is about 100 µm thick), so it should be expected that mechanical resistance to ion movement by the porous electrode structure affects the device speed significantly.

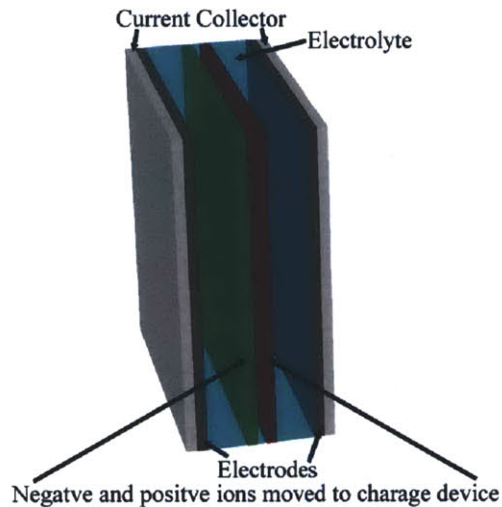

Figure 6.15: A simple model of a parallel-plate carbon nanotube ultracapacitor with 1 cm^2 plate area showing effective ion movement when the device charges (spacing between plates is drawn larger than scale)

Based on these results, two important conclusions can be drawn. First, a device was produced on silicon with an extremely high specific capacitance of 52.6 $\mu F/mm^2$, the highest known to have been reported for capacitors on silicon. Second, the finger spacing and time constant of the devices appears to scale linearly with finger width, demonstrating that the time constant of an ultracapacitor can be varied by simply adjusting the mechanical structure of the electrode, which so far as is known, is a novel concept.

Chapter 7

Conclusion

The goal of this project was to explore the possibility of designing miniature ultracapacitors that can compete favorably with electrolytic capacitors in select applications. Specifically, the following questions were asked.

- What should the electrodes be made of?
- What mechanical structure should be used?
- What electrolyte should be used?
- How should the device be packaged?
- What can be achieved using the time, money, and technology available right now?

Work was performed on each of these and the results are summarized below.

What should the electrodes be made of?

It was determined that vertically aligned carbon nanotubes are a good choice to create miniature high-speed ultracapacitors for two main reasons. First, because they create an electrode with uniform pore spacing, the devices can operate faster than those created with activated carbon electrodes. Second, because the nanotubes are grown on silicon using low pressure chemical vapor deposition, this technique also opens the possibility of high-value integrated (on-die) capacitors.

What mechanical structure should be used?

A structure consisting of interdigitated fingers appears to be promising, again for two separate reasons. First, it allows the entire device to be created on a single piece of silicon with no mechanical assembly. This should help make manufacturing inexpensive and compatible with existing semiconductor technology. Second, and most significantly, using an interdigitated structure allows the ionic resistance, the main limiting factor to ultracapacitor speed, to be reduced without sacrificing their large capacitance, by making the fingers very narrow.

What electrolyte should be used?

This remains partially answered. In exploring different electrolytes, it was determined that aqueous electrolytes are safe, inexpensive, and work well, but they cause the nanotubes to delaminate from the substrate in a matter of hours, excluding their use in a practical device. On the other hand, non-aqueous electrolytes (acetonitrile and propylene carbonate) do not appear to cause long-term device deterioration of the nature seen with aqueous electrolytes, but they do, for unknown reasons, cause the thin nanotube fingers to bend over and short to each other, preventing a usable device from being constructed. However, overall they seem more promising than the aqueous electrolytes and the shorting problem may be able to be overcome by forming a gel out of the electrolyte, which would provide mechanical stability to the fingers.

How should the device be packaged?

This question is also open, though there are two attractive possibilities. The first is to cover the device with a glass coverslip adhered on the edges with epoxy. The advantage of this is that the relative impermeability of glass would prevent significant electrolyte evaporation. However, it is not clear how to easily fabricate and fill the structure. The second possibility is to find a material in which the entire device can be potted successfully. This has the advantage that it should be extremely simple to manufacture the devices. However, the requirements on a sealing compound are severe in that it must be compatible with the electrolyte both in liquid and cured form and not allow for significant electrolyte evaporation.

What can be achieved using the time, money, and technology available right now?

Specific achievements of this project are as follows. First, it was demonstrated that the interdigitated electrode concept does indeed produce a usable capacitor. Second, a device was produced on silicon with an extremely high specific capacitance of 52.6 $\mu F/mm^2$, the highest known to have been reported for capacitors on silicon. Finally, and most significantly, strong evidence was presented that the ionic resistance of devices with sufficiently narrow fingers scales with the finger width, meaning that the time constant of an ultracapacitor can be varied by simply adjusting the mechanical structure of the electrode, an idea not known to have been explored before. Together, these advancements enable the construction of faster ultracapacitors with lower internal resistance and higher capacitance values.

Future Work

This project only attempted to explore the concept of a miniature integrated high speed ultracapacitor. There is still much more to be done to create a working, and ultimately manufacturable, device. Specifically:

- The growth process needs to be adapted to commercial equipment and optimized for reliable growth of nanotube electrodes with a high yield rate.
- An electrolyte needs to be found that does not cause the device to short or disintegrate over time.

- An easily manufacturable and sufficiently reliable packaging system needs to be devised.

Based on the promising results of this early work, I am confident that this technology will continue to evolve toward its ultimate goal of displacing unreliable and space-wasting electrolytic capacitors in a variety of applications.

Bibliography

[1] C. Pollak. "Electrolytic Current Rectifier and Condenser." U.S. Patent 672 913, Mar. 7, 1901.

[2] A. Douglas. *Radio Manufacturers of the 1920s*, vol. 1, Vestal, NY: Vestal Press, 1988, p. xii.

[3] R. A. Rightmireus. "Electrical Energy Storage Apperatus" U.S. Patent 3 288 641, Nov. 29, 1966.

[4] J. Miller. "Graphene Double-Layer Capacitor with ac Line-Filtering Performance". *Science*, vol. 329, pp. 1637-1639, Sep. 2010.

[5] R. L. Signorelli, "High energy and power density nanotube-enhanced ultracapacitor design, modeling, testing, and predicted performance," Ph.D. thesis, Dept. Elect. Eng., Mass. Inst. of Tech., Cambridge, MA, 2009.

[6] M. Cantoro et al, "Catalytic Chemical Vapor Deposition of Single-Wall Carbon Nanotubes at Low Temperatures," *Nano Letters*, vol. 6, pp. 1107-1112, May 2006.

[7] Y. Q. Jiang et al, "Planar MEMS Supercapacitor using Carbon Nanotube Forests," *J. Microelectromech. Syst.*, pp. 587-590, Jan. 2009.

[8] H. J. In et al, "Microfabrication Of Electrochemical Capacitors With Vertically-Aligned Carbon Nanotube Electrodes," *Proceedings Power MEMS*, pp. 137-140, 2009. Available: `http://cap.ee.imperial.ac.uk/~pdm97/powermems/2009/pdfs/papers/037_0179.pdf`.

[9] D. C. Ku, "Methodology, Morphology, and Optimization of Carbon Nanotube Growth for Improved Energy Storage in a Double Layer Capacitor," M.S. thesis, Dept. Elect. Eng., Mass. Inst. of Tech., Cambridge, MA, 2007.

[10] K. B. P. Jinesh, "Dielectric properties of atomic-layer-deposited $La_yZr_{1-y}O_x$ and $Er_yHf_{1-y}O_x$ thin films," Ph.D. dissertation, University of Twente, Enschede, Netherlands, 2010.

[11] J. G. Kassakian et al., *Principles of Power Electronics.* Reading, MA: Addison-Wesley Corp., 1992.

[12] Panasonic Corp., "Reliability of Aluminum Electrolytic Capacitors." Available: `http://industrial.panasonic.com/www-data/pdf/ABA0000/ABA0000TE4.pdf`.

[13] *CRC Handbook of Chemistry and Physics*, 88th ed., CRC Press, Boca Raton, FL, 2008, sec. 12, pp. 39-40.

[14] Y. Murakami et al, "Detachment of vertically aligned single-walled carbon nanotube films from substrates and their re-attachment to arbitrary surfaces," *Chemical Physics Letters*, vol. 422, pp. 575-580, Dec. 2006.

[15] Q. Qu et al, "Study on electrochemical performance of activated carbon in aqueous Li_2SO_4, Na_2SO_4 and K_2SO_4 electrolytes," *Electrochemistry Communications*, vol. 10, pp. 1652-1655, Aug. 2008.

[16] M. Motoyoshi, "Through-Silicon Via (TSV)," *Proceedings of the IEEE*, vol. 97, pp. 43-48, Jan. 2009.

[17] J. Schindall, "The Charge of the Ultra - Capacitors," *IEEE Spectrum*, vol. 44, no. 11, pp. 42-46, Nov. 2007.

[18] J. Schindall, "What's in a Name? A New Model for Regenerative Electrical Energy Storage," *IEEE Power Electronics Society Newsletter*, vol. 20, no. 1, pp. 32-34, 1st Qtr., 2008.

[19] G. Zhong et al., "Acetylene: A Key Growth Precursor for Single-Walled Carbon Nanotube Forests," *The Journal of Physical Chemistry*, vol. 113, no. 40, pp. 17321-17325, Sep. 2009.

[20] J. Chmiola et al., "Anomalous Increase in Carbon Capacitance at Pore Sizes Less Than 1 Nanometer," *Science*, vol. 313, pp. 1760-1763, Sep. 2006.

[21] J. Reed, *Principles of Ceramics Processing*, 2nd Edition, John Wiley and Sons, New York, 1995, pp. 157-167.

[22] M. D'Asaro et al., "Miniature Carbon-Nanotube-Enhanced Ultracapacitor for Electronics Applications," 21st International Seminar on Double Layer Capacitors & Hybrid Energy Storage Devices, Fort Lauderdale, FL, USA, Dec. 2011.

Appendix A

Raw Impedance Spectroscopy Data

In this appendix are the unprocessed impedance spectrograms for both sets of data used to find the relationship between finger width and ionic resistance. These are included because there is significant ambiguity in describing the location of the inflection point that determines the transition between the regions of the plot representing the ionic resistance and the ultra capacitance, and thus in determining the value of the ionic resistance. Thus, should the reader have any doubts about the accuracy with which the analysis was performed, the calculations may be repeated using the plots shown here.

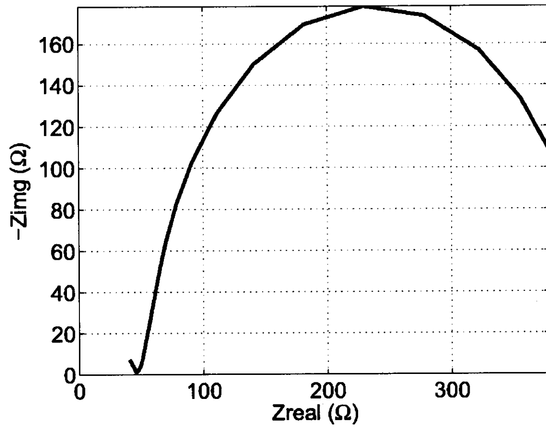

(a) Sample with 10 μm wide fingers

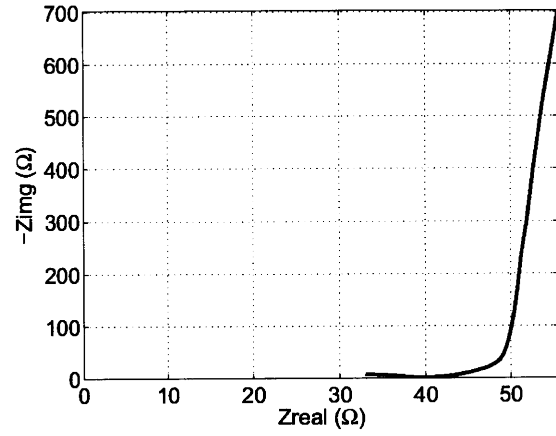

(b) Sample with 20 μm wide fingers

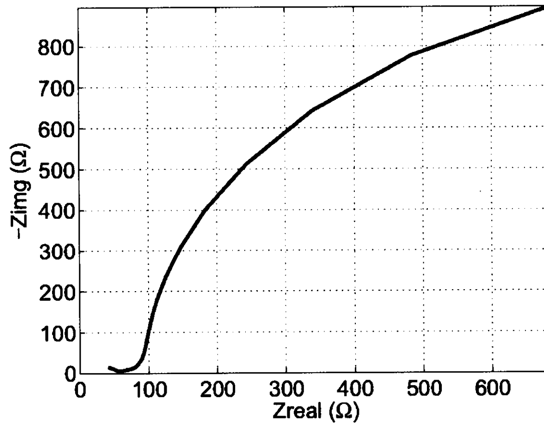

(c) Sample with 40 μm wide fingers

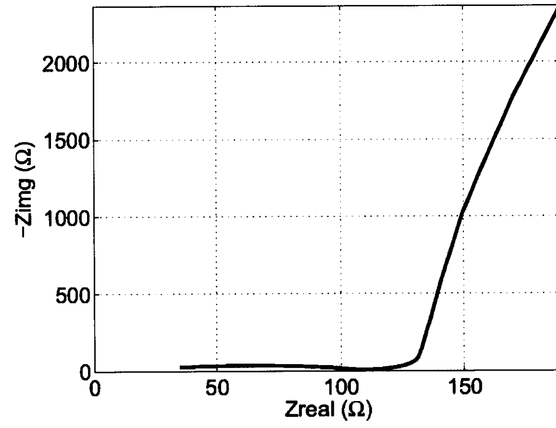

(d) Sample with 80 μm wide fingers

Figure A.1: Unscaled impedance spectrograms for first set of samples

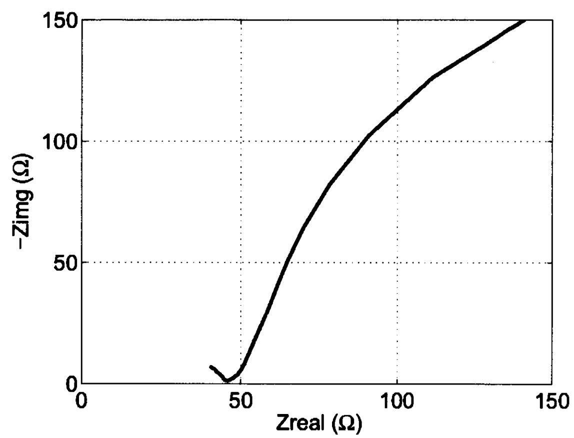

(a) Sample with 10 μm wide fingers

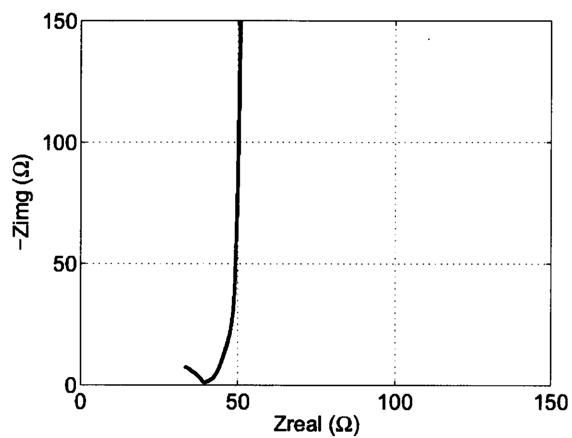

(b) Sample with 20 μm wide fingers

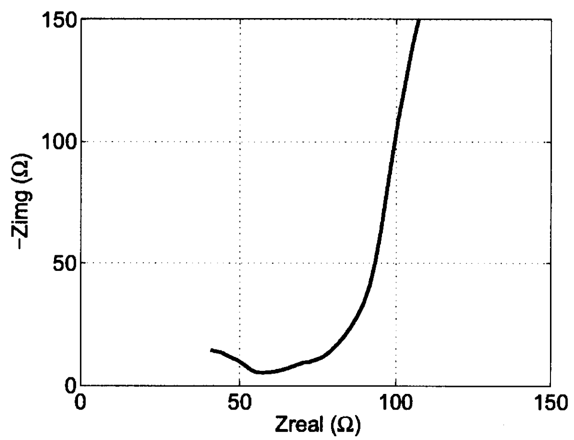

(c) Sample with 40 μm wide fingers

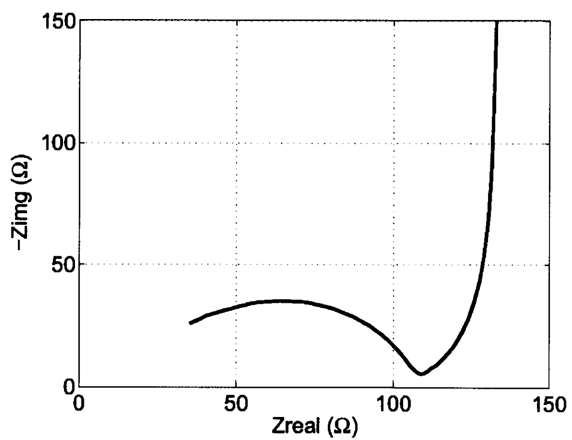

(d) Sample with 80 μm wide fingers

Figure A.2: Impedance spectrograms scaled to show ionic resistance for first set of samples

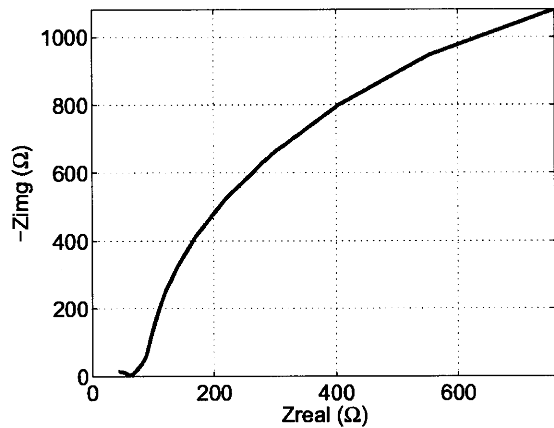

(a) Sample with 10 µm wide fingers

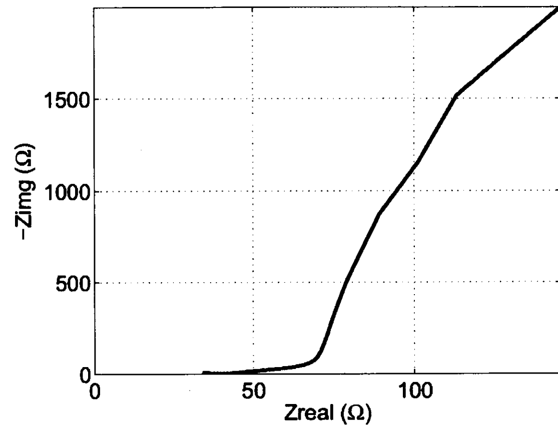

(b) Sample with 20 µm wide fingers

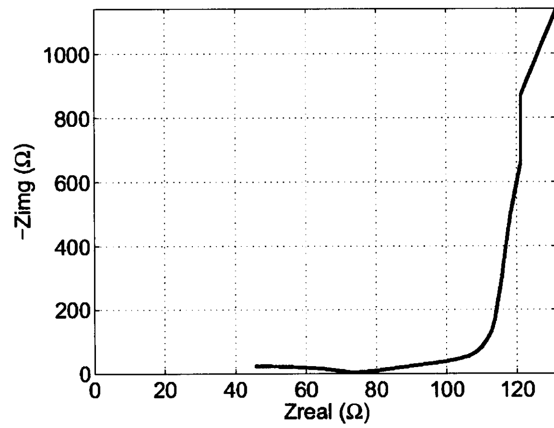

(c) Sample with 40 µm wide fingers

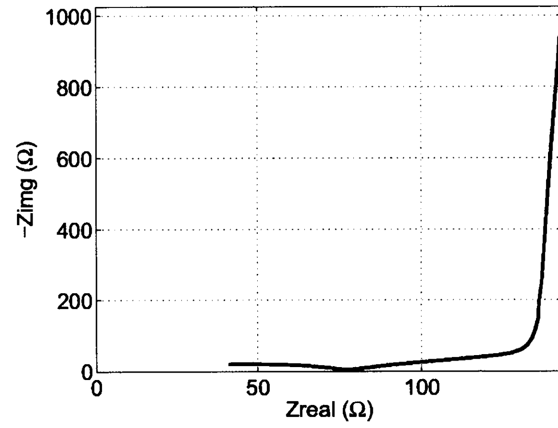

(d) Sample with 80 µm wide fingers

Figure A.3: Unscaled impedance spectrograms for second set of samples

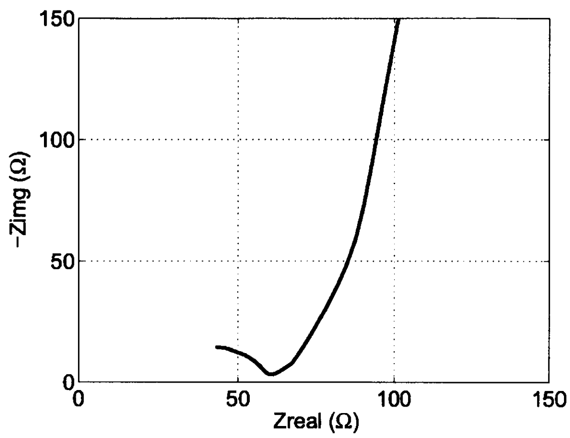

(a) Sample with 10 μm wide fingers

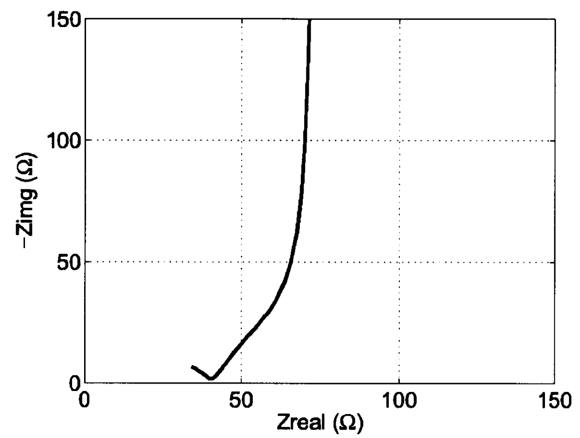

(b) Sample with 20 μm wide fingers

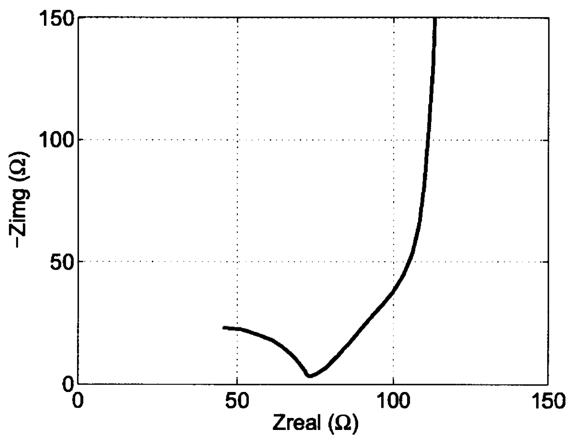

(c) Sample with 40 μm wide fingers

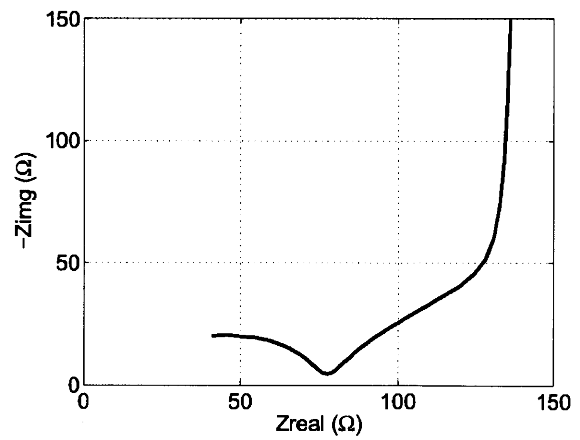

(d) Sample with 80 μm wide fingers

Figure A.4: Impedance spectrograms scaled to show ionic resistance for second set of samples

Appendix B

Growth Procedure

Reproduced below is the procedure used to operate the LPCVD growth equipment used to grow nanotubes on all of the samples discussed in this document.

Phase 1: Preheating

1. Turn on the power to all the equipment and open the cylinder shutoff valves and auxiliary gas shutoff valves (blue valves) if they are closed.
2. Open the large and small vacuum valves and the chamber isolation valve.
3. Turn the heater current up to 9 A.
4. Wait until the metal ends of the vacuum chamber are warm to the touch (about 15 min).

Phase 2: Sample Insertion

1. Close the large and small vacuum valves and the chamber isolation valve.
2. Move the Air/Vac valve to the Air position and wait for the chamber to vent.
3. Loosen the stainless steel ring on the end of the chamber until it comes off the threads.
4. Pull the end of the chamber off and push it to the side.
5. Using tweezers carefully place the sample on the center of the heater.
6. Adjust the rotation of the heater support so that it is parallel to the work surface.
7. CAREFULLY replace the end of the chamber and tighten the stainless steel ring. If this is not done with sufficient care the sample will slide off the heater from the vibrations.
8. Align the infrared thermometer so that it is directly underneath the center of the sample.
9. Move the Air/Vac valve to the Vac position.

Phase 3: Evacuation

1. Slowly open the small vacuum valve.
2. Wait until the chamber pressure is under 10 torr.
3. Open the large vacuum valve.

4. Open the chamber isolation valve.
5. Turn on the argon and set it to the value listed in the reduction column of Table B.1. If this is the first run of the day, also turn on the other gases to their growth flow rates.
6. Wait five minutes.
7. Turn off the argon.
8. Wait until the pressure in the chamber is under 10 mtorr.

Phase 4: Reduction

1. Set the gas flow rates to the values listed in the reduction column of Table B.1.
2. Close the large vacuum valve.
3. Wait 5 minutes.

Phase 5: Growth

1. Set the gas flow rates to the values listed in the growth column of Table B.1.
2. Adjust the heater control unit for a growth temperature of 780 °C, turn on the heater power, and wait until the green light illuminates, indicating that the heater is up to temperature. Refer to Appendix D for instructions regarding the use of the heater control unit.
3. Wait 15 minutes for growth.
4. Turn the heater current off using the heater control unit.
5. Turn off the hydrogen and acetylene gas.
6. Open the large vacuum valve.
7. Allow argon to continue flowing for 5 minutes and then shut it off as well.

Phase 6: Sample Removal

1. Close the large and small vacuum valves and the chamber isolation valve.
2. Move the Air/Vac valve to the Air position and wait for the chamber to vent.
3. Loosen the stainless steel ring on the end of the chamber until it comes off the threads.
4. Pull the end of the chamber off and push it to the side.
5. Carefully remove the sample using tweezers.
6. If another sample is being grown, return to step 5 of phase 2; else continue to phase 7.

Phase 7: System Shutdown

1. Replace the end of the chamber and tighten the stainless steel ring.
2. Move the Air/Vac valve to the Vac position.
3. Open the large vacuum valve and allow the chamber to evacuate to 1 torr or lower.
4. Close the large vacuum valve tightly and check that the small vacuum valve and the chamber isolation valve are both closed tightly.
5. Turn off all of the equipment.
6. Close the cylinder valves.
7. Make sure that the door to the gas cabinet is closed tightly.

Table B.1: Gas flow rates for reduction and growth

Gas	Reduction	Growth	Notes
Argon	642 sccm	642 sccm	
Hydrogen	88 sccm	66 sccm	
Acetylene	0 sccm	28 sccm	Keep regulator pressure under 15 psi when gas is flowing

Appendix C

Mass Flow Controller System

The mass flow controller system is responsible for dispensing a regulated flow of gas into the growth chamber during reduction and growth procedures.

C.1 Mass Flow Controller Calibration

At the beginning of this project all of the mass flow controllers (MFCs) used in the system were replaced, as the existing ones were worn out and not correctly calibrated for the gases they were being used with. However, a substantial amount of data had been taken with the equipment while the old MFCs were installed, so it was desirable to cross-calibrate the new MFCs against the old MFCs. To accomplish this two techniques were used. The first was to use the new MFCs in metering mode to measure the flow rates of the old MFCs for several flow rates. However, this does not verify the calibration of the new MFCs. To accomplish that, a bubble test was performed as shown in Figure C.1.

A bubble test begins by filling a tub with water. A graduated cylinder is then filled and placed entirely under the water in the bucket before being inverted and brought to the surface such that the top of the cylinder is just at the level of the water in the tub but not so high that the vacuum is broken and the water pours out of the cylinder. A tube, connected to the output of the mass flow controller under test is then inserted into the underside of the graduated cylinder such that the end of the tube is just level with the water surface. The tube end is placed level with the water surface to assure that the back pressure from the water is equal to atmospheric pressure. To test the MFC with negative pressure on its output, the tube is inserted further into the cylinder, while to test with positive pressure, the tube is lowered into the tub of water. With the system setup, the MFC is turned on and allowed to run for a fixed amount of time before being shut off again. This causes bubbles to enter the graduated cylinder, displacing some of the water. After the MFC is shut off, the graduated cylinder is lowered into the water such that the water level inside the cylinder is exactly level with the water in the tub. This is to equalize the pressure inside the cylinder to that of atmospheric pressure. The amount of gas in the cylinder can then be read. This amount of gas, measured in cubic centimeters, divided by the time, in minutes, that the MFC was allowed to run is the flow rate of the MFC in sccm (standard cubic centimeters per minute.) This assumes, of course, that the water is at room temperature and that the test is performed at sea level.

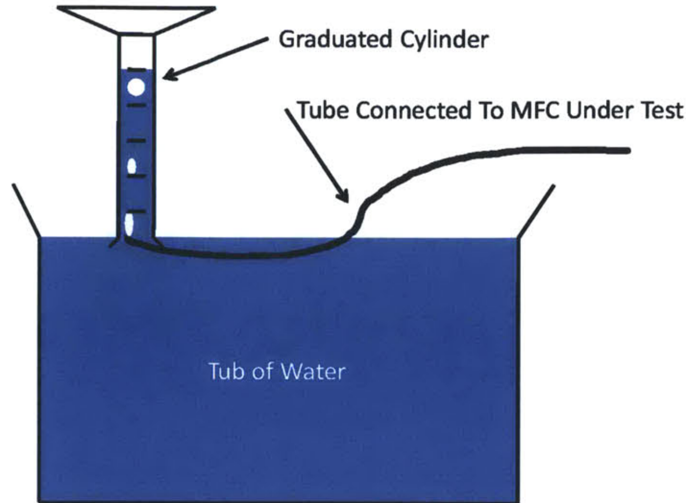

Figure C.1: Diagram of the setup required for a bubble test.

Per the above procedures, each of the mass flow controllers (for each of hydrogen, argon, and acetylene) were tested, and the flow rates measured for the new mass flow controllers were plotted against the flow rates registered by the old mass flow controllers. The resulting plots can be seen in Figure C.1.

As can be seen in Figure C.1, the old mass flow controllers proved to be quite linear, but exhibited significant offset and calibration (slope) errors as compared with the new mass flow controllers. Furthermore, the new mass flow controllers agreed with the results from their bubble tests. Thus, the equations of the lines fitted to these plots can be used to convert flow rates taken with the old mass flow controllers to the corresponding flow rates on the new, correctly calibrated mass flow controllers. Equations C.1, C.2, and C.3 relate the old flow rates to the new.

$$\text{New Flow Rate} = 0.24 * \text{Old Flow Rate} - 0.077 \tag{C.1}$$

$$\text{New Flow Rate} = 1.03 * \text{Old Flow Rate} + 13.063 \tag{C.2}$$

$$\text{New Flow Rate} = 0.44 * \text{Old Flow Rate} - 0.87 \tag{C.3}$$

C.2 Documentation for MIT / LEES Mass Flow Controller Interface Box

The Mass Flow Controller Interface Box was custom built for this project and is responsible for interfacing the mass flow controllers with the National Instruments LabView based software that controls the experiment. It should not require any maintenance, but in case of trouble, the schematic is shown in Figure C.4. The label on the back of the unit that describes the connector pinout, wire colors, and wire functions is reproduced

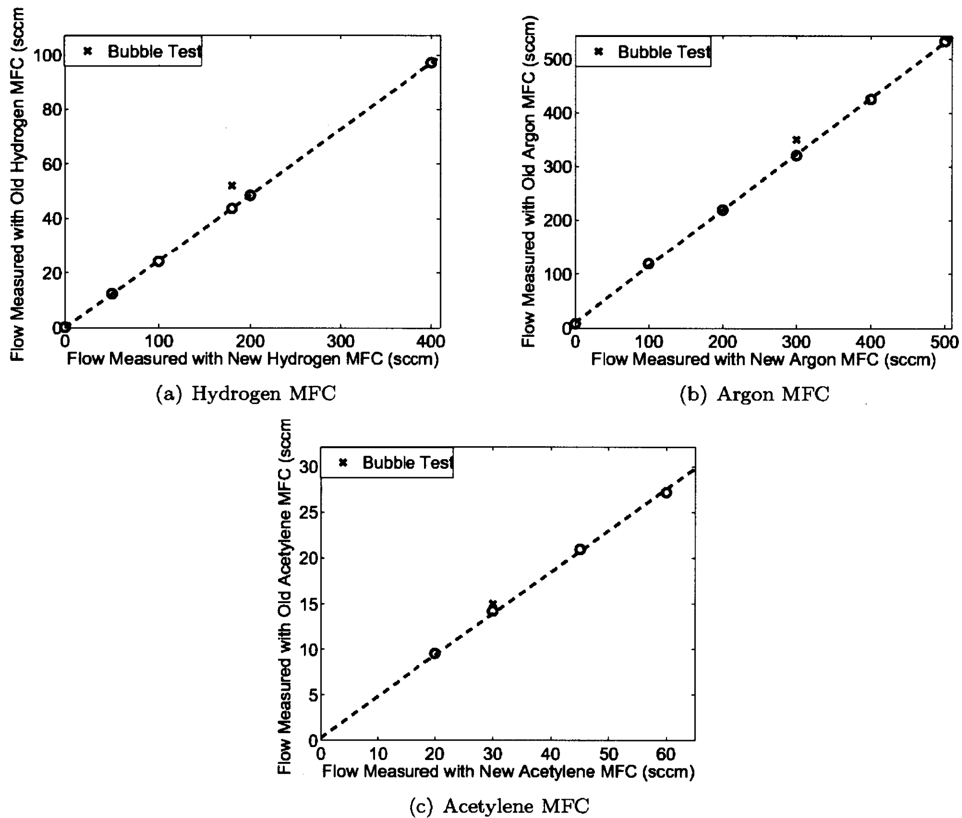

(a) Hydrogen MFC

(b) Argon MFC

(c) Acetylene MFC

Figure C.2: Flow rates of the new mass flow controllers plotted against the flow rates registered by the old mass flow controllers.

below in Figure C.3.

Mass Flow Controller Cable Description								
Function	Gas 1		Gas 2		Gas 3		Gas 4	
	Pin Number	Wire Color	Pin Number	Wire Color	Pin Number	Wire Color	Pin Number	Wire Color
MFC Flow Output	1	Purple – Black	10	Brown	19	Red – Black	28	Blue – White
Open MFC Valve	2	Brown – Green	11	Pink	20	Green – Red	29	Green – White
Close MFC Valve	3	Brown – Black	12	Yellow	21	White – Red	30	Purple – White
Power Common	4	Gray – Green	13	Blue	22	Orange – Red	31	Yellow – White
-15V	5	Purple – Green	14	Green	23	Yellow – Black	32	Red – White
+15V	6	Orange – Black	15	Tan	24	Brown – Red	33	Brown – White
MFC Flow Input	7	White – Black	16	Purple	25	Blue – Black	34	Orange – White
Signal Common	8	Orange – Green	17	Orange	26	Yellow – Red	35	Blue – Red
Chassis Ground	9	Red – Green	18	Black	27	Green – Black	36	White – Green

Figure C.3: Reproduction of the pinout label on the back of the MIT / LEES mass flow controller interface box.

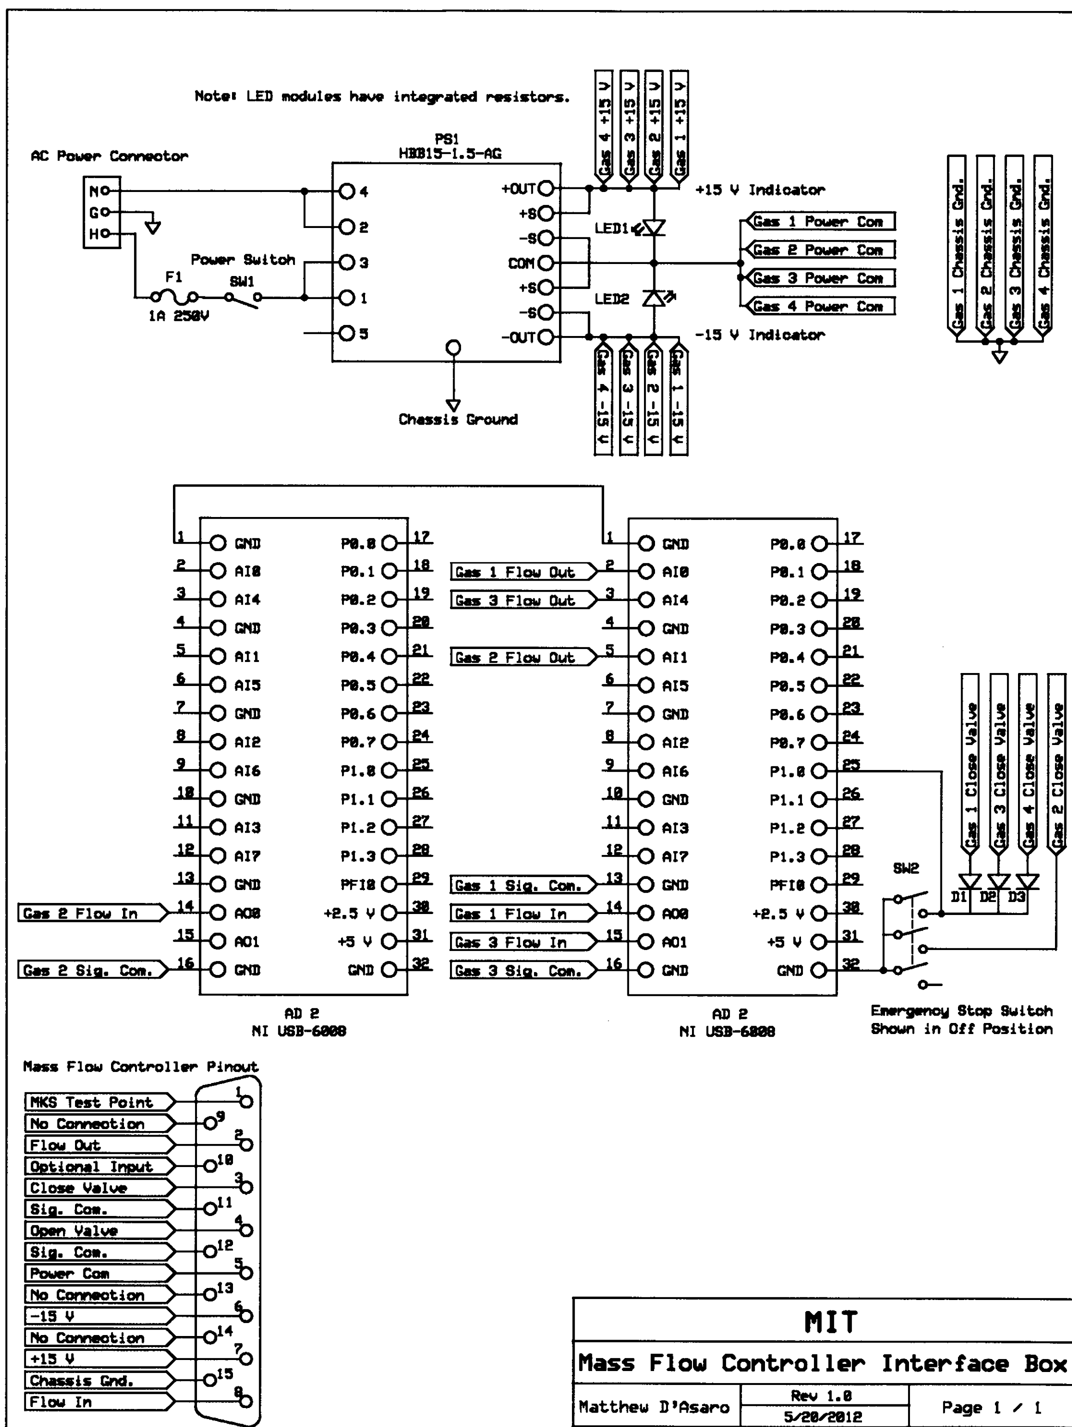

Figure C.4: Schematic of the MIT / LEES mass flow controller interface box. Note that in the present setup, gas 1 is argon, gas 2 is acetylene, gas 3 is hydrogen, and gas 4 is not connected.

Appendix D

Heater System

This appendix describes the heater system, including use of the heater control unit, maintenance procedures for the heater itself, and instructions for replacing the quartz tube and the silicon heater element.

D.1 Growth Chamber

The growth chamber, shown in cross-section in Figure D.1, consists of a quartz tube sealed at either end with o-rings. The o-rings are held against the glass and the end of the vacuum fittings (which are not shown in Figure D.1) by compression nuts. Compression rings inside the compression nuts force the o-rings against the quartz tube. To protect the ends of the tube, rubber cushion washers (square rings) were installed which, in turn, rest against stainless support washers inside the vacuum fittings.

Maintenance of this system is straightforward and consists of occasionally disassembling the entire system, cleaning all of the parts, replacing the o-rings if needed, and re-lubricating with Dow Corning vacuum grease. Note however that only the o-rings and threads on the compression nut should be lubricated. Vacuum grease should be applied sparingly as excess attracts dirt and may contaminate the system. This procedure should be performed whenever trouble is had attaining a 10 mtorr vacuum or when excess debris accumulates in the chamber.

Note that despite the cushion washers, the end of the quartz tube will become chipped over repeated use and it will eventually need to be replaced. For this reason, care should be taken to insert the tube into the vacuum fitting carefully and fully - it may catch before seating completely, causing chipping when vacuum is applied. In the case that the chamber does crack, it is easily replaced. The tubing is standard 2" quartz laboratory tubing and can be cut to length with a diamond saw or a waterjet cutter.

D.2 Heater Support and Electrical Interface

The heater assembly rests inside the growth chamber, supporting the heater horizontally and providing electrical connections to it. It consists of a ceramic substrate onto which stainless blocks (the lower heater clamps) are mounted with hex machine screws driven from below. Under the heads of these screws are fastened the wire terminals. Two wires are connected to each end of the heater. One pair provides electrical

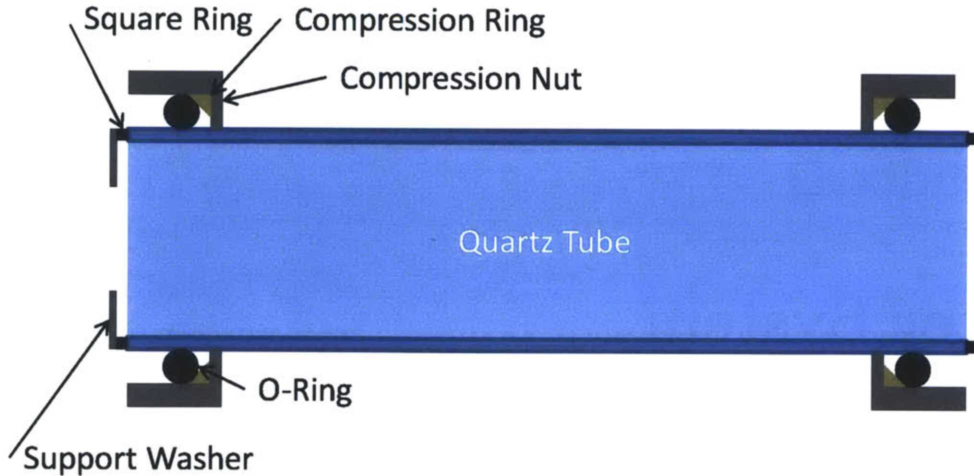

Figure D.1: Cross section of the growth chamber without heater support installed. Note that the vacuum fittings that this assembly mounts to are not shown (not to scale).

power to the heater, while the other pair is used by the heater control unit to monitor the voltage across the heater element. On top of this first set of blocks is a second set, the upper heater clamps, with the heater between them. A second set of hex machine screws fastens the lower and upper heater clamps together, providing electrical connection and mechanical support to the heater. Under the heads of this second set of hex machine screws are split lock washers intended to provide a spring with which to hold the heater between the clamps without breaking it. Slid over the ends of the machine screws mounting the lower heater clamps are a pair of flat washers. These were added to assure uniform pressure on the end of the heater by the upper heater clamps. Without them the heater is much more prone to breakage. Figure D.2 is a cross-sectional drawing of the heater assembly.

D.3 Replacing the Heater

The heater, due to thermal stresses, will eventually crack and require replacement. When this occurs, the following procedure should be used to install a replacement.

1. Unplug the heater control unit.
2. Allow the heater assembly to cool for at least an hour after power has been shut off to avoid the risk of burns.
3. Loosen the left-hand compression nut on the growth chamber until it disengages from its threads, then slide the entire growth chamber to the right until the heater assembly is completely exposed.
4. Loosen and remove the heater mounting screws. Be sure not to lose the lock washers. Then remove the upper heater clamps.
5. Remove and discard the remains of the old heater. Be sure that the flat washers (now visible on top of the lower heater clamps) remain in place.

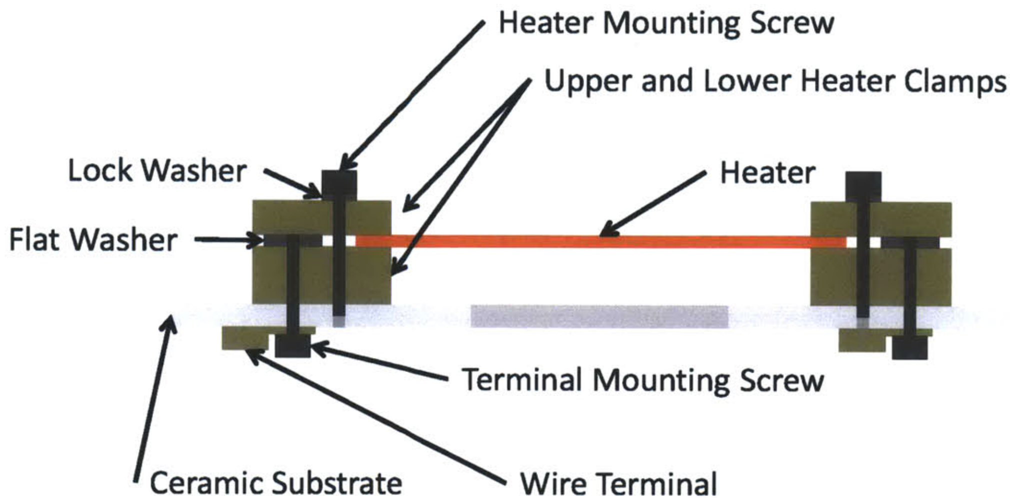

Figure D.2: Cross section of the heater assembly. Note that connecting wires are not shown (not to scale).

6. Using a diamond scribe, cut a 10-cm wide strip from a polished, highly doped silicon wafer. The exact thickness and conductivity are unimportant so long as the resulting heater has a resistance on the order of 1 Ω. Then cut this strip to a length such that when set on the lower heater clamps each end will be about 1 - 2 mm from the holes for the heater mounting screws. The exact length depends on the alignment of the lower heater clamps with respect to the ceramic substrate.
7. Check for excessive corrosion on the upper and lower heater clamps where they meet the heater surface. Any excessive corrosion should be removed with fine sand paper.
8. Place the heater on the lower heater clamps. Assure that it is exactly centered and that the flat washers are in place.
9. Install the heater mounting screws. DO NOT OVERTIGHTEN. They should be tightened just enough to fully compress the split lock washers and no more.
10. Slide the growth chamber closed, tighten the compression nut, and pull a vacuum.
11. Plug in the heater control unit and turn on the heater to 800 °C. The heater should glow red. The peak intensity of this glow should be in the center of the heater. If it is not, turn off and unplug the heater control unit. Allow the system to cool. Open the growth chamber as described in Step 1. Loosen the heater mounting screws and slide the heater slightly in the direction that the glow was stronger, then re-tighten the heater mounting screws. Recheck for even glow and repeat until the glow is exactly in the center of the heater.
12. Normal operations may now resume.

D.4 Heater Control

The attached pages document the use, calibration, and internal operation of the heater control unit designed by David Otten.

Instruction Manual
CNT Heater Control

Written by David Otten
Started January 12, 2012
Revised May 22, 2012

Table of Contents

Introduction........ 89
Connection Instructions........ 91
  Sensor Connections........ 91
  Power Supply and Enable Input Connections........ 92
  Heater Connections........ 93
  Serial Interface Connections........ 94
Manual Interface using LCD and Keyboard........ 95
  Reset........ 96
  Default Display........ 96
  Top Menu........ 97
  Setpoint........ 97
  Calibration........ 98
  Control........ 100
  Emergency Stop........ 100
Computer Interface using Serial Port........ 101
  [Computer control........ 101
  0 Zero heater PWM........ 101
  + Increment Heater PWM........ 101
  - Decrement Heater PWM........ 102
  D Disable Temperature Feedback Control........ 102
  E Enable Temperature Feedback Control........ 102
  L List commands........ 102
  P Pulse Width........ 102
  S Set temperature setpoint........ 103
  V Version of firmware........ 103
  W Set diagnostic counter........ 103
  X Enable diagnostics........ 103
  Y Disable diagnostics........ 104
Hardware........ 104
Performance........ 105
Schematic........ 107
End of Manual........ 111

Introduction

The carbon nano tube (CNT) heater control is a stand alone device to control the temperature of the heater in a CNT furnace based on the signal from an infrared thermocouple (IRt/c). The heater is a rectangular piece of doped silicon through which a

current of approximately 9 amps is passed to heat it up to a temperature of 600 to 900 °C. Because of the somewhat metallic surface of the heater, a low emissivity (LoE) IRt/c is used to sense the heater temperature.

Features:

- Temperature can be set from 200 °C to 990 °C with a resolution of 1 °C
- Simple menu-driven user interface to set all parameters
- Manual and/or computer interface to heater control
- Dedicated blinking red LED to indicate heater control is active
- Dedicated green LED to indicate heater control is within 1.0 degree of setpoint
- High speed operation – temperature typically settles to within 1 °C of setpoint within 1 second
- Separate optically isolated enable input can be used to shut off heater current in an emergency
- Isolated high speed serial interface to computer (9600 – 115200 baud)
- Dedicated 4 line x 20 character LCD display to show temperature setpoint and measured temperature.
- 24-bit ΔΣ A/D converter sampling at 55 Hz to monitor temperature and other signals
- 10-bit 20 kHz PWM amplifier with onboard LC filter to supply current to heater.
- 16-bit general purpose analog output is available for monitoring or control

The pictures below show the front and back of the heater control printed circuit board.

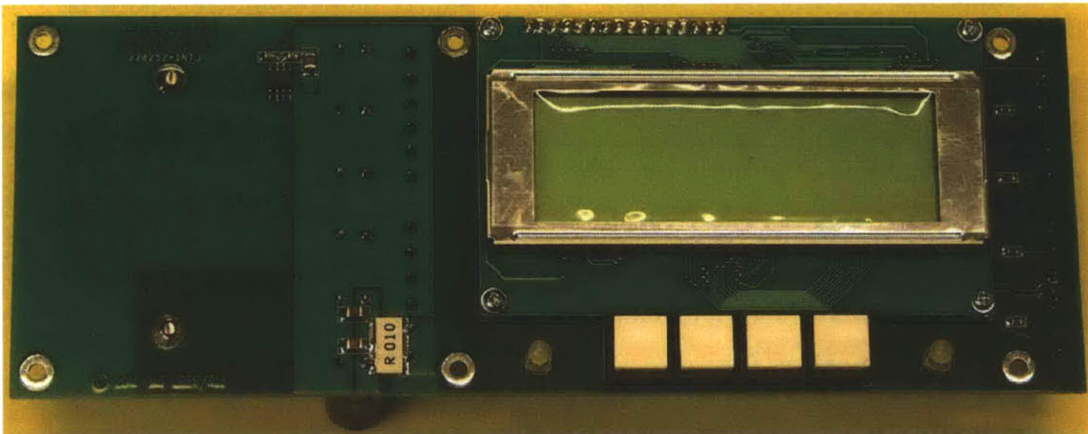

Front View of Controller with Display and Menu-selection Buttons

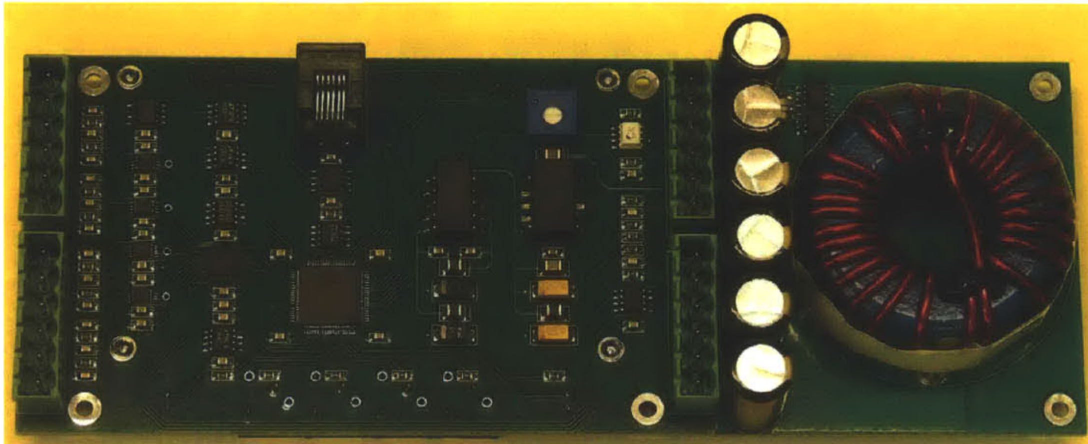

Rear View of Controller with Connectors and Circuitry

Connection Instructions

This section of the manual deals with the various connections to the Heater Controller.

Sensor Connections

The heater controller has inputs for up to 4 external sensors. Two of these can be Type J thermocouples and two can be any signal source. The thermocouples are connected to amplifiers with integral cold junction compensation in hardware. The arbitrary signals are connected to the inputs of programmable gain amplifiers (PGA) with differential inputs. Gains of 1, 10, 100, and 1000 can be selected by the uP. The temperature of the PGA inputs is monitored to allow cold junction compensation in software if the signals come from thermocouples. *It is important to connect the unused inputs of all channels to ground to prevent saturation of the amplifiers and overloading of the A/D converter inputs.*

A digital output is included to control the power to a pair of low power lasers to help identify the sensing position of the IRt/c sensors on the CNT heaters. This function is not currently implemented in the software.

An analog output is included to allow for analog monitoring of the temperature sensors or controller operation. This function is not currently implemented in the software.

The sensor input connections are located on the left edge of the rear of the controller. The connectors are shown in the photograph below. The top pin of each connector is pin 1.

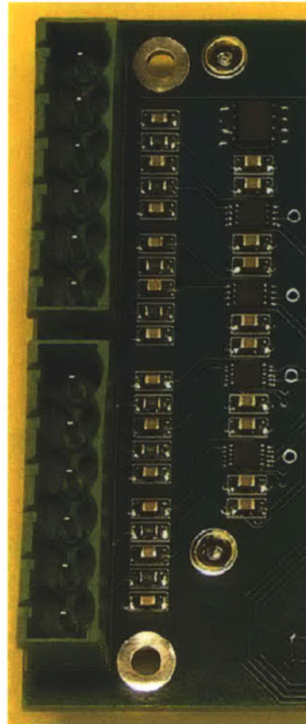

Sensor Connectors

The table below shows the signals on each pin.

Sensor Connector Signals	
Pin Number	**Signal**
Top 1	Thermocouple 1 Input +
Top 2	Thermocouple 1 Input -
Top 3	Laser Power
Top 4	Ground
Top 5	Thermocouple 2 Input +
Top 6	Thermocouple 2 Input -
Bottom 1	PGA 1 Input +
Bottom 2	PGA 1 Input -
Bottom 3	D/A Ouput
Bottom 4	Ground
Bottom 5	PGA 2 Input +
Bottom 6	PGA 2 Input -

The software is configured to have the LoE IRt/c signals connected to PGA 1 and the HiE IRt/c signals connected to PGA 2. The software currently only uses the LoE IRt/c sensor for control but monitors the HiE IRt/c signal.

Power Supply and Enable Input Connections

The heater controller operates from a 12 – 24 volt DC power supply. The heater may require currents in excess of 10 amps so two pins on the connector are used for the positive supply and two pins are used for the negative supply.

Provision is also made for an external enable signal. This signal is optically isolated from the ground of the controller. A DC voltage of 5 – 24 volts applied between pins 1 and 2 of the connector will enable the controller. The polarity of the enable signal is not important.

The power supply and enable input connections are located on the upper right side of the rear of the controller. The connector is shown in the photograph below. The top pin of the connector is pin 1.

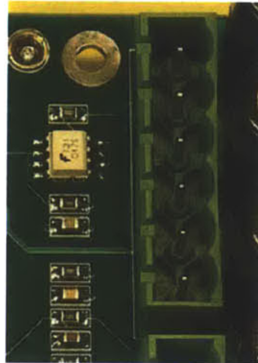

Power Supply Connector

The table below shows the signals on each pin.

Power Supply Connector Signals	
Pin Number	**Signal**
1	Enable Input
2	Enable Input
3	+ Supply
4	+ Supply
5	- Supply
6	- Supply

If no enable signal is provided, the LCD on the controller will display the EMERGENCY STOP message.

Heater Connections

The heater controller includes a high current PWM amplifier. This amplifier may provide currents in excess of 10 amps so two pins are included on the connector to supply current to each end of the heater. An integral LC filter is included in the design to minimize the high frequency voltage and current signals on the heater.

A separate differential heater voltage input is also provided on the heater controller. This allows for a 4-terminal measurement of the heater power and resistance.

The heater connections are located on the lower right side of the rear of the controller. The connector is shown in the photograph below. The top pin of the connector is pin 1.

Heater Connector

The table below shows the signals on each pin.

Heater Connector Signals	
Pin Number	**Signal**
1	Heater Voltage -
2	Heater Voltage +
3	Heater Current +
4	Heater Current +
5	Heater Current -
6	Heater Current -

The heater voltage, power, and resistance are not available on the LCD display but the voltage, current, and power are available on the computer interface which allows the resistance to be calculated and all variables monitored.

Serial Interface Connections

The heater control has a serial interface that connects to a computer. The interface allows computer control and monitoring of the controller. The serial interface connections are located on the top left side of the rear of the controller. The connector is shown in the photograph below. The left pin of the connector is pin 1.

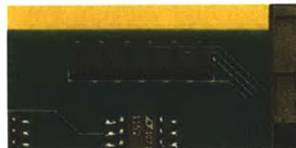

Serial Interface Connector

The table below shows the signals on each pin.

Serial Interface Connector Signals		
Pin Number	**Signal**	**Color**
1	GND	black
2	CTS	brown
3	+5V	red

4	TXD	orange
5	RXD	yellow
6	RTS	green

The serial interface follows the conventions of an RS-232 interface except for the voltage levels. Instead of traditional RS-232 voltage levels (+/-12 volts), only logic levels (0 and 5 volts) are available. A special USB cable is supplied with the transmitter that works with these logic levels and allows the transmitter to be connected to the computer using a USB port. (Today USB ports are generally more available on laptops and other computers than traditional serial ports.)

A special driver is required for the USB to TTL Serial Cable from FTDI Chip. This is available on their web site at http://www.ftdichip.com/Drivers/VCP.htm. Once the driver is installed, HyperTerminal, a terminal emulation program available on Windows operating systems, may be used to test the interface.

The baud rate may be set to 9600, 19200, 57600, or 115200 via the LCD and menu system described below.

Manual Interface using LCD and Keyboard

All the options of the heater control can be selected from the manual user interface shown in the photograph below.

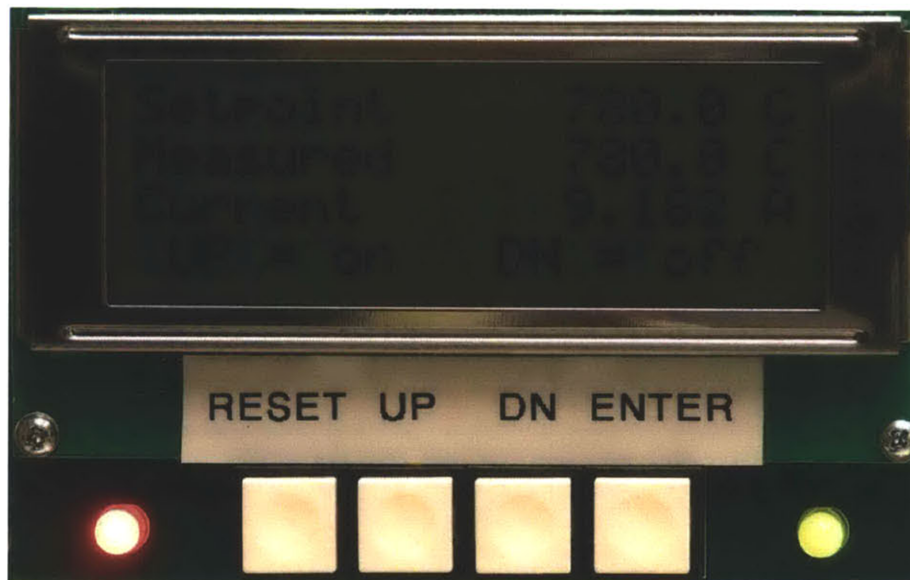

Normal Display showing Setpoint and Measured Temperature, and Heater Current

The red LED on the left shows that temperature control is active. The green LED on the right shows that the measured temperature is within 1 °C of the setpoint.

Diagnostic Display showing IRt/c mV Signal, equivalent Temperature, Setpoint Temperatue, equivalent mV Signal, Reference (ambient) Temperature, Heater Voltage, Heater Current, and Amplifier Pulse Width

Menus are also displayed on the 4-line LCD display and 4 buttons, labeled RESET, UP, DN, and ENTER are used to select options from the menu. The function of each button and all menus is described below.

Reset

When the unit is first turned on, the version number of the software will be displayed for 2 seconds.

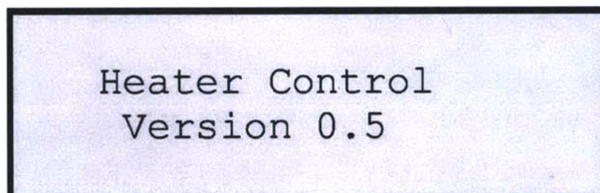

At any other time, while the RESET button is pressed, the same information will be displayed.

Default Display

Following the initial power-on or after the RESET button is released, the following information is available on the LCD.

```
Setpoint      780.0 C
Measured      780.0 C
Current       9.182 A
 UP = on    DN = off
```

The first line is the setpoint temperature in °C. The second line is the measured temperature in °C. The third line is the heater current in amps. The fourth line is a prompt to instruct the user on how to turn the feedback on and off using the UP and DN switches below the display. To exit the default display and enter the menu system, the ENTER button must be pressed. Unfortunately there is no room on the display for this information without leaving something else off.

If the diagnostic mode is enabled from the computer, the default display is changed to include the following information.

```
LoE  13898 RT  26.83
LoE  780.0 HV  6.760
Set    780 HI  9.117
mV   13897 PW    737
```

The left signal on the first line is the raw mV signal from the LoE IRt/c sensor. The left signal on the second line is the equivalent temperature from the LoE IRt/c sensor. The left signal on the third line is the setpoint temperature in °C. The left signal on the fourth line is the equivalent mV signal corresponding to the setpoint temperature. The right signal on the first line is the reference or ambient temperature in °C. The right signal on the second line is the heater voltage in volts. The right signal on the third signal is the heater current in amps. The right signal on the fourth line is the output amplifier pulse width.

Top Menu

When the ENTER button is pressed from the default display, the top menu is displayed on the LCD screen. This is shown below.

```
     PARAMETERS
█Setpoint
 Calibration
 Control
```

The black block to the left of the S in Setpoint will be a blinking cursor. The position of the cursor can be moved up or down with the UP or DN buttons. Once the cursor is on the line of the desired parameter, the ENTER button should be pressed. At any point the RESET button may be pressed to return to the default display without changing the current parameter. Each of the sub menus below the top menu will be discussed in turn.

Setpoint

The setpoint temperature is set with this menu. The display will initially contain the current setpoint.

```
      SETPOINT
      980 deg C
```

```
UP = inc, DN = dec
ENTER = next digit
```

The blinking cursor will initially be located over the first digit of the setpoint temperature. This digit can be increased with the UP button or decreased with the DN button or remain at its current value. Each digit must be between 0 and 9 inclusive. When the current digit is correct, the cursor can be advanced to the next digit with the ENTER button. There is no provision for moving the cursor back to a previous digit. If this is required, the RESET button should be pressed and the process started over from the default display. After the last digit has been confirmed with the ENTER command, the new value will be checked to make sure it is within range. The range for the setpoint temperature is 0 – 990 °C. If it is out of range, it will be modified to be within range. The setpoint temperature is then updated with the new value and the program returns to the top level and displays the default menu.

Calibration

A series of menus are used to set the calibration constants for both the LoE and HiE sensors. At each level, the UP or DN button can be used to select the next level and then the ENTER button can be used to move to that level.

The first level selects whether parameters for the LoE or HiE sensor will be selected.

```
      CALIBRATION
LoE
HiE
```

If the LoE sensor is selected, the next level selects whether the Offset or Gain parameter will be selected.

```
  LoE PARAMETERS
Offset
Gain
```

If the LoE Offset parameter is selected, the following screen will be displayed.

```
      LoE OFFSET
        -0.0057
UP = inc, DN = dec
ENTER = next digit
```

The blinking cursor is initially placed over the first digit of the offset. The UP and DN buttons can be used to modify that digit. Each digit must be between 0 and 9. When the

ENTER button is pressed, the cursor advances to the next digit. There is no provision to go back to a previous digit. If this is required, the RESET button should be pressed, and the process repeated from the beginning. After the last digit is confirmed with the ENTER command, the LoE Offset is updated with the new value and the program returns to the top level and displays the default menu.

If the LoE Gain parameter is selected, the following screen will be displayed.

```
    LoE GAIN
     1.0000
UP = inc, DN = dec
ENTER = next digit
```

The UP, DN, and ENTER buttons are used to modify the LoE Gain in the same way as the other parameters, the new value is saved, and the program returns to the top level and displays the default menu.

A similar menu tree exists for the HiE Parameters. This level selects whether the Offset or Gain parameter will be selected.

```
  HiE PARAMETERS
Offset
Gain
```

If the HiE Offset parameter was selected, the following screen will be displayed.

```
    HiE OFFSET
     -0.0032
UP = inc, DN = dec
ENTER = next digit
```

The UP, DN, and ENTER buttons are used to modify the HiE Offset in the same way as the other parameters, the new value is saved, and the program returns to the top level and displays the default menu.

If the HiE Gain parameter was selected, the following screen will be displayed.

```
    HiE GAIN
     1,0000
UP = inc, DN = dec
ENTER = next digit
```

The UP, DN, and ENTER buttons are used to modify the HiE Gain in the same way as the other parameters, the new value is saved, and the program returns to the top level and displays the default menu.

Control

Two menus are used to set up the interface to the heater control. For both menus, the initial position of the cursor is on the currently selected entry. This entry is also identified with a dot in the first column. The UP and DN buttons can be used to select a different entry. The new entry is saved as the current entry when the ENTER button is pressed.

The first menu selects the type of control available on the serial interface. In manual mode, all the commands listed in the next section (Computer Interface using Serial Port) of this manual are available. In computer mode, only the first command is available. The other commands are disabled because they are designed for a human operator and either contain syntax not required by the computer, or are too powerful and would allow a single character from the computer to significantly change the operation of the controller.

```
      CONTROL
•Manual
 Computer
```

The second menu selects the baud rate of the serial interface. The new baud rate should be available immediately after it is set.

```
 9600    BAUD RATE
 19200
 57600
•115200
```

Emergency Stop

The heater control is provided with an independent optically isolated enable input that can be used to allow controller operation or shut off the heater current in an emergency. A voltage signal of 5 to 12 volts DC is required to enable the heater. If this signal is not present, the following display is provided. The expectation is that this input would be used with an emergency stop button to coordinate all shut down functions of the CNT furnace.

```
      EMERGENCY
        STOP
```

Computer Interface using Serial Port

The screen dump below shows a typical display when the transmitter is first turned on and the serial interface is connected.

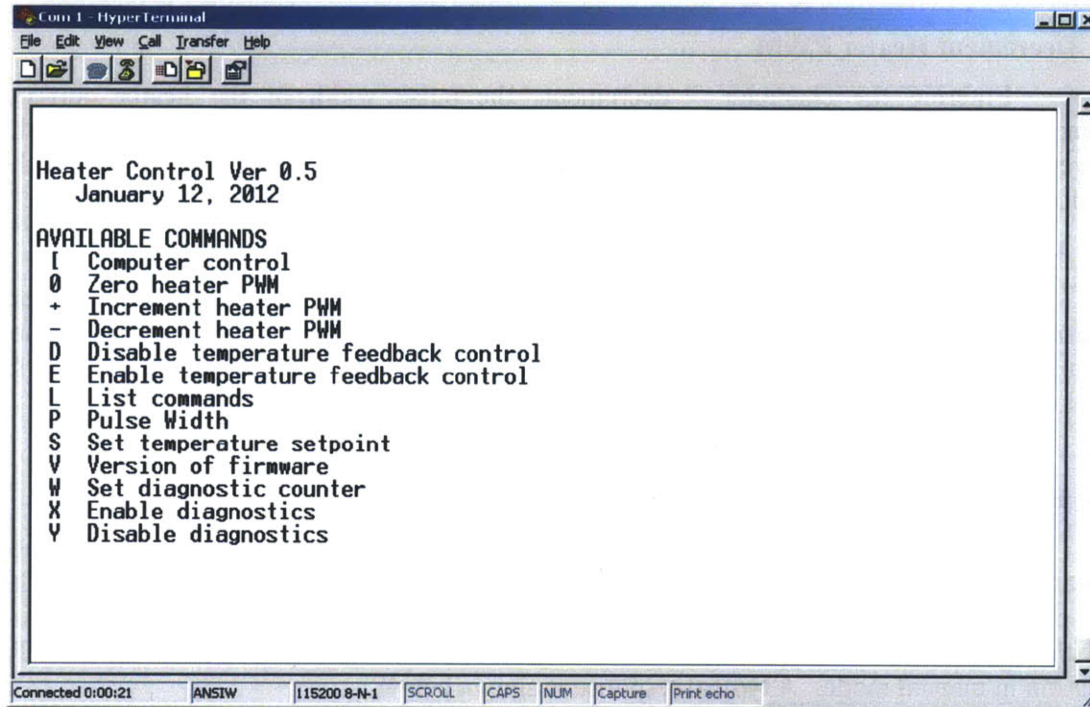

The first two lines show the software version number and date. This is followed by a list of the available commands. All commands can be entered with a lower case or upper case letter. Each of the commands will be explained in turn.

[Computer control

Entering the character [signals the start of a computer command. Three commands are supported. [E] is used to enable the temperature feedback control. [D] is used to disable the control. [xxx] is used to set or update the setpoint temperature where the xxx represents any number between 0 and 990. The syntax of this command is very specific which allows the heater controller to check it carefully and ignore all commands that do not have the correct format. This makes the protocol robust against noise and helps to prevent communication errors. There is no direct response from this command but the flag and setpoint information is transmitted back to the computer when the diagnostic display is enabled.

0 Zero heater PWM

Entering the number 0 sets the pulse width of the heater current amplifier to zero. There is no direct response from this command, but the pulse width information is transmitted back to the computer when the diagnostic display is enabled.

+ Increment Heater PWM

Entering the character + increments the pulse width of the heater current amplifier. If the pulse width is incremented from its maximum value of 639, it will roll over to 0. There is no direct response from this command, but the pulse width information is transmitted back to the computer when the diagnostic display is enabled.

- Decrement Heater PWM

Entering the character – decrements the pulse width of the heater current amplifier. If the pulse width is decremented from its minimum value of 0, it will roll over to 639. There is no direct response from this command, but the pulse width information is transmitted back to the computer when the diagnostic display is enabled.

D Disable Temperature Feedback Control

Entering the letter D disables the feedback control. The firmware sets the heater current to zero, sets the left LED to green, and turns off the right LED. This is the equivalent to pressing the DN button in manual mode. A typical response is shown below.

```
Disable temperature feedback control
```

E Enable Temperature Feedback Control

Entering the letter E enables the feedback control. The firmware controls the heater current to match the measured heater temperature to the setpoint temperature. It also turns on the left LED to blinking red and the right LED to green if the temperature is within 1 deg C of the setpoint or red if it is not. This is equivalent to pressing the UP button in manual mode. A typical response is shown below.

```
Enable temperature feedback control
```

L List commands

Entering the letter L displays a list of the available commands. A typical response is shown below.

```
AVAILABLE COMMANDS
  [  Computer control
  0  Zero heater PWM
  +  Increment heater PWM
  -  Decrement heater PWM
  D  Disable temperature feedback control
  E  Enable temperature feedback control
  L  List commands
  P  Pulse Width
  S  Set temperature setpoint
  V  Version of firmware
  W  Set diagnostic counter
  X  Enable diagnostics
  Y  Disable diagnostics
```

P Pulse Width

Entering the letter P is used to set the pulse width of the heater current amplifier. If the value entered is not within range, the prompt will be issue again. A typical response is shown below.

```
Enter pulse width (0 - 1023) 100
```

S Set temperature setpoint

Entering the letter S is used to specify the setpoint temperature. If the value is higher than the allowed range, it will be reset to the maximum value. A typical response is shown below.

```
Enter setpoint temperature (0-990)  780
```

V Version of firmware

Entering the letter V displays the current version of the firmware and the date it was created. A typical response is shown below.

```
Heater Control Ver 0.5
    January 12, 2012
```

W Set diagnostic counter

Entering the letter W is used to specify the diagnostic counter. The diagnostic counter is used to control the number of lines of diagnostic data that are displayed when diagnostics are enabled. If the counter is set to 1 every line is displayed. The A/D sample rate is 55 Hz and the temperature is measured every other sample so the data rate will be 27.5 Hz. If the counter is set to 2, every other sample is displayed. If the counter is set to 3, every third sample is displayed. If the counter is set to 0, every 256th sample is displayed. A typical response is shown below.

```
Enter diagnostic counter (0-255)  0
```

X Enable diagnostics

Entering the letter X enables the diagnostic display on the LCD and starts transmission of diagnostic data over the serial interface. The data is transmitted in comma separated value (CSV) format suitable for importing into an Excel file. The number of data samples transmitted is controlled by the diagnostic counter. A typical response is shown below.

```
Enable diagnostics
D,780,26.62,54,-12,87.8,0,-0.001,0.000,0.00,-14
D,780,26.63,54,-12,87.8,0,-0.001,-0.001,0.00,-15
D,780,26.63,54,-13,87.8,0,0.000,0.000,0.00,-15
```

The following data is transmitted. The first field is a D or E to indicate if the temperature feedback is enabled or disabled. The second field is the setpoint temperature in °C. The third field is the temperature of the LoE IRt/c connector on the heater controller board in °C. The fourth field is the thermocouple voltage a Type J thermocouple would generate at that temperature if the reference junction was at 25.0 °C. This signal is in A/D counts. The fifth field is the LoE IRt/c signal in A/D counts. The sixth field is the equivalent temperature measured by the LoE IRt/c in °C. The seventh field is the pulse width of the heater current amplifier. The eighth field is the heater

voltage in volts. The ninth field is the heater current in amps. The tenth field is the heater power in watts. The eleventh field is the HiE sensor signal in A/D counts.

Y Disable diagnostics

Entering the letter Y disables the diagnostic display on the LCD and stops the transmission of diagnostic data over the serial interface. A typical response is shown below.

```
Disable diagnostics
```

Hardware

The CNT Heater Controller is built around an Infrared Thermocouple (IRt/c) manufactured by Exergen. When this passive device is pointed at a heated surface, it generates a voltage related to the temperature. The output is also a function of the emissivity of the surface. Sensors for high emissivity (HiE) surfaces (not reflective) and low emissivity (LoE) surfaces (somewhat reflective) are available. The heater in the CNT furnace is a silicon substrate with low emissivity. The sensor selected for this system is the IRt/c.3AMF which has a medium focus of 70 mm and a spot size of 3 mm. This allows the sensor to be mounted below the heater to measure the temperature while the CNT sample is placed on top of the heater.

Input RC filters followed by programmable gain instrumentation amplifiers are included in the design to amplify the input signals. An absolute temperature sensor can measure the temperature of the input connectors if cold junction temperature compensation in software is required. Additionally special thermocouple amplifiers with integral cold junction compensation in hardware are also included.

The signal for the IRt/c is digitized with a 24-bit ΔΣ A/D converter from Linear Technology (LTC2449). The sample rate of the converter can be adjusted to trade off speed vs noise performance in any given application. In the heater controller the A/D samples at 55 Hz. For good feedback performance, the LoE IRt/c is sampled every other conversion and all the other signals are sampled at a lower rate.

All the operations of the heater controller are coordinated by a PIC 18F8722 uP. This 8-bit chip runs at a clock rate of 32MHz using an 8 MHz internal oscillator and an integral 4X PLL. At this frequency it executes most instructions in 125 nsec.

The uP drives a switching amplifier with its onboard PWM generator. The 20 kHz output of the amplifier is filtered with an LC filter to reduce the switching noise before it is sent to the heater. The amplifier routinely supplies currents in excess of 10 amps to the heater. Provision is made to measure the current through the heater and the voltage across it with a 4-terminal measurement.

Also included in the design of the heater controller is an LCD display and switches to provide a user interface. The display is used to show the heater setpoint and measured temperature and to change parameters with a convenient menu system. An isolated serial interface allows the controller to be monitored or controlled from a computer without causing any ground loops.

An important function of the uP is to convert the non-linear output of the IRt/c from mV to °C. This is done with a lookup table. Tables are provided to convert from both mV to °C and vice versa. Tables are also provided for a HiE IRt/c sensor and for type J thermocouples.

For safety, an optically isolated enable input is provided to the controller. The uP turns off the heater current when no enable signal is present.

Though it is not used, a 16-bit D/A convert is designed into the system. This can be used to provide an analog temperature signal or any other diagnostic function if desired.

The heater temperature is controlled with a PID loop calculated by the uP. To simplify the calculations, the control is done with the raw A/D signals and the results are converted from A/D counts to °C outside of the feedback loop. To make this work, the setpoint temperature is also converted to A/D counts when it is first entered into the system or if it is changed.

Performance

The graph below shows the response of the system to a series of setpoint temperatures between 300 and 700 °C.

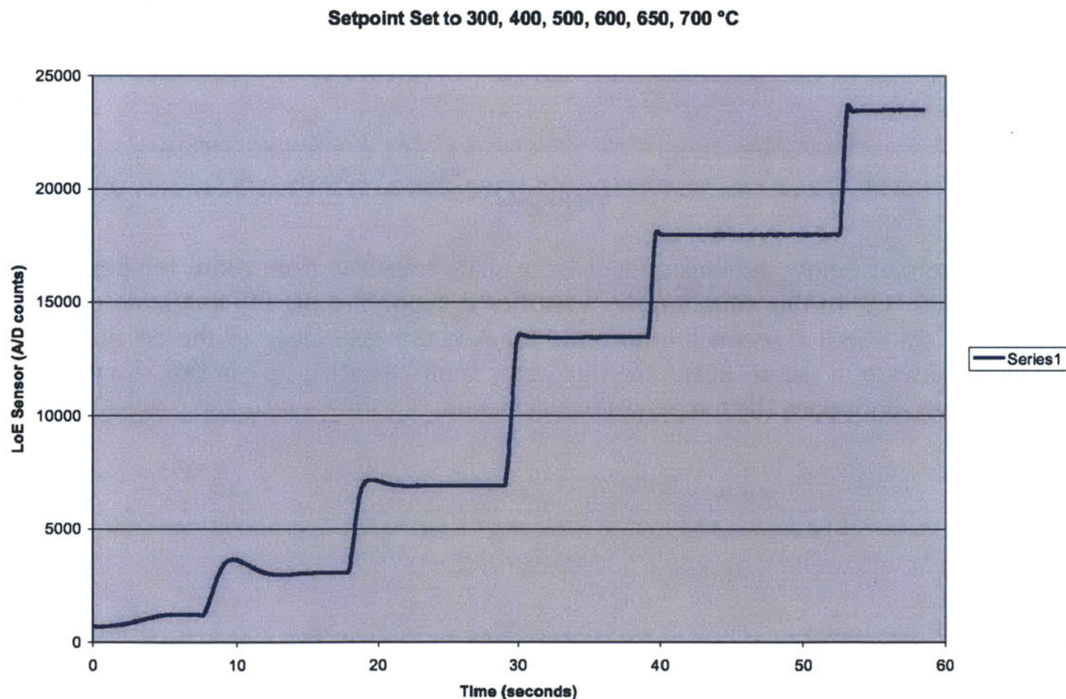

Note that the response changes with temperature. This is because the sensitivity of the sensor changes with temperature. At low temperatures the sensor is not very sensitive (low mV/°C) and the gain around the control loop is low. This results in low bandwidth. At high temperatures the sensor has much higher sensitivity (high mV/°C)

resulting in more gain around the loop and much higher bandwidth. Care is taken to make sure that at the highest temperatures anticipated, the system is still very stable.

The graph below shows a step input from 650 to 700 °C. This demonstrates the response time and stability of the system at the higher temperatures. Note that the signal settles in approximately 1 second and the rise time is much faster than that.

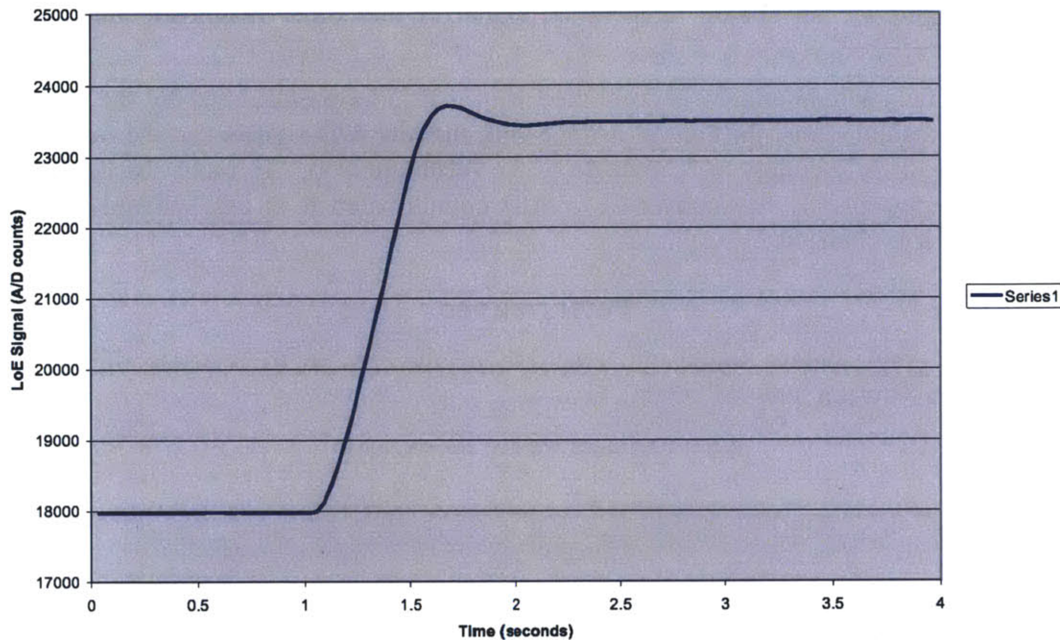

The plot below demonstrates a large scale transient from room temperature to 650 and 700 °C. In this situation the amplifier is supplying the full available power to the heater for about 3 seconds before the temperature gets close to the desired value. Special software in the uP keeps the integrator from "winding up" in this situation and causing overshoot once the temperature is achieved.

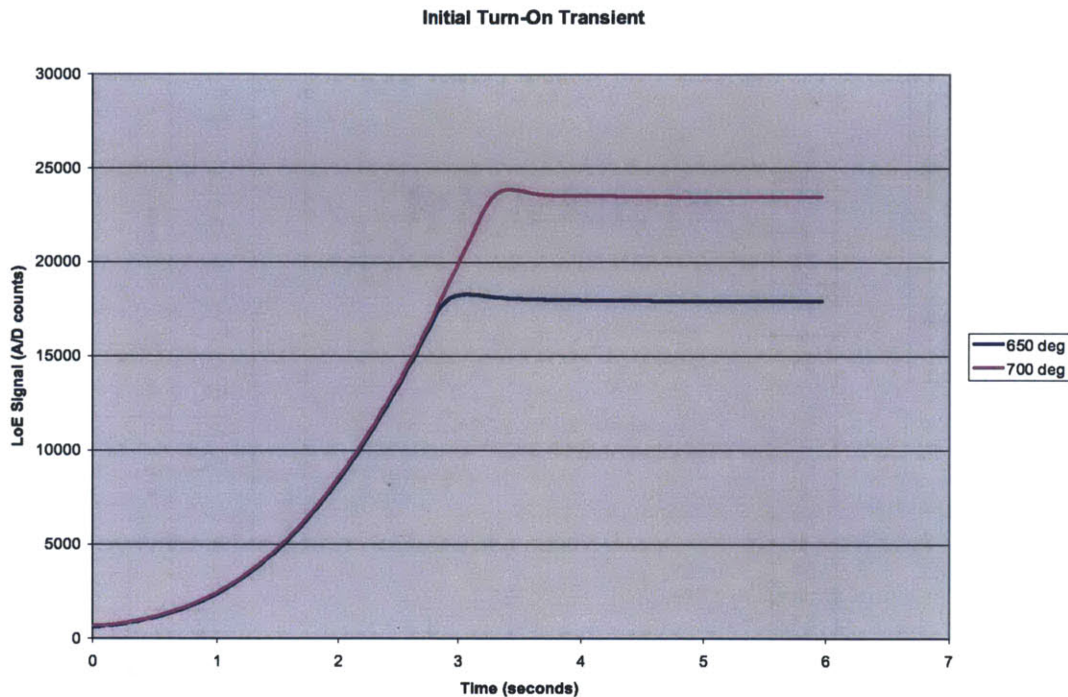

Calibration tests were also done to determine the accuracy of the system. A special wax was purchased from Omega that melts at a specific temperature. This was applied to the heater and the temperature setpoint varied to determine the temperature at which the wax melted. The results are within ±5 °C. Three different waxes were used. The results are shown in the table below.

Wax Melting Point (°C)	Setpoint (°C)	Error (°C)
649	575	74
704	685	19
816	805	11

The results are very encouraging, particularly at the higher temperatures.

Schematic

The following three pages contain the schematic for the heater controller.

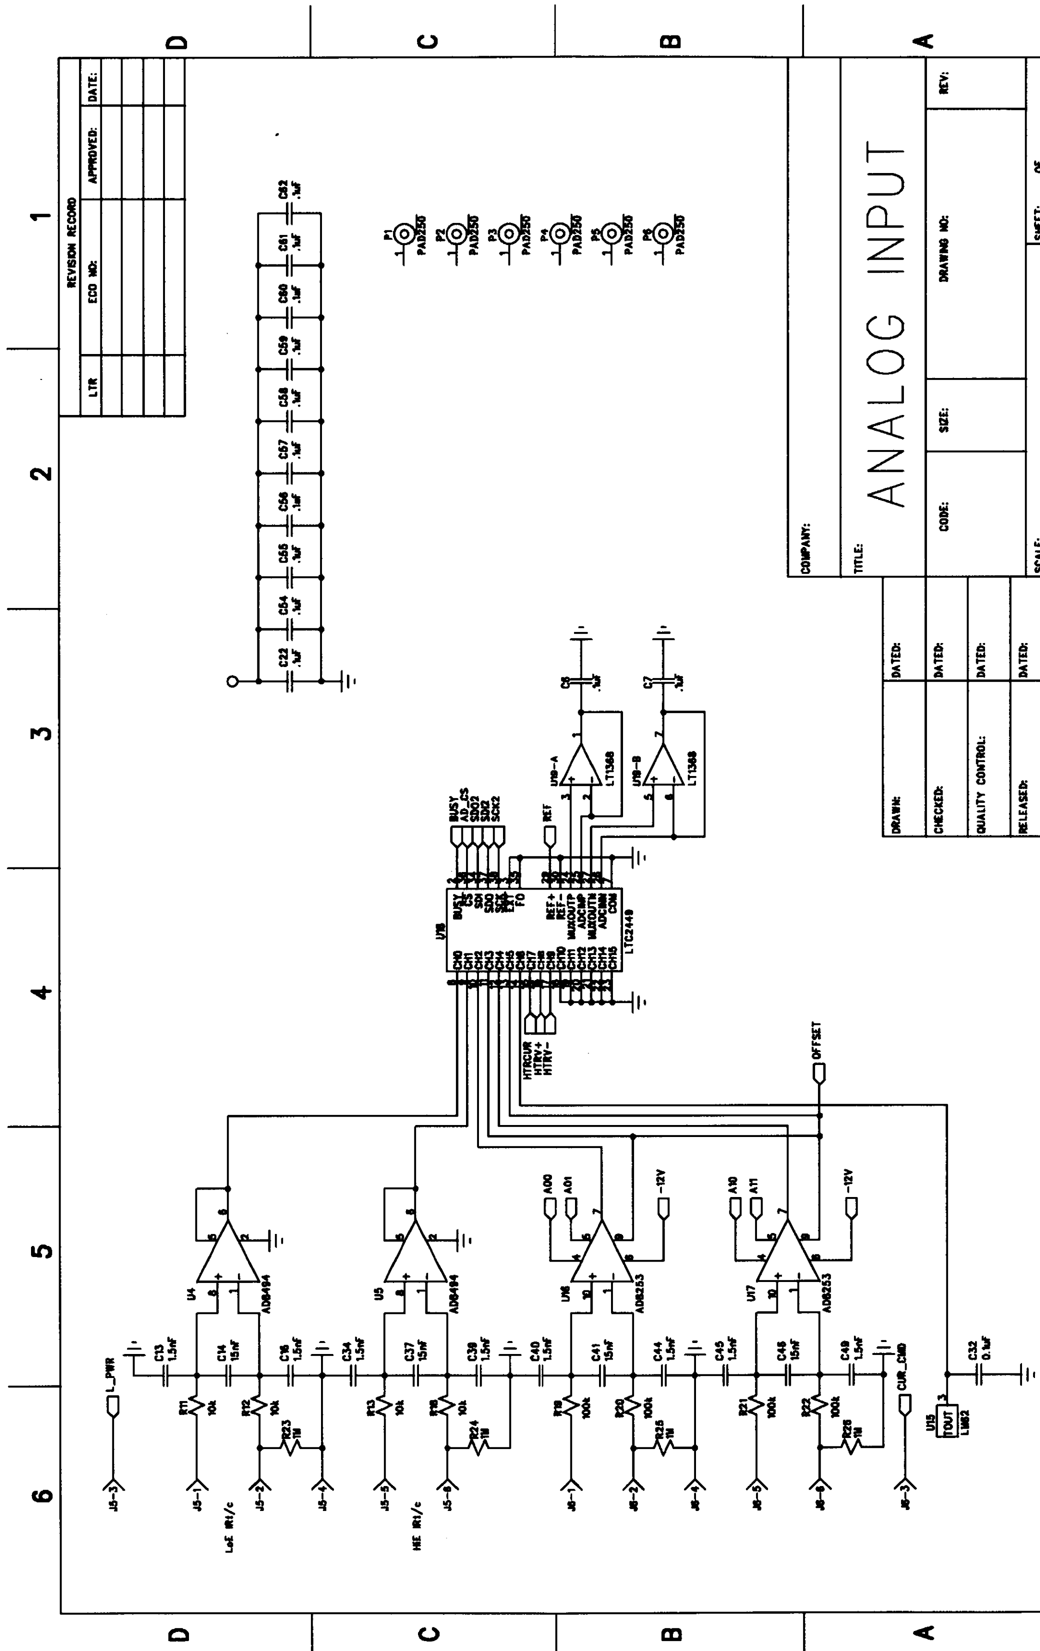
ANALOG INPUT
REVISION RECORD
LTC2440
LT1368
AD8494
AD8253
OFFSET
HTRCUR
HTRV+
HTRV-
L_PWR
CUR_CMD
REF
BUSY
-12V

MICROPROCESSOR

ANALOG OUTPUT

End of Manual