MEMS-BASED FABRICATION OF POWER ELECTRONICS COMPONENTS FOR ADVANCED POWER CONVERTERS

A Thesis
Presented to
The Academic Faculty

by

William Preston Gallé III

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

Georgia Institute of Technology
December 2012

MEMS-BASED FABRICATION OF
POWER ELECTRONICS COMPONENTS
FOR ADVANCED POWER CONVERTERS

Approved by:

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

Dr. Deepak Divan
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Oliver Brand
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Linda Milor
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Suresh Sitaraman
School of Mechanical Engineering
Georgia Institute of Technology

Date Approved: August 21st, 2012

Warm hugs straighten my ear muffs

and got me through metal detector mornings;

and even if never to be seen again,

you're in my air.

-from The Poetry of Ajax

ACKNOWLEDGEMENTS

First and foremost, a superlative debt of gratitude and recognition is owed to my advisor, Dr. Mark G. Allen. His strong support from beginning to end, pragmatic yet open-minded approaches towards everything, and endless patience have shown me that there are very few limits on the ultimate boundaries of achievement.

I would like to than Dr. Deepak Divan for serving as the chairman of my dissertation committee, along with the other committee members, Dr. Oliver Brand, Dr. Linda Milor, and Dr. Suresh Sitaraman. Their attention and willingness to participate in my undertaking is greatly appreciated.

As well, numerous members of the MSMA research group have provided essential education and empowerment to me over the years. These include: Mr. Richad Shafer, an open tome of knowledge and experience of all flavors; Dr. Florian Herrault, Dr. Seong-Hyok Kim, and Dr. Yong-Kyu Yoon, research mentors *par excellence*; and Dr. Jin-Woo Park, Dr. David Arnold, Dr. Chang-Hyeon Ji, Dr. Swaminathan Rajaraman, Dr. Sharon Wu, Mr. Ed Birdsell, and Dr. Seung-Keung Choi, magnanimous collaborators, critics, and teachers.

In a broader scope, many others at Georgia Tech have been continuous sources of assistance, guidance, and perspective. These include Dr. Oliver Brand, Mr. Gary Spinner, Dr. Alan Doolittle, Mrs. Purnima Sharma, Mrs. Laureen Rose, Dr. Muhannad Bakir, and Dr. Amanda Brooks.

As well, I must acknolwedge the extraordinary flexibility and dedication to my success that has been demonstrated by my employer, SiTime Corporation, during the final stages of my graduate studies. In particular, Mr. Pavan Gupta, Mrs. Nivi Arumugam, and Dr. Fari Assaderaghi deserve my sincerest gratitude for their support.

Finally, the people that make everything possible – those closest to me – must be recognized. My dear sister, Julie Baggenstoss, and her lovely family have demonstrated the purest unconditional support throughout this entire experience. My parents, Dr. Bill Gallé and Julia Gallé, have always been my most generous and most essential supporters. And last, but not least, my wife, Houri Johari-Gallé, continues to show me the very meanings of commitment, *joie de vivre*, and kindness.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

Fabrication technology, based on MEMS processes, for constructing components for use in switched-mode power supplies are developed and presented. Capacitors, magnetic cores, and inductors based on sacrificial multilayer electroplating are designed, fabricated, and characterized. A coherent set of microfabrication processes based on sacrificial multilayer electroplating followed by a selective, sacrificial etch is utilized and extended.

Construction of capacitors using this technique leverages its intrinsic ability to produce high surface area structures, creating a path to very high-density capacitors. Its capability of producing very thin layers of magnetic alloys is used to create magnetic cores with reduced eddy current losses.

This technique can enable radical performance improvements in power magnetics devices as well. Although magnetic alloys offer high saturation density and low hysteresis losses, eddy current losses remain as a critical obstacle to the employment of these materials in high-speed power converters. By demonstrating that sacrificial multilayer electroplating offers a reliable method for controlling eddy current losses, this work demonstrates a path to bringing the natural benefits of magnetic alloys to bear in power electronics systems.

Accordingly, highly laminated magnetics cores are fabricated and used to build power inductors. In-depth characterization of the produced inductors is presented to confirm this improvement. As well, the demonstration of the produced inductors in DC/DC converters at high frequencies is presented.

Initial work addressing the top-down development of a fully-integrated DC/DC converter is presented. This work focuses on the integration of an advanced power inductor on a silicon substrate, along with provisions for substrate contact as well as performance optimization.

Supporting results detailing the advancement of the core process of sacrificial multilayer electroplating are presented. These results focus primarily on the development of a second-generation automated multilayer electroplating system, as well as on experiential learning and practical packaging of the microfabricated structures. The combination of these advances allowed for the sacrificial

multilayer electroplating process to be extended to produce structures not only with thinner layers than previous results, but also larger total thicknesses.

The initial application of the advanced sacrificial multilayer plating process to produce microfabricated capacitors is presented. These capacitors, which achieved in excess of 1.5 nF/mm^2 capacitance density, served both to confirm the ability of the process to produce high surface area structures as well as demonstrate its feasibility for producing high density capacitors.

Next, the fabrication of highly-laminated magnetic cores and power inductors based on sacrificial multilayer electroplating is presented, along with the design and development of a system for characterizing inductor behavior at high-frequency, high-flux conditions. Experimental results for devices of a number of lamination thicknesses are presented and analyzed, and a trend of suppressed eddy current loss is demonstrated.

Finally, the design and operation of DC/DC converters built around these highly-laminated-core inductors are presented. Both buck and boost topologies were implemented and operated from 2 MHz to 6 MHz. For comparison, an off-the-shelf ferrite inductor was also used in place of the laminated-core inductors. While the ferrite inductor dissipated less power than the microfabricated inductors, its magnetic core was also drastically larger than that of the microfabricated inductors.

The principal contributions of this work are: the demonstration of a MEMS-based microfabrication process for the fabrication of both inductors and capacitors; the advancement of this microfabrication process; and the ability of this microfabrication process to produce inductors with minimized eddy current losses and low core volume requirements.

CHAPTER I

INTRODUCTION

1.1 Motivation

The marvels of modern electronics depend every bit as much upon power supplies as they do upon microprocessors, optical displays, and nonvolatile memories. However, the former are not as widely researched the latter, in large part because power converters generally do not play a large role in determining overall system performance. This has resulted in power electronics receiving considerably less design, fabrication, and packaging research attention than the other, more outwardly visible system components have.

This arrangement is economically sound, but it is bound for change nonetheless. Two continuing trends in the electronics industry – mobility and energy efficiency – are now exerting major pressure on power converters. Responses to this pressure have come from circuit design innovations like soft switching and multiphase topologies, or from advances derived from mainstream electronics, chiefly in the form of faster power transistors and more powerful control circuits.

More radical solutions, namely advanced *power* systems packaging and *power* passive components fabrication, bear the mark of this history; they are only weakly accommodated by conventional packaging research, and power electronics-specific approaches are few in number. But, such radical solutions may one day become necessary for supporting continued advances in power electronics circuitry, if not for meeting overall electronics system size, weight, and efficiency demands.

Even if progress in power electronics is not critically slowed by limited innovation in packaging and components, modern electronics can benefit substantially from improved power electronics integration. Certainly, many conventional techniques for achieving smaller size or higher efficiency, such as soft switching or faster switching, could be facilitated by the reduced parasitics and electromagnetic interference (EMI) that fully integrated power passives would provide. As well, microfabrication-based packaging approaches are very apt to favor planar constructions, as opposed to the cubic and cylindrical constructions found in conventional discrete power magnetics

and capacitors. This will better match form factor demands not only of other circuit boards within the system but also of the entire system.

This document presents research directed to address the relative lack of power electronics packaging platforms. The essence of the approach is to bring to bear the rapidly growing body of fabrication knowledge found in the field of microelectromechanical systems (MEMS) upon this problem. Previous, proven research results in micromagnetics and high-speed microfabricated magnetic cores are extended and adapted to meet the practical needs of advanced power converters.

1.2 History

1.2.1 The Evolution of Power Converters

As the "Electric Age" debuted at the turn of the 20th century, power converters – systems to convert electricity from one form to another – were not required. Early power distribution grids were well-tuned to the loads they served, not only as a consequence of electricity's narrow application domain at the time – lights and motors – but also a consequence of electricity providers' vigorous competition. In many ways, the power grid as it now exists was specified by the needs of electricity's first end users. For instance, multi-phase AC distribution saw adoption not because of its large-scale efficiency savings, but because the first robustly self-starting motors required two or more phases.[43]

Inevitably, devices and uses for electricity came about which required power in forms that the contemporary grid could not provide, and the need for power converters thus materialized. An example of this was incandescent light bulbs. When powered by 25 Hz AC electricity, the bulbs produced a visible flicker. The solution to this problem, a motor with a directly coupled DC generator, can be considered as one of the first widely-adopted power converter applications.

Figure 1 shows the main eras of power electronics. With each new generation of devices, and each new era, fundamentally differing power supply needs were presented. Transformers were deployed relatively early within the power grid, and have stood until only recent decades as the primary means for AC/AC power conversion. The shift from magnetic amplifiers, fundamentally AC devices, to vacuum tubes, fundamentally DC devices, drove a pronounced increase in demand demand for AC/DC converters.

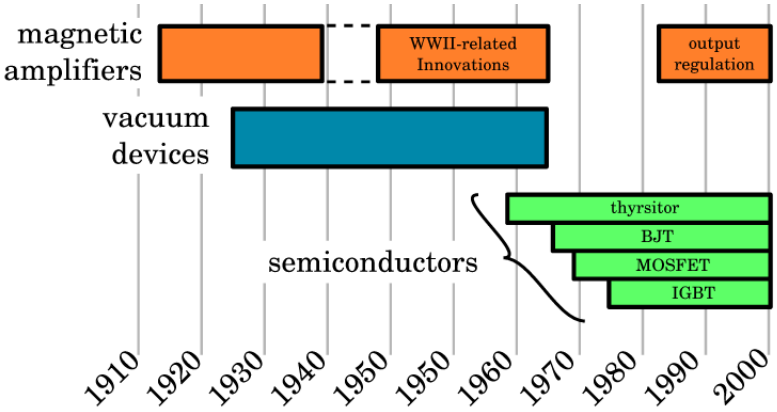

Figure 1: Power Electronics History (after [71])

In the 1970s, transistors, again fundamentally DC devices, became commercially feasible, and electronics became increasingly pervasive in consumer markets. Consequently, AC/DC converters became just as pervasive. [45] As battery technology improved and electronics continuously required less power, the field of portable electronics grew at a fast rate, and interest in DC/DC converters grew accordingly. A similar dynamic has driven interest in high-capacity AC/AC converters for use in power distribution grids. [74]

1.2.2 Emphasis on Power Converter Packaging

The proliferation of electronics among mass markets fueled a push towards faster and lighter electronics. As mobility and portability are universally on the rise in consumer electronics, this force is also beginning to bear directly on power supplies, since they still contribute a disproportionately large amount of size and weight.

As a result, power converters have been garnering more and more attention from system designers. Market predictions bear the same conclusion; most of the growth areas for power electronics are in sectors where size and weight are at a premium. [57] The generalities of this trend are not new in that its basic drivers are external to the field of power electronics itself.

Tightening regulations governing efficiency, power grid pollution, and electromagnetic compatibility (EMC) provide one broad impetus, while the ever-evolving needs of electron devices provide another. Among the former, [39]

- the need for harmonics-limiting power factor correction,

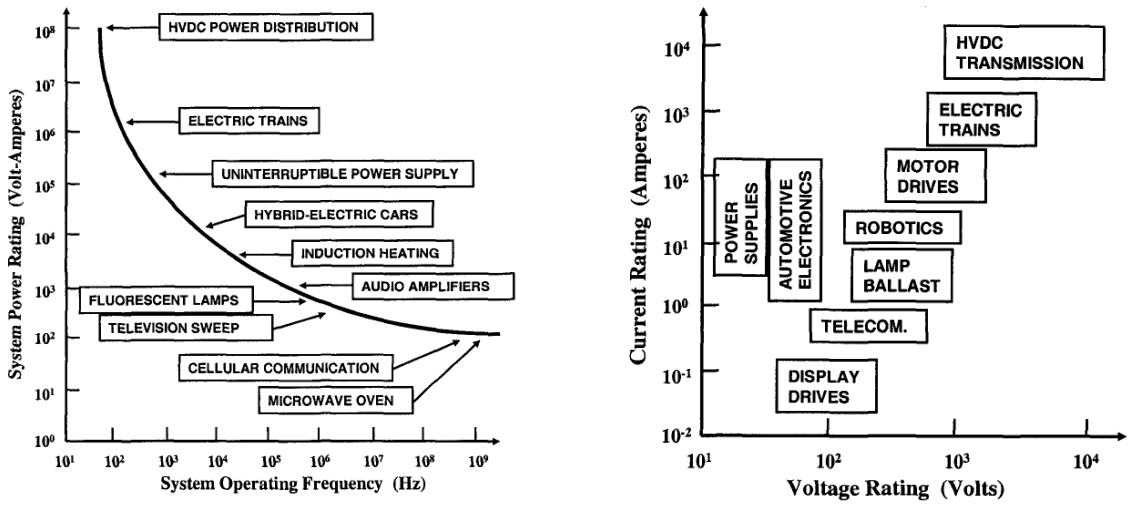

(a) (b)

Figure 2: Operating Frequency and Power for Electronics Systems[3]

- critically slow progress in battery energy densities, requiring ultra-efficient battery-sourced converters,
- growing energy costs, leading to regulations requiring more efficient conversion,

stand out, while the more pronounced aspects of the latter include

- digital integrated circuits (ICs) moving towards low-voltage, high-current supply rails, projected to approach 0.8 V and 200 A by 2012,
- widely diverging supply rails within a single system, often including 1.5 V and lower for processors, 3.3 V and 5 V for interface logic, higher voltages for flash memories, and potentially much higher voltages for displays, and
- continuing shrinkage of electronics systems, to the point that the power supply contributes an "application-killing" quantity of size and weight.

1.3 MEMS Microfabrication for Power Electronics

1.3.1 MEMS Microfabrication Processes

1.3.1.1 Thick polymer films

Driven in large part by the electronics packaging industry, photoresists that can be spun into thick layers and patterned with high aspect ratios are under continuous development. Examples include

Shipley SPR20 [41] and Futurrex NR-21 [18] The distinct advantage of photoresists for formation of thick polymer films is the relative ease which with they can be removed. Alternative thick-film materials, such as SU-8 or polyimide, are very difficult to remove.

In addition to photoresists, other polymer materials can be used to create and pattern thick films. Common examples include Microchem's UV-patternable epoxy SU-8 and any of a wide variety of photosensitive polyimides. These materials can enable thicker structures and higher aspect ratios than are generally achievable with photoresist films. [51]

LIGA, a German-language acronym for Lithographie (photolithography), Galvanoformung (electroplating) and Abformung (molding), is a well-known MEMS fabrication process. LIGA relies not only on thick, high-aspect ratio polymer processing – exposure of PMMA photoresist with X-ray radiation – but also on electrodeposition through the resulting PMMA layer. [19]

1.3.1.2 Structural electrodeposition

Structures with microscale features can be formed through the electrodeposition of metal, as is done in LIGA. Electrodeposition can be electroless, in which metal deposits spontaneously onto a seed layer from solution, or electrolytic, in which an external source of electricity is connected between the solution and the seed layer to drive the deposition reaction forward. Popular electrodeposited materials include copper, nickel, silver, chromium, and gold, and alloys of these with other metals including iron, cobalt, tin, lead, and zinc.[20]

Electrodeposition is typically done with a photoresist plating mold, although other polymers or even silicon can be used as a plating mold. Figure 3 illustrates this process. The partially removed mold was not a photopatterned polymer, but a laser-drilled polymer layer.

1.3.1.3 Sacrificial Release

Sacrificial release is a venerable MEMS fabrication processes, used to achieve free-standing structures. The essential process is illustrated in Figure 4. Sacrificial release processes of one form or another are used in virtually all MEMS fabrication processes. From this wide variety of practical applications, a number of generic needs and challenges can be identified. [21]

Sacrificial Material Systems

The materials that can be used in a sacrificial release process are subject to certain constraints.

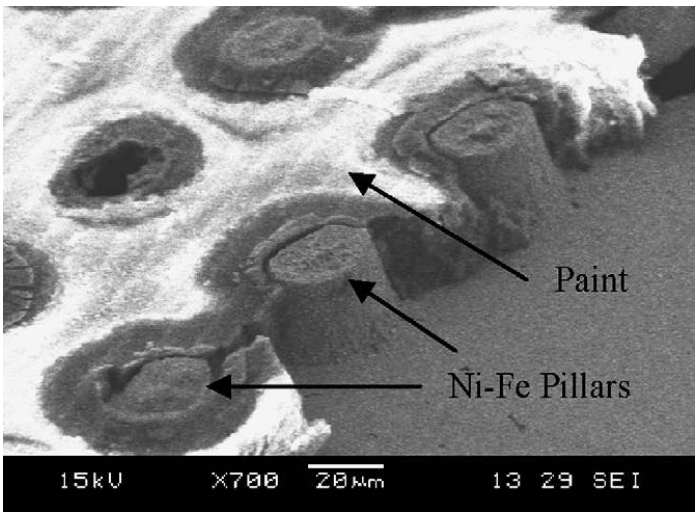

Figure 3: Micromachining Based on Electrodeposition. The pillars are formed by electrodeposition of nickel-iron, and the layer of paint, with laser-drilled holes, serves to dictate the shape of the plated pillars.[55]

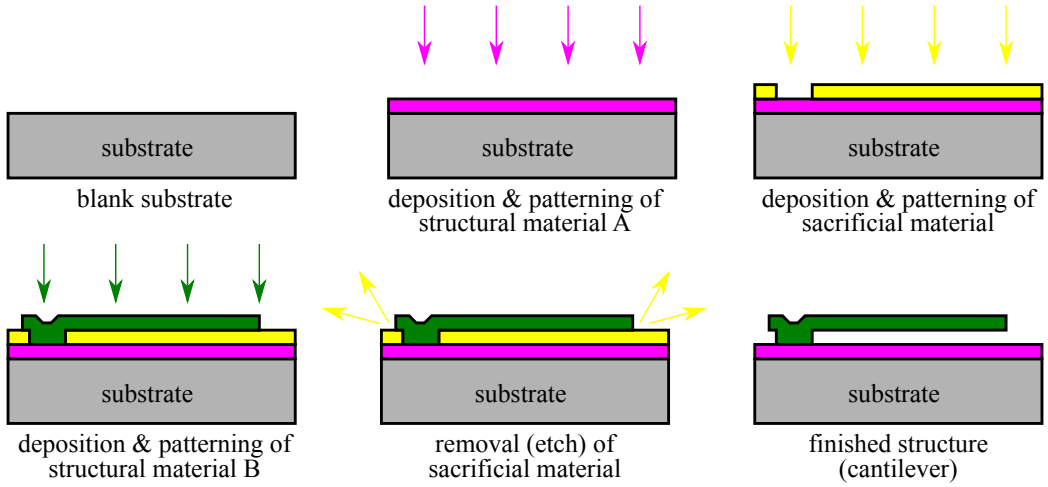

Figure 4: Generic sacrificial release process. Layers are sequentially deposited onto a substrate, followed by the removal of selected intervening layers. Typically the selective removal is accomplished by chemically selective etch processes.

Figure 4 suggests the principal ones. First, the structural and sacrificial materials must be compatible in that they can be deposited atop one another. Each deposition process imposes its own thermal budget, handling requirements, and chemical environment that preceding layers may not be compatible with. As well, there can be considerable practical advantage in being able to deposit all materials with the same equipment. For instance, if all three materials in Figure 4 are deposited by the same PECVD system, the entire unreleased structure can be formed from a single run of the PECVD machine. The second major constraint on sacrificial release material systems is that a robust sacrificial etchant, capable of removing the sacrificial material(s) without damaging the structural material(s), must be available. Accordingly, the etchant must be considered as an integral part of the material system.

A large number of material systems have been identified: SiN/Si/KOH, Si/SiOx/HF, and metal/photoresist/acetone

Sacrificial Etch Mass Transport

The sacrificial etch step, particularly the removal of the sacrificial material, imposes some fabrication challenges that are not strongly dependent on the material system in use. One of these is the difficulty of transporting chemical reactants to and from the reaction sites. This is alleviated by the inclusion of etch holes throughout the released structure. These essentially enforce a maximum distance that reactants must travel.

Another approach to solving this issue is the use of vapor (gas-phase) etchants. In particular, vapor-phase HF for the sacrificial removal of silicon dioxide is very effective.

Stiction Control

A major difficulty with sacrificial release processes is the phenomenon known as stiction. Stiction, illustrated in Figure 5, is the permanent collapse of the void created by the removal of the sacrificial material. Typically, this collapse is brought on and sustained not by gravity or any bulk forces, but rather by adhesive forces. In the case of wet sacrificial etches, the combination of adhesion between the released structure and whatever liquid remains in the sacrificial void after the etch – generally a rinse solvent – and the surface tension (self-adhesion) of the liquid initiates the collapse. Adhesion between the structure's surfaces perpetuates the collapse.

Several techniques to prevent stiction have been developed. In general, it is possible to alter the

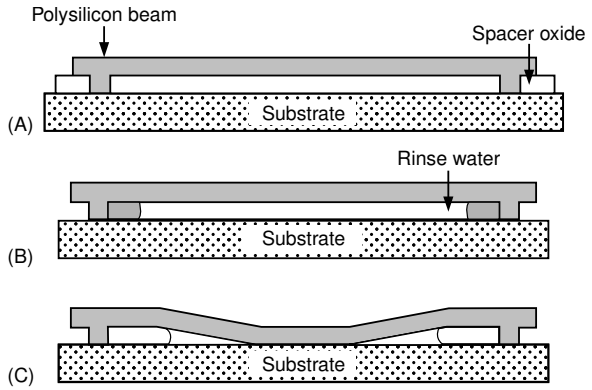

Figure 5: Stiction in sacrificial oxide etch [19]

composition of the liquid through the addition of other, miscible liquids, which then diffuse into the sacrificial void. In this manner, a higher-tension liquid such as water can be largely replaced by a lower-tension solvent such as methanol, to facilitate the subsequent drying step.

The physical behavior of the final rinse liquid can be modified in a number of ways to help avoid stiction. The liquid can be frozen and removed by sublimation; this removes the action of the liquid's surface tension. Alternatively, supercritical drying, where a gas such as CO2 is brought to a supercritical phase and circulated, can be used to remove the remaining solvent without imparting the forces that would be created by evaporation of the liquid.

Another approach to stiction control involves modification of the structural surfaces themselves. So-called "anti-stiction bumps" prevent contact between structural surfaces over a large area. As well, low-adhesion coatings can be applied to the structural surfaces so that any contact that does occur entails less stress. [19]

1.3.1.4 CMOS Integration

Synergistic Combination

MEMS fabrication techniques generally offer two compelling advantages in terms of commercialization. First, they are often based primarily on batch-oriented processes, not only promising the economies of scale that have defined the dynamics of the CMOS industry, but also allowing a great deal of cost insensitivity across a given device's design space.

Second, MEMS processes are typically highly compatible with silicon wafer substrates practically identical to the ones used in CMOS integrated circuit manufacture. This enables MEMS wafer

foundries to take advantage of the well-developed semiconductor capital equipment markets, reducing the scope of the required custom tool development and allowing effective use of the relatively immense body of knowledge and experience that has come from decades of CMOS development and manufacture.

Each of these two advantages stem from the historical development of MEMS as an outgrowth of integrated circuits. They can, in fact, deliver even greater benefits if a MEMS fabrication flow can be merged with CMOS fabrication flow, producing at the end a wafer with fully-integrated CMOS and MEMS structures. In addition to enabling new types of CMOS-MEMS hybrid platforms and allowing interconnect with minimized parasitics, such process integration can also potentially yield enormous cost savings.

Tautologically, the compatibility between CMOS and MEMS fabrication processes depends on both processes. Practically, CMOS foundries do not offer a great array of options to achieve compatibility with MEMS fabrication. This is due not only to the numerous challenges intrinsic in producing modern ICs, but also due to the relatively small size of the market for CMOS-MEMS devices, relative to the market for CMOS-only (*i.e.* normal ICs) devices.

Post-CMOS MEMS Integration

There are a number of ways in which it is possible to perform the merging of the MEMS and CMOS fabrication flows, but they can be classified according to the chronological ordering of the two flows, engendering the three self-explanatory terms Pre-CMOS, Intra-CMOS, and Post-CMOS integration of MEMS processing. [9]

However, the CMOS fabrication flow makes Post-CMOS integration the most practical. CMOS fabrication imposes very high thermal budgets, with diffusion and oxidation processes that easily approach 1000 C. Many MEMS materials, particularly metal films and polymers, simply cannot tolerate the CMOS thermal budget. As well, planarization requirements for CMOS fabrication, due to lithographic systems that sacrifice depth-of-focus in favor of greater resolution, can also preclude the prior use of substantial micromachining steps.

Correspondingly, in this research, the vision of Post-CMOS integration is pursued. This requires the observation of a thermal budget of approximately 400 C, and the use of processes that would be chemically compatible with typical CMOS passivation materials. Fortunately, this is readily

achievable with the processing presented in this thesis, all of which occurs at less than 250 C, and involves only weak wet etch (*ex.* 1:49 HF:water) treatments.

Most importantly, the targeting of Post-CMOS integration enables the ultimate integration of the presented devices onto finished CMOS wafers, at the wafer level, to produce fully-integrated, single-chip power converters.

1.3.2 Integrated Power Systems

Using the processes presented in the previous section, the implementation of fully-integrated power converters and, more generally, power management systems, has been achieved. The maturity and robustness of these microfabrication processes, as well as proven techniques for fabrication of the components within the power systems, play a central role in their success.

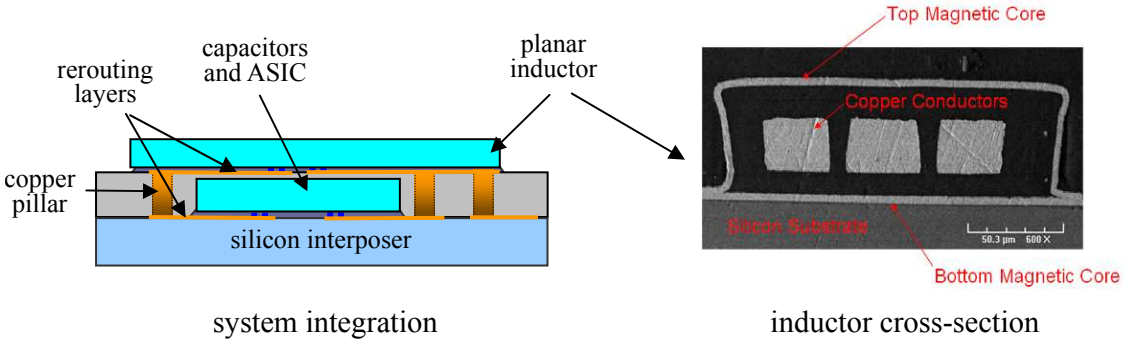

Figure 6: Integrated DC/DC Converter for Wireless Sensor Nodes.[27]

Figure 6 shows one such system, a single-chip DC/DC converter running at 15 MHz, incorporating a 1 to 4 μH inductor, 100 nF capacitors, a CMOS circuit fabricated with a 0.35 μm process, and achieving an efficiency of over 65 %. Silicon-based polymer and electroplating techniques were used to achieve the full integration of the separately- processed inductor and IC substrates. As well, as suggested by the inductor cross-section in Figure 6, a complex sequence of microfabrication steps was required to fabricate the inductor. [27]

Another example of successful, microfabrication-based power system integration is shown in Figure 7. In [14], a switched-capacitor power converter is integrated with a single-walled carbon nanotube (SWNT)-based supercapacitor to provide comprehensive power support as might

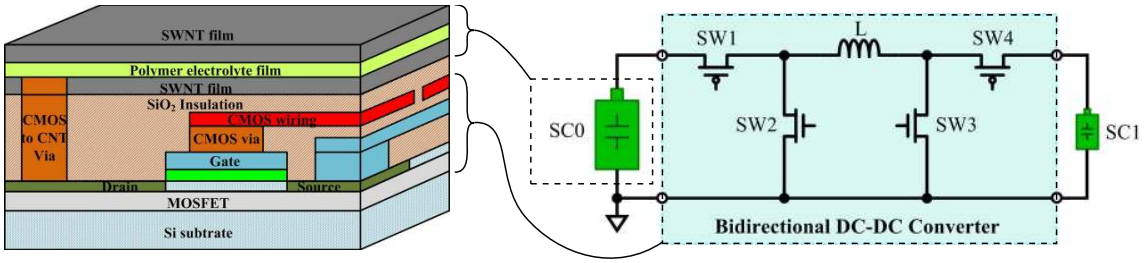

system implementation system schematic

Figure 7: Integrated Power Storage and Conversion System[14]

be needed by a microharvester-powered wireless sensor node. Not only are thick-regime, low-temperature insulator and electroplating steps used to achieve effective system integration, but a nanofabricated energy storage component – the SWNT supercapacitor – plays a central role in the practical importance of the system.

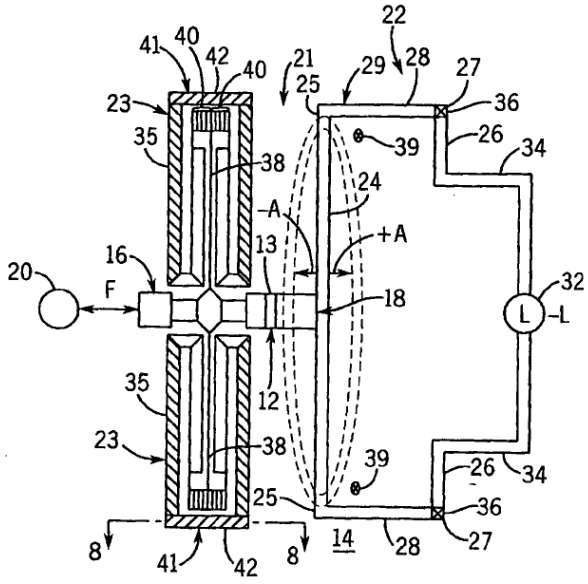

Figure 8: Isolating, all-MEMS Power Converter [25]

Less conventional examples of a fully-integrated power converters are found in [47] and [25]. Each of these patented devices implements DC/DC power conversion in a single microdevice that carries out electrical-to-mechanical and subsequent mechanical-to-electrical conversion. While a pair of co-fabricated, engaged rotary machines are used in [47], a resonant, flexure spring-mounted boss is used in [25]. Both of these applications are based on the extensive use of standard bulk and surface micromachining processes.

1.3.3 Microfabricated Power Components

1.3.3.1 Integrated Passive Electronics Modules

Even in cases where full system integration is not easily reached, microfabrication-based fabrication of individual components and multi-component modules can still deliver great benefit.

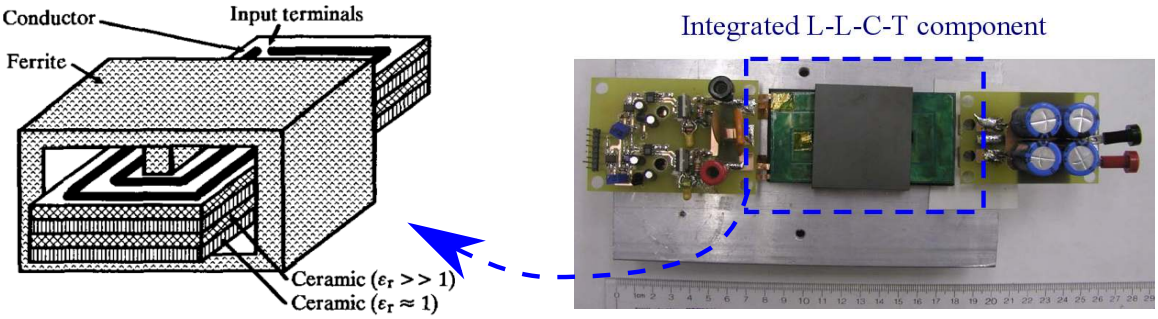

Figure 9: Integrated passive electronics modules (IPEMs) generic construction[73] and application in 1 kW DC/DC converter [36])

The integrated passive electronics module (IPEM) approach to system integration, exemplified in [36], simultaneously creates a number of practical advantages, and is enabled by existing microscale fabrication processes. An IPEM consists of any number of co-fabricated passive components, and is typically fabricated with a low-temperature co-fired ceramic (LTCC) multilayer process. This approach delivers the standard integration benefits of reduced component costs, facilitated system assembly, and improved interconnect parasitics. However, one of its particular advantages is that it implements the grouping of components with similar thermal needs, and the ceramic substrate, with its high thermal conductivity, facilitates the removal of heat.

Remarkably as well, the LTCC-based IPEM approach enables the practical deployment of new types of passive networks, which can in turn potentially enable new types of system architectures. One such passive network is the LCT (inductor- capacitor-transformer) module, shown in [32]. This type of module is used in high-efficiency, quasi-resonant power conversion circuits such as LLC and LCC circuits, and can be seen as an enabling technology for these advanced designs.

1.3.3.2 PCB-based Integration

Even well-studied system integration approaches such as printed circuit boards (PCBs) and other printed-circuit interconnect technology can be leveraged in new, power electronics-centric ways.

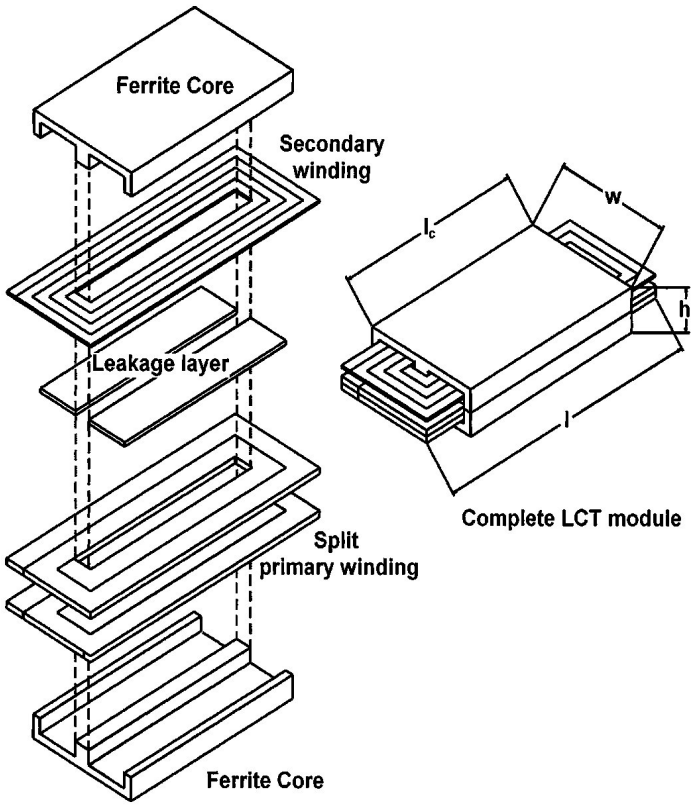

Figure 10: LCT Module Construction (exploded)[32]

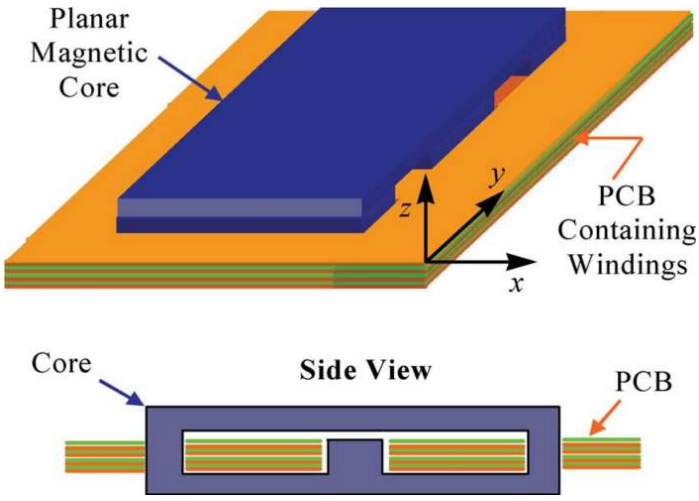

Figure 11: PCB-wound Planar Magnetics[60]

Planar magnetics, illustrated in [60] are now widely used in commercial power converters [61]. In planar magnetics, PCB traces form the coils of magnetic devices such as inductors and transformers, and the magnetic cores are assembled about the PCB coil structures.

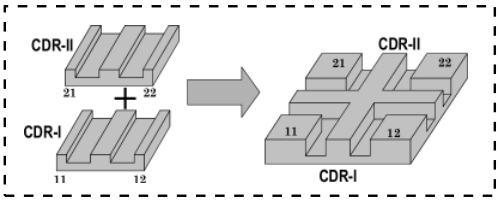

(a) magnetics integration

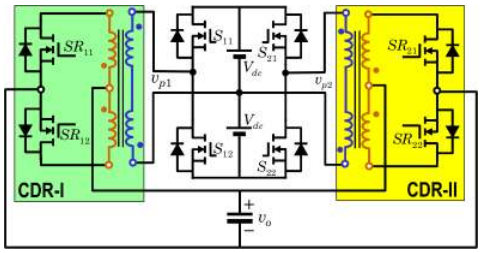

(b) system schematic

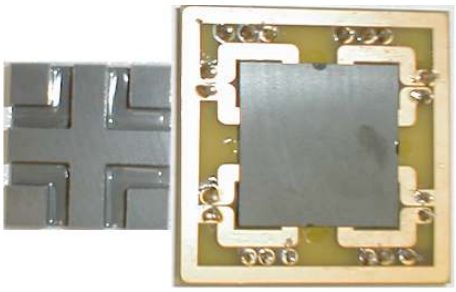

(c) magnetics implementation

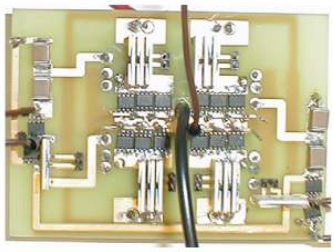

(d) system implementation

Figure 12: Matrix-Integrated Magnetics [12]

The planar magnetics approach has been extended to create new devices, as well as new circuit topologies, in the matrix-integrated magnetics approach. In this miniaturization-minded research, a regular fabric (matrix) of magnetic structures is envisioned, with PCB coils providing the application-specific configuration of the magnetic circuits.

In the first place, this approach would justify the large-volume production of magnetic cores in the appropriate matrix configurations. This would deliver concomitant economies of scale in the manufacture of the cores, as well as the assembly of the converters. This approach, given a suitable body of knowledge about the corresponding converter circuits, could be scaled, even across fabrication technologies.

The in-place fabrication of components during PCB manufacture is also a widely followed approach to improved power electronics integration. In [35], shown in Figure 13, an extra sequence of steps is added to the processing of the outer layers of a PCB to add passive components. These passive components are significantly thinner than the PCB's dielectric layers, and this work demonstrated that this can be distinctly advantageous, if a suitable integration process is available.

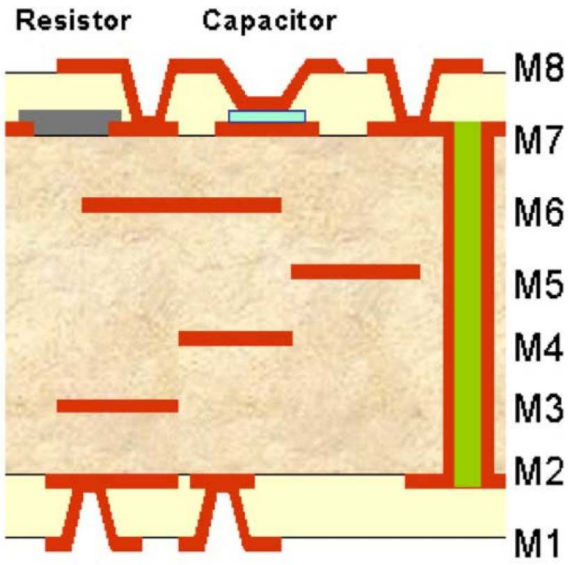

Figure 13: PCB-embedded Capacitors[35]

The thin capacitors enable the achievement of high capacitance densities, as needed power supply filtering. The thin resistors enable the use of high-conductivity resistive materials, so that low-range resistors can be produced with laterally compact dimensions. Finally, the large difference in thickness between the passive components and the PCB dielectric facilitates the photopatterning, etching, and plating of the outer copper traces; thicker passives might produce nonplanarities in the outer surfaces, on which patterned copper films must be created.

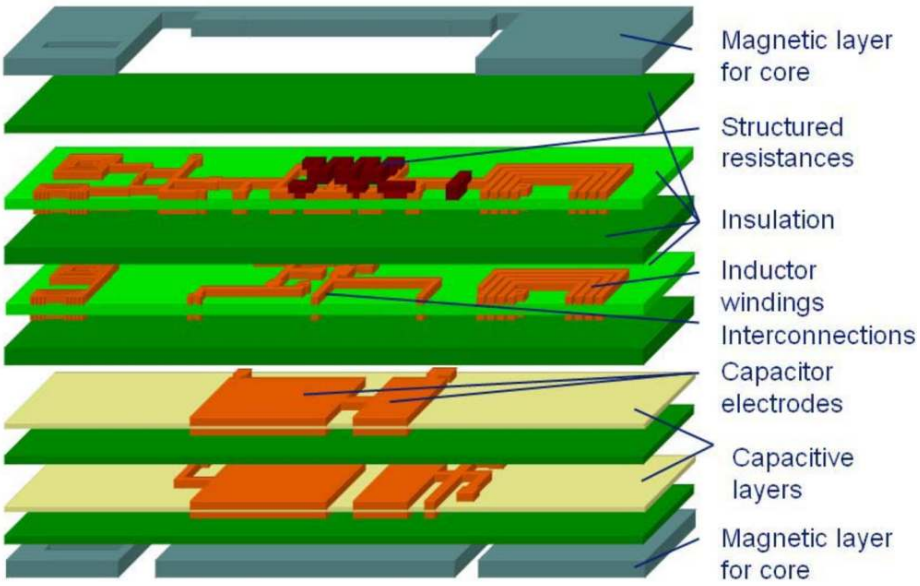

Figure 14: Comprehensive PCB-Integrated Passives Platform[69]

Figure 14 shows a more comprehensive approach to PCB-embedded passives. In this work, magnetic core materials are provided for, enabling the incorporation of high-density magnetic devices.

1.4 Contributions of This Work

The research and development presented in this thesis falls well into line with this process-oriented, bottom-up pursuit of higher-performance, better-integrated power systems.

1.4.1 MEMS-based Process for Power Electronics

The microfabrication process of sacrificial multilayer electroplating is further developed in this work, specifically for practical use in power system packaging. This advanced process serves as the basis for the devices produced herein, as well as to establish the great potential of its future applications.

1.4.2 Microfabricated Power Components

The primary point of demonstrated impact for the sacrificial multilayer plating process is in the creation of high-performance devices for power electronic applications. A high-density capacitor as well as highly-laminated, high-speed magnetic cores were produces. These two component classes typically furnish not only the majority of the size and weight of power conversion systems, but also the largest integration challenges.

1.4.3 Deployment into Power Converters

The demonstration of these devices in high-efficiency power converters is important not only as an ultimate verification of their functionality, but also as a basis from which more advanced integration methods can be pursued in future work.

1.5 About this Thesis

This thesis is organized into six chapters. The relationships between these chapters, as well as significant auxiliary work, is shown in Figure 15. Chapter 1 presents a history of the field, a motivation for the research presented, and related results from other authors.

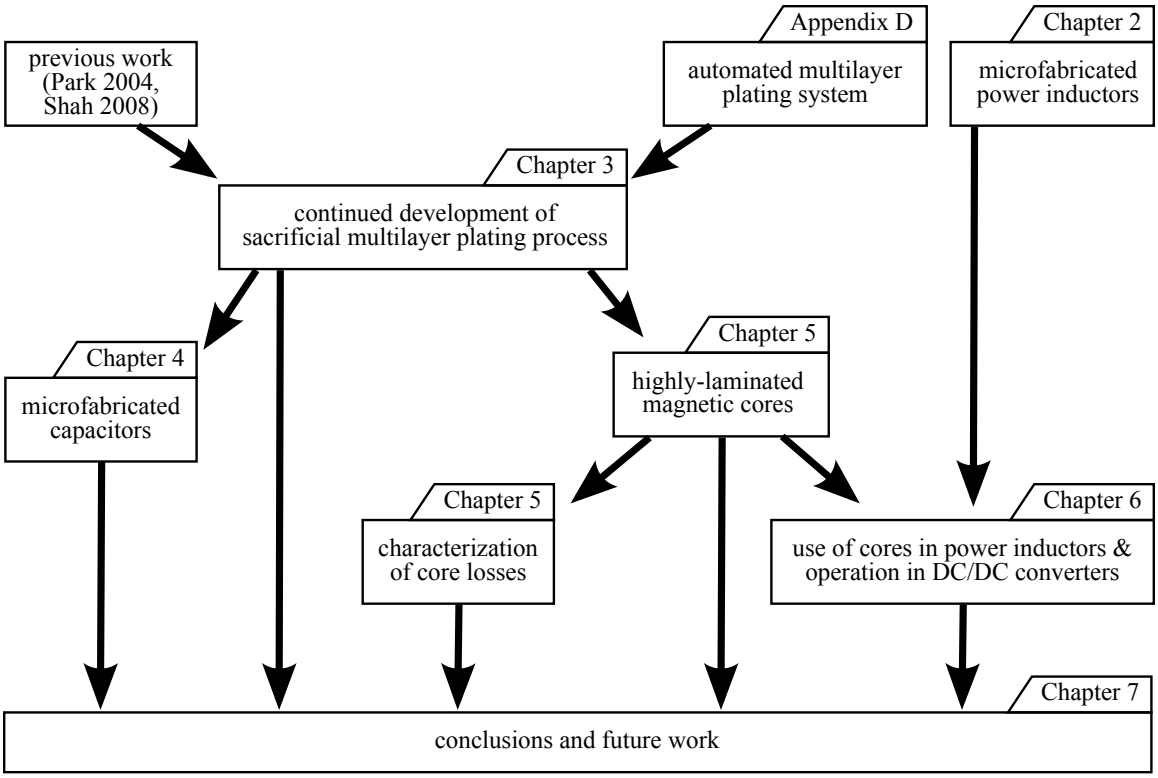

Figure 15: Conceptual Map of this Thesis. The work detailed in Chapter 3 enabled the concrete developments presented in Chapters 4,5, and 6

Chapter 2 presents the fabrication and characterization of an integrated power inductor. This work is conceptually important, as it shows the viability of a high-efficiency, single-die DC/DC converter based on the same magnetic materials used in later work. As well, it was the vehicle for the development of a substantial body of practical knowledge related to the integration of microfabricated components into power converters.

Chapter 3 presents the chief advances and results in the continued refinement of the sacrificial multilayer plating process. This includes not only an account of the difficulties encountered and the solutions developed accordingly, but also a review of competing approaches and a brief exposition of potential future applications for this promising technology. This topic also encapsulates the development of an advanced robotic multilayer electroplating system, which is discussed more comprehensively in Appendix C.

The initial application of the newly advanced sacrificial multilayer plating process, a high-density on-silicon capacitor, is presented in Chapter 4. This device proved the effectiveness of

the fabrication flow at producing high surface areas. It also establishes the potential for future on-chip integration of filter capacitors to complement the integration of inductors, producing a truly self-contained single-chip DC/DC converter. Appendix A provides some supporting mathematical results that are used in the calculations in this chapter, as well as in Chapter 5.

The development, fabrication, and characterization of submicron-laminated magnetic cores, the most substantial application of the sacrificial multilayer plating process, are presented in Chapter 5. The characterization of the cores entailed the development of a complex characterization system, as well as the marshaling of substantial a mathematical framework; this is all covered in Appendix B.

Chapter 6 presents the ultimate step in demonstrating the value of this research: the operation of the submicron-laminated cores in high-frequency SMPSs. The design of the inductors and power supplies, as well as the results constitute this final chapter.

CHAPTER II

MICROFABRICATED POWER INDUCTORS

This chapter presents the fabrication, characterization, and in-system operation of a microfabricated power inductor. A general background on inductors is presented first, followed by details of the device's design and fabrication. Finally, the device-level characterization and system-level characterization of the device is presented.

2.1 Background

2.1.1 Inductor Theory

An inductor is one of the fundamental lumped two-terminal electrical components. It generates a voltage $v(t)$ that is proportional to the time derivative of its current $i(t)$ by its inductance L:

$$v(t) = L\frac{d}{dt}i(t) \tag{1}$$

This definition covers only behavior at the inductor's terminals, allowing for any conceivable implementation, utilizing any electronic or physical means.

2.1.1.1 Electromagnetic Induction

However, in practice, inductors are typically implemented based solely on electromagnetic transduction. The electromagnetic transduction pathway that creates inductive behavior is the same in all such inductors, and is described precisely by Maxwell's equations. An electrical current $\vec{J}$ gives rise to a magnetic field according to a spatially-distributed magnetic field $\vec{H}$:

$$\vec{\nabla} \times \vec{H} = \vec{J} \tag{2}$$

This magnetic field $\vec{H}$ drives a magnetic flux $\vec{B}$ according to the constitutive relation

$$\vec{B} = \mu\vec{H} \tag{3}$$

and the magnetic flux $\vec{B}$, as it changes over time, creates an electric field $\vec{E}$ according to Faraday's Law

$$\vec{\nabla} \times \vec{E} = -\frac{\partial B}{\partial t} \tag{4}$$

This electric field, integrated between two points in space

$$\Delta V = \int_A^B \vec{E} \cdot d\vec{s} \tag{5}$$

creates a potential difference, or voltage, V.

For purposes of modeling inductive behavior, the preceding development implies solving for the magnetic and electric field distributions at each instant in time. While this is the approach used in numerical simulation programs such as HFSS (http://www.ansys.com/hfss), it is unwieldy for manual analysis and design guidance. A useful approximation can be reached by taking into account the fact that, in many cases where inductance is a principal concern, the conductor forms a virtually closed path. In fact, the conductor typically makes many passes around approximately the same closed path; each pass is referred to as a turn, and the collection of turns is referred to as a coil [66].

Revisiting Equation 4 with this closed loop in mind, Stokes' Law can be applied to yield

$$\oint_{P_1} \vec{E} \cdot d\vec{s} = -\frac{\partial}{\partial t} \int_A \vec{B} \cdot d\vec{A} \tag{6}$$

where the contour integral is evaluated along the path of a single turn P_1, and the area integral is evaluated over the path's interior area A. Meanwhile, the terminal voltage v, as given in Equation 5, evaluated along the conductor's total path P_N (*i.e.* N_t turns), is given by

$$v = \oint_{P_N} \vec{E} \cdot d\vec{s} \tag{7}$$

$$= N_t \oint_{P_1} \vec{E} \cdot d\vec{s} \tag{8}$$

$$\tag{9}$$

Capturing the right hand side of Equation 6 in a new quantity Φ, known as the total flux,

$$\Phi = \int_A \vec{B} \cdot d\vec{A} \tag{10}$$

the structure's inductance is modeled by the relation

$$v = N_t \oint_{P_1} \vec{E} \cdot d\vec{P} \tag{11}$$

$$= -N_t \frac{\partial \Phi}{\partial t} \tag{12}$$

This transduction pathway acts any place where current is flowing, but a variety of analytical constructions are used to model inductance in different cases.

2.1.1.2 Self Inductance of a Straight Conductor

An isolated conductor will develop a voltage fluctuation due to changes in the current in carries; thus, it is said to exhibit self inductance. In this case, the magnetic field responsible for the inductance is located in the space immediately surrounding the wire. By Ampere's Law, the intensity of the magnetic field is symmetric about the conductor's axis, and varies inversely with distance from the conductor. In the case of a straight, cylindrical the magnetic field developed has a simple structure. This allows systematic evaluation of the conductor's self-inductance. For a cylindrical conductor of length l and radius r, the self-inductance L is [23]:

$$L = 0.002l\left[\ln\left(\frac{2l}{r}\right) - 1 + \frac{r}{l} + \frac{\mu_R}{4}\right] \tag{13}$$

$$\approx 0.002l\left[\ln\left(\frac{2l}{r}\right) - 0.75\right], \qquad l >> r \text{ and } \mu_R = 1 \tag{14}$$

, where μ_R is the relative magnetic permeability, L is in nanohenries, l and r are in centimeters.

For conductors with rectangular cross section, calculation of self-inductance is more complicated, but closed-form expressions are still available. For a straight conductor with rectangular cross section described by with a and height b, the self-inductance L is [23]:

$$L = 0.002l\left[\ln\left(\frac{2l}{a+b}\right) + 0.50049 + \frac{a+b}{3l}\right] \tag{15}$$

$$\approx 0.002l\left[\ln\left(\frac{2l}{a+b}\right) + 0.50049 + \frac{a+b}{3l}\right], \qquad l >> r \text{ and } \mu_R = 1 \tag{16}$$

where L is in nanohenries and l, a, and b are in centimeters.

2.1.1.3 Mutual Inductance of Parallel Conductors

Similarly, if two conductors are located relatively close to one another, the magnetic field generated by one conductor's current can affect both conductors. Specifically, a change in the current of either

conductor will generate a voltage fluctuation *in both conductors*. In this case, the mathematical description of the inductive behavior is a generalized version of Equation 1:

$$\begin{bmatrix} v_1(t) \\ v_2(t) \end{bmatrix} = \begin{bmatrix} L_1 & L_m \\ L_m & L_2 \end{bmatrix} \begin{bmatrix} \frac{d}{dt} i_1(t) \\ \frac{d}{dt} i_2(t) \end{bmatrix} \tag{17}$$

where i_1 and v_1 are the current and voltage in one conductor, i_2 and v_2 are the current and voltage in the other conductor, L_1 and L_2 are the self-inductances of each conductor, and L_m is said to be the the mutual inductance between the conductors.

For two parallel inductors, when $\mu_R = 1$, the mutual inductance M, in nanohenries, is given by

$$M = \pm 2l \left[\ln\left(\frac{l}{D_G} + \sqrt{1 + \left(\frac{l}{D_G}\right)^2} \right) - \sqrt{1 + \left(\frac{D_G}{l}\right)^2} + \frac{D_G}{l} \right] \tag{18}$$

where l is the length of the conductors in and D_G is the geometric mean distance between the conductors, in cm. The sign of the mutual inductance is positive if the two conductors' currents are flowing in the same direction, and negative otherwise. D_G is approximately equal to the spacing between conductor centers, but also incorporates geometric effects in cases where the conductors' widths occupy a significant portion of the distance between their centers [23].

The mutual inductance expression in Equation 18 and the self-inductance expression in Equation 15 can be combined to accurately estimate the inductance of a planar spiral inductor. This is accomplished by dividing the inductor up into a number of segments and summing, across all segments and all combinations of two segments, the self-inductance and mutual inductance, respectively.

2.1.1.4 Magnetic Circuits

The previous discussion applies to cases where conductors are surrounded by free space (*i.e.* $\mu_R = 1$), and carefully accounts for geometric effects. However, in cases where the conductors are interlinked with high-permeability magnetic material, there is another approach to modeling of magnetic devices that allows straightforward analysis.

In this approach, illustrated in Figure 16, the magnetic fields are assumed to be confined within the magnetic material, forming a mostly-closed path, known as a magnetic circuit. The geometry of the magnetic material defines the geometry of the magnetic circuit, and the assumptions thus

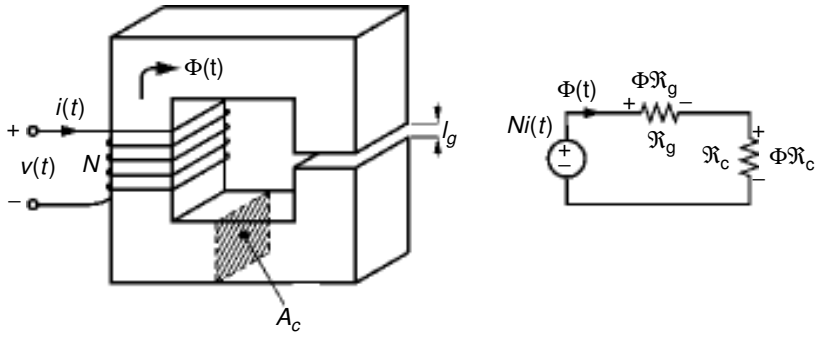

Figure 16: Generic magnetic circuit. A coil with N turns and current i creates a flux Φ that is distributed across the circuit's cross section A_C. [66]

enabled create a precise mapping of the magnetic system to an equivalent electrical circuit. Voltage is mapped to magnetmotive force, MMF, which can arise either from a current i in a winding with N turns ($MMF = Ni$) or from a magnetic field within the magnetic circuit ($MMF = -\int H \cdot dl$). Ampere's Law,

$$\int H \cdot dl = i \tag{19}$$

then transforms into Kirchoff's Voltage Law ($\sum V = 0$ for any closed loop within the circuit). Likewise, Gauss's Law ($\int_V B \cdot dA = 0$) is transformed into Kirchoff's Current Law ($\sum i = 0$ at any node in the circuit). The articulated set of linear relationships between dynamic quantites such as current, flux, field intensity, and voltage are merged into a single quantity, reluctance ($\Re$ in Figure 16, which acts analagously to resistance:

$$\Re = \frac{l}{A_C \mu_0 \mu_R} \tag{20}$$

where l is the path length and A_C is the cross-section area. With these transformations applied, relatively complex magnetic circuit structures may be analyzed using simple, familiar circuit analysis. [66].

Note that the inductor presented in this chapter, due to its unique structure, would require an analysis combining elements of magnetic circuit modeling as well as free-space coupling between parallel conductors. Detailed modeling of this inductor was not undertaken, but the above results provide a qualitative understanding of its behavior and a basis for comparing relative performance of the different device designs.

2.1.2 Switching Power Supplies

The inductor presented in this chapter, as well as the inductors in Chapter 5 and the capacitors in Chapter 4 are treated with respect to their potential application in power converters. In particular, these components are designed and evaluated relative to the demands of switched-mode power supplies (SMPSs). SMPSs provide conversion from an input power source to an output power source based on the combination of

- very fast modulation of internal voltages and currents, using electrically-controlled switches, such as MOSFETs and IGBTs, and
- subsequent removal of high-frequency modulation artifacts using substantial energy-storage components such as inductors and capacitors.

There are a large number of topologies that implement this combination of techniques. While only the most basic SMPS topologies were explored in this thesis, there are a small number of design concerns that are common across the vast majority of SMPS designs.

2.1.2.1 Common Design Goals

Most importantly, the SMPS's conversion efficiency, defined as the output power divided by the input power, must be maximized. This is not only because of a fundamental desire to minimize wasted energy (particularly in battery-powered systems), but also because any power wasted within a SMPS is heat that the system must then dissipate.

There are many mechanisms for power loss within a SMPS, including switching loss due to the finite transition time of the electrically-controlled switches and conduction loss in the wires and traces constituting the system. However, power dissipation internal to the energy-storage components used to filter out the modulation artifacts can easily dominate the total system loss. While circuit-domain solutions such as high-capacity MOSFET gate drivers or simply widened copper traces can alleviate the former loss sources, addressing filter component losses requires engineering of the energy storage materials (*ex.* magnetic core or capacitor dielectric) and/or optimization of the devices' internal structures. [72].

Another common goal in SMPS design is to minimize system size. This is generally accomplished by properly sizing circuit components to the required voltage and current levels. However, increasing the SMPS's switching frequency is an additional way to decrease system size. Increasing the frequency of the modulation artifacts reduces the capacity of the energy storage components needed to attenuate them. Increasing the switching frequency also allows the SMPS to respond more quickly to changes in output loading. However, increasing the switching frequency typically comes at the cost of increased switching losses.

2.1.2.2 Buck Converter

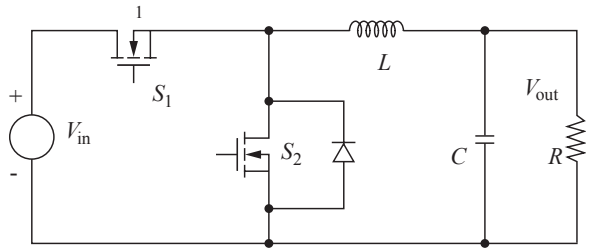

Figure 17: Buck converter circuit, showing an implementation with a synchronous power pole. [26]

One of the simplest SMPS topologies is the buck converter. The basic circuit of a buck converter is shown in Figure 17. A power pole is used to switch one terminal of the inductor L between V_{in} and ground. The inductor's other terminal is connected to the SMPS output. The inductor and the capacitor act as an LC filter which attenuates the AC components of the switched-node voltage, delivering only the DC component to the output. Note that this type of converter is only capable of delivering an output voltage that is less than the input voltage.

While the switched node voltage is at V_{in}, the inductor is being charged, and its current increases at a rate equal to $(V_{in} - V_{out})/L$. When the switched node voltage is at ground, the inductor's current decreases at a rate equal to V_{out}/L. This dynamic behavior is shown in Figure 18. It can be seen that the current ripple Δi_L is determined, to a first order, by the input and output voltages and the inductance L. This current ripple, in turn, is a primary determinant of the output voltage ripple.

2.1.2.3 Distributed Power Conversion

The importance of miniaturizing power converters, as targeted with this work, is emphasized with the development of distributed power conversion architectures. In most contemporary electronic

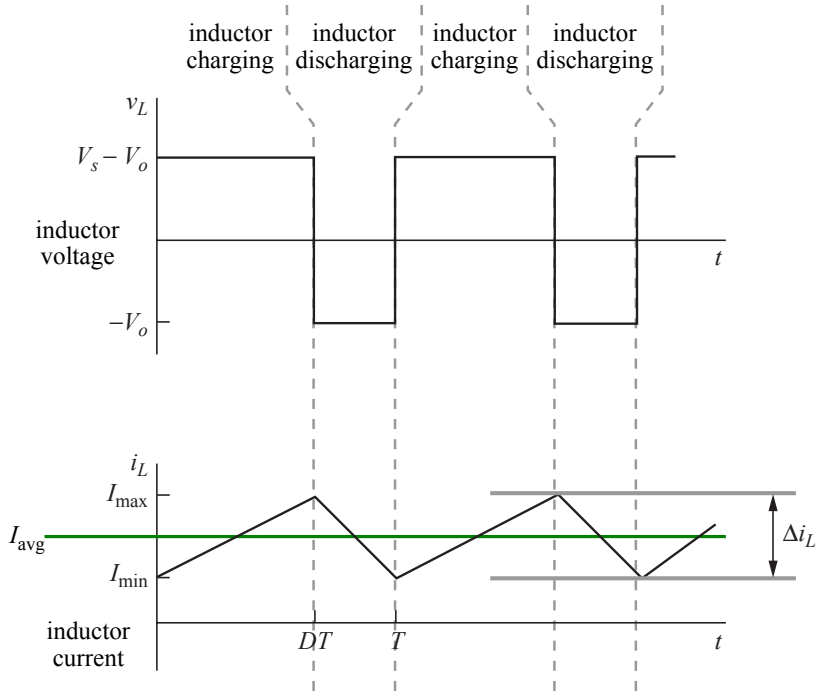

Figure 18: Buck converter waveforms, showing the voltage and current through the inductor. The assumption is made that the output voltage V_O remains constant, so that the switched-node voltage waveform is identical to the inductor voltage waveform, except for a DC offset equal to V_O. after [26]

systems, each subsystems requires a particular power supply voltages. In distributed power conversion, these power supply voltages are generated physically close to the subsystems that use them. This is in contrast to the more traditional centralized power conversion architecture, where a single power converter generates all of the required voltages, and relatively long wires or traces are used to deliver these voltages to the corresponding subsystems.

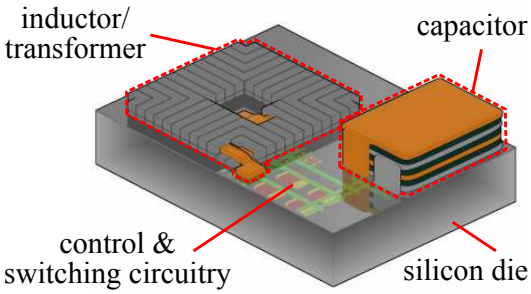

Figure 19: Implementation of a complete DC/DC converter on a single silicon die, showing the integration of the principal passives – capacitor and inductor – onto the die

2.2 *Design and Fabrication*

As is the case with many microfabricated devices, the design of this inductor is closely coupled with its fabrication process. Accordingly, these two aspects of this device are presented in interleaved fashion.

2.2.1 Full Device Construction

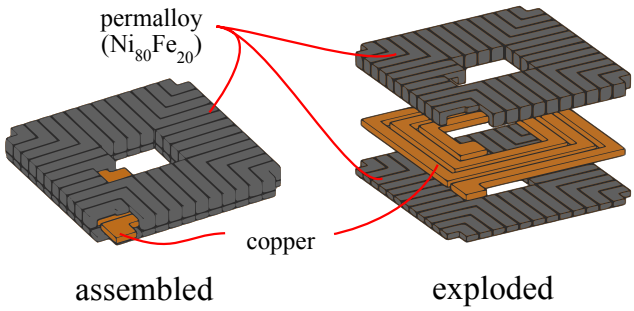

Figure 20: The microfabricated inductor consists of a copper spiral sandwiched between two magnetically-coupled nickel-iron layers

Figure 20 shows the full construction of the inductor. The device's winding is a single-layer planar spiral inductor, formed by electrodeposition of copper. This winding is augmented by a distributed magnetic core that enhances the electromagnetic coupling between parallel conductor segments. A cross-section of this magnetic core is shown in Figure 21, along with the expected cor-

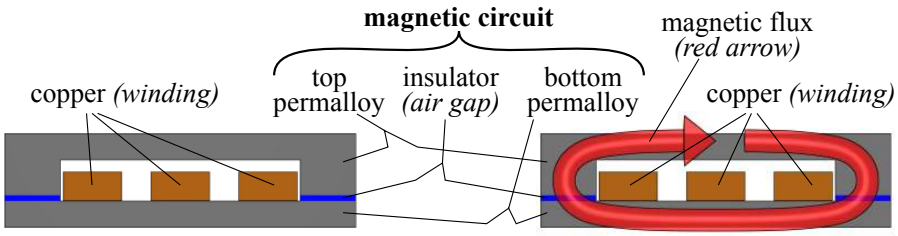

Figure 21: Full Inductor Cross-Section. The inductor features a full-length, gapped magnetic core enveloping adjacent conductor traces.

responding flux circuit. Two permalloy layers, both electrodeposited, form the bulk of the magnetic core. In between the permalloy layers, a thin layer of insulating material implements an air gap, completing the magnetic core.

The magnetic core is distributed in the sense that, although it is continuous in the direction of

magnetic flux, it is discontinuous along the direction of current flow. These discontinuities along the conductor segments, or lateral laminations, can be seen in figure Figure 20 or Figure 23.

Not shown in Figure 20 or Figure 21, various polymers are applied to provide insulation between structural members as well as mechanical support and passivation.

2.2.1.1 Design Highlights

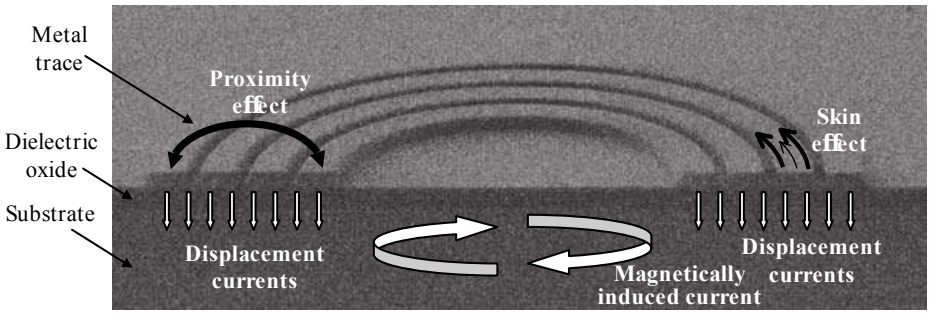

Figure 22: Loss Mechanisms in a Conventional Integrated Inductor. [54]

Figure 22 shows the main loss mechanisms in a typical on-chip integrated inductor [54]. The device presented in this chapter has a number of design features designed to address these losses, as can be seen in Figure 20 and Figure 21.

Most importantly, the design has provisions to substantially lower eddy current losses. First of these is the inclusion of a high-permeability magnetic circuit. This provides confinement of the inductor's magnetic field so that its intersection with the presumed substrate – conductive silicon – is minimized. Magnetic field lines that travel through the substrate can not only create eddy currents, but also introduce noise into the circuits contained in the silicon.

In addition, this inductor's structure will reduce eddy current losses within the magnetic core. As described in Appendix B, any magnetic circuit with nonzero conductivity that carries a nonconstant magnetic field will experience eddy currents and the concomitant power dissipation. For metallic magnetic materials, such as nickel-iron, lamination along planes parallel to the magnetic flux lines is a well-known way to limit these losses [66].

And, as can be seen, the division of the magnetic circuit into a number of laterally isolated pieces achieves this lamination. This is a principal advantage offered by this inductor over similar art.

Another advantage of this construction is the configurability of the magnetic circuit. The formation of the circuit by two separate electrodeposition steps allows the insertion of a thin nonmagnetic layer between the circuit's halves, creating an air gap. The ability to precisely control the air gap in an inductor's magnetic circuit is necessary to optimize the device's performance in a given application.

This design's exclusive use of electrodeposition allows very thick conductor and magnetic layers, as limited only by the capabilities of modern thick photoresists. Sputtering could have been used to deposit a wide variety of metal and alloy films of high quality, but it would have imposed two major limitations. First, sputter deposition to the thicknesses required for this design – approximately 25 μm for permalloy and 50 μm for copper – would take large amounts of time, driving up production costs. Second, sputtered layers of the required thicknesses would be difficult to pattern.

Finally, this inductor design, by not requiring any high-temperature (*i.e.* over 300 C) steps, presents minimal obstacles to post-CMOS wafer-level integration. Completed CMOS wafers are typically not able to withstand temperatures higher than 250 C. At these temperatures, not only do the metals forming the interconnect layers start to interdiffuse and degrade, but also the thermal expansion mismatch between the many post-gate-oxide films can create prohibitively high levels of shear stress. Chemical vapor deposition (CVD) or plasma-enhanced CVD (PECVD) of insulators can easily require temperatures in excess of 300 C, and ferrite-based magnetic materials can require firing temperatures over 900 C.

2.2.2 Simplified Device Constructions

It was found that, based on the full device design, shown in Figure 20, several interesting variants could be fabricated These variants all offered potential performance advantages or at least potentially appealing trade-offs between fabrication complexity and performance. These are enumerated in Figure 23, and each offers its own advantages and disadvantages.

2.2.2.1 Air Core

The air core inductor was of practical interest due to its extremely simple fabrication, but its main utility in this work was as a baseline, against which the other variants' performance could be evaluated. The air core variant not only guaranteed zero magnetic core losses, but also served as a control

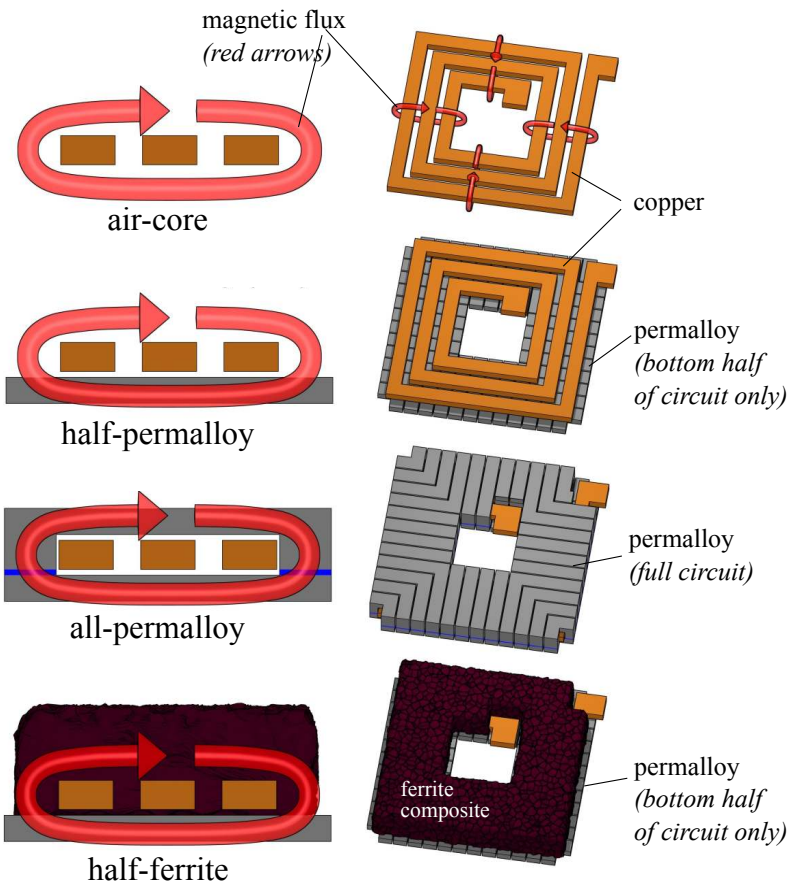

Figure 23: Variants of the Microfabricated Inductor

for any number of unknown or anomalous factors governing the performance of all inductors.

2.2.2.2 Half Permalloy

The half-permalloy inductor was selected for inclusion due primarily to its simple fabrication. As will be discussed, the fabrication steps following the deposition of the copper winding contributed a large amount of the difficulties seen in process development as well as in ultimate fabrication yield. The half-permalloy variant avoided these troublesome steps while still delivering a functional inductor.

At the same time, the half-permalloy inductor included many performance-enhancing features from the full (all-permalloy) design. The lower half of the magnetic core would not only substantially improve magnetic coupling between conductors, but also serve to steer magnetic flux away from the substrate.

Table 1: Microfabricated Inductor Variants

Variant	Advantages	Disadvantages
Air-core	• Simple fabrication • Zero core losses	• Lower inductance • Decreased separation from substrate
Half-permalloy	• Simple fabrication • Increased inductance • Decreased magnetic field reaching substrate	• Sub-maximal inductance • Nonzero core losses
Half-ferrite	• Simplest full magnetic circuit fabrication • No risk of shorting coil turns • Packaging compatibility	• Hysteresis losses in ferrite • Composite is unpatternable
All-permalloy	• Highest flux containment • Maximal control over magnetic circuit • Highest flux containment	• Eddy current losses • Difficult fabrication

2.2.2.3 Half Ferrite

The half-ferrite construction was, like the half-permalloy construction, attractive due largely to its simple fabrication. The half-ferrite inductor was formed by simply stencil-printing a polymer-matrix ferrite composite atop the finished half-permalloy inductor. Since the polymer/ferrite composite was nonconductive, there was no need to insulate the copper beforehand.

The half-ferrite construction also offered potentially improved performance due to the closed magnetic core provided by the ferrite. As well, by incorporating a modern ferrite material, the half-ferrite construction served as a point of reference for validating the overall electrodeposition-based approach against conventional alternatives.

2.2.3 Fabrication

The fabrication process for the microfabricated inductor is outlined in Figure 24. A test-grade 100mm-diameter silicon wafer, with approximately 1 μm of thermal oxide, is used as the beginning

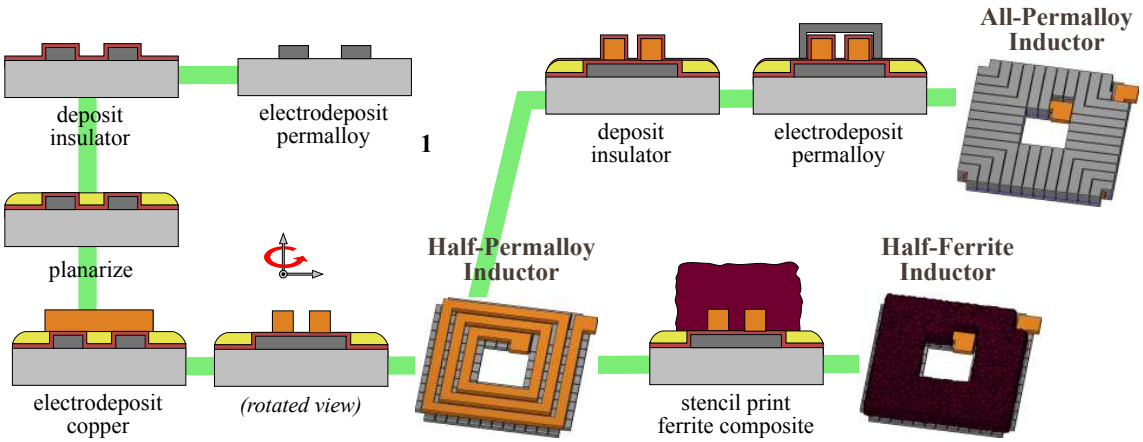

Figure 24: Inductor Microfabrication Process A two-branch process flow accommodated all inductor variants.

point. A 25 μm layer of permalloy (80% nickel, 20% iron) is anodically electroplated, using a 50 μm layer of Futurrex NR-2 thick negative photoresist as the plating mold.

After the plating mold and seed layer are removed, two polymer layers are deposited on top of the permalloy: one layer provides electrical insulation between the permalloy and the subsequent copper layer, and the other layer provides planarization to facilitate all remaining lithography steps. On top of the now-planarized and now-insulated permalloy layer, a 50 μm layer of copper is deposited, forming the inductor coil, again using NR-2 for the plating mask. At this point, the "half permalloy" inductor variant is complete.

At this point, continued fabrication can be executed to create either a "half-ferrite" inductor or an "all-permalloy" inductor.

In either case, after the copper plating mold and seed layer are removed, a layer of parylene is deposited on top of the copper to provide electrical insulation. After all subsequent fabrication steps, and immediately before device characterization, the parylene is removed from the coil ends (*i.e.* device terminals) to allow electrical test.

The final step for a half-ferrite device is the addition of the magnetic composite, to act as the topmost half of the magnetic circuit. A stencil mask, fabricated from a thick (200 μm) Kapton sheet with a carbon dioxide laser, is used to pattern a magnetic composite material. The magnetic composite was a 95% weight dispersion of ball-milled nickel-zinc ferrite powder (from Steward Inc) in an SU-8 matrix. SU-8 is a photopatternable epoxy from Microchem, Inc. (www.microchem.com),

although in this application the high density of opaque ferrite particles made photoexposure impractical. Instead, an extended elevated-temperature bake, at 150 C, was used to cure the SU-8. At the time of stencil-printing, the composite was highly viscous. After stencil-printing, the composite was cured with an extended 150 C hotplate bake.

The final step for an all-permalloy device is the electrodeposition of the topmost half of the permalloy magnetic circuits. As with previous steps, NR-2 thick negative photoresist was used as a plating mold for the deposition of a 25 μm layer of permalloy. After plating, the mold and seed layer are stripped, and the device is complete.

Achieving good ahesion between the electrodeposition seed layer and the parylene insulator covering the copper coils was an enduring challenge in this work. It was found that a brief oxygen etch in an RIE chamber immediately before sputter deposition of the seed layer enhanced the adhesion. It was also found that mechanical articulation between the top permalloy layer and the parylene layer greatly enhanced the durability of the finished devices.

2.3 Characterization

The microfabricated inductors were characterized both at the device level and at the system level. The device-level characterization provided proof of operation as intended, and also provided insight into the performance of the inductor variants. The system-level characterization was necessary to establish the ultimate utility of the inductors in a well-functioning power conversion system.

2.3.1 Device-level Characterization

2.3.1.1 Procedure

All device data was taken with a Hewlett-Packard 4194A impedance analyzer, with a 41941A High-Frequency Impedance Probe. Figure 25 illustrates the measurement arrangement. 16 AWG copper wire, attached to the 41941A probe head, was used to form probe needles.

All impedance measurements were taken as series inductance and series resistance ("Ls-R" mode) and the inductor's Q was calculated from inductance L and resistance R:

$$Q = \frac{2\pi f L}{R} \tag{21}$$

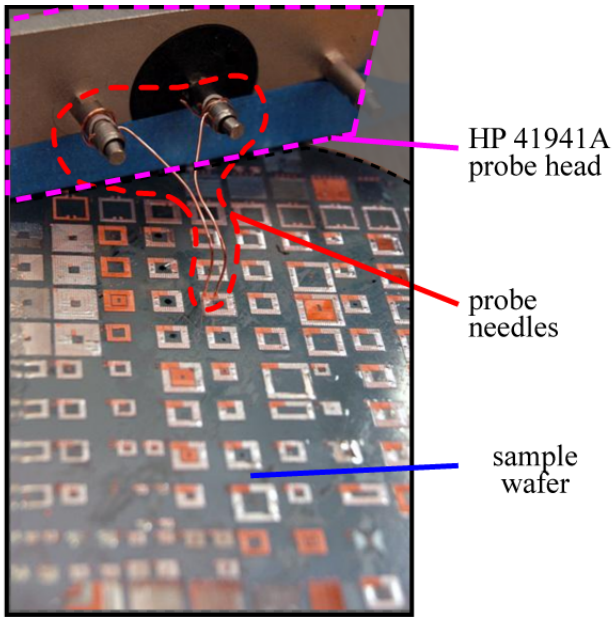

Figure 25: Device-level probing of microfabricated inductors. 16 AWG copper wire was used as probe tips.

where f is the frequency of the measurement. f was swept from 100 kHz to 40 MHz, to completely address the potential frequencies at which the inductor might see employment in a SMPS.

2.3.1.2 Results

The device-level characterization results for all four device variants are shown in Figure 26. These results confirm the successful operation of the devices. All devices showed appreciably higher inductance over the air core inductor in the targeted frequency range of 1-5 MHz. As well, the devices showed high Qs over the same range, approximately 1.7 for the full-circuit variants (all-permalloy and half-ferrite), but going as high as 5 for the half-permalloy device.

There are many unsurprising features in the results shown. Foremost, the all-permalloy device showed the highest inductance, as would be expected due to the fully-closed, high-permeability magnetic circuit. Likewise, the other full-circuit variant, the half-ferrite inductor, showed the next highest inductance. The lower permeability of the ferrite material, relative to that of electrodeposited permalloy, as well as the *de facto* air gaps introduced into the composite by the polymer matrix, would explain the half-ferrite device's lower inductance relative to the all-permalloy device.

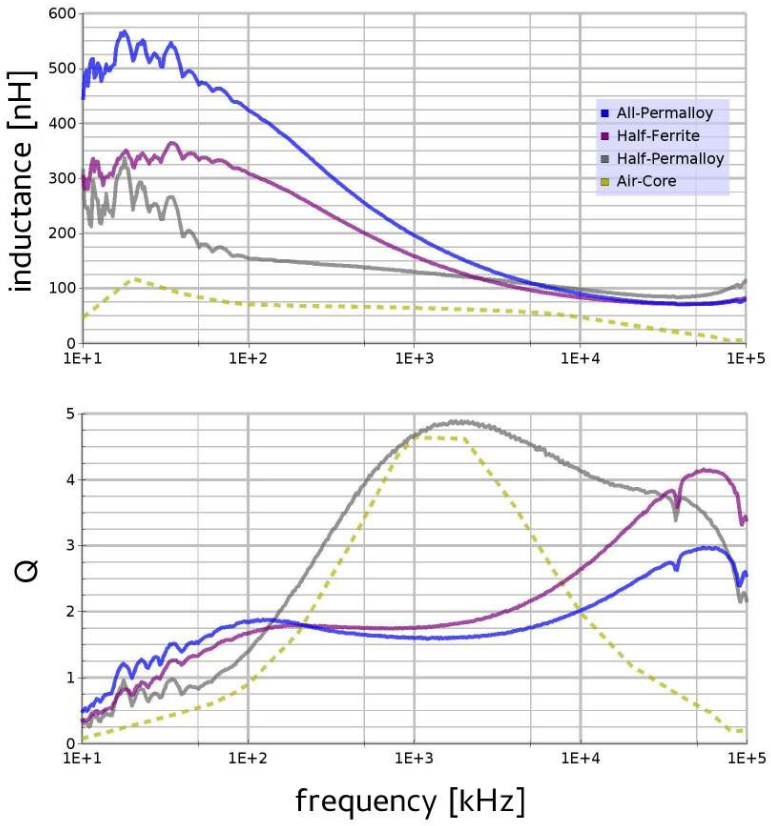

Figure 26: Microfabricated Inductor Device Characterization

Another expected behavior is the falloff of inductance in the full-circuit varieties in the 0.1 MHz to 10 MHz regime. The skin depth of permalloy can be expected to approach the thickness of the deposited layers (25 μm) at approximately 0.5 MHz. The apparent immunity of the half-permalloy device to this trend can be readily explained by its lower inductance at low frequencies.

It is expected that the air core device would offer better Q than the other variants, since it lacked any core material to incur substantial losses. However, one of the more surprising results seen in this chart is the high Q shown by the half-permalloy device. Since the half-permalloy device only offers modest inductance, it must be concluded that it experiences very little energy loss.

The half-ferrite offered a distinct tradeoff relative to the all-permalloy inductor, giving a simultaneous increase in Q and decrease in inductance. The attractive options within this tradeoff would depende heavily on the application. But, the considerably simpler fabrication of the half-ferrite inductor would clearly shift any such tradeoff considerations in favor of the half-ferrite inductor.

2.3.2 System-level Characterization

2.3.2.1 Overview

The intended application for this inductor was a single-die switched-mode power supply, and demonstration of its suitability in such a role is thus called for. However, complete implementation of a single-die system would entail the burdens of not only design and fabrication of a CMOS IC, but also successful handling of all related post-CMOS integration steps. The latter burden was avoided by employing a PCB-based prototype DC/DC power converter, shown in Figure 27.

The former burden – design, fabrication, packaging, and verification of a custom power converter IC – was handled by another doctoral student at Georgia Tech, Luke Milner, in Dr. Gabriel Rincon-Mora's research group. The CMOS IC provided not only basic system controls, but also a synchronous output stage incorporating optimized output driver transistors and tuned gate drivers to minimize total switching losses. The close coupling of this IC's design with the larger vision of a single-die power converter ensured that the characterization results obtained could be used to infer performance bounds for a truly-integrated converter.

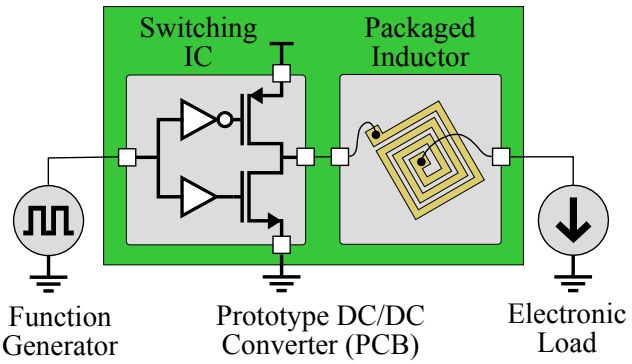

Figure 27: Prototype DC/DC Converter Utilizing Microfabricated Inductor. The inductor was packaged separately from the switching electronics, and a PCB provided overall system integration.

2.3.2.2 Device Packaging

For robust integration into the prototype DC/DC converter, the microfabricated inductor was packaged into a ceramic DIP "side braze" package. Figure 28 shows the packaged device. Without removal or damage to the microfabricated inductor, the silicon wafer was diced, and glued into the ceramic package. Aluminum bondwires were used – 6 in parallel for each inductor terminal – to

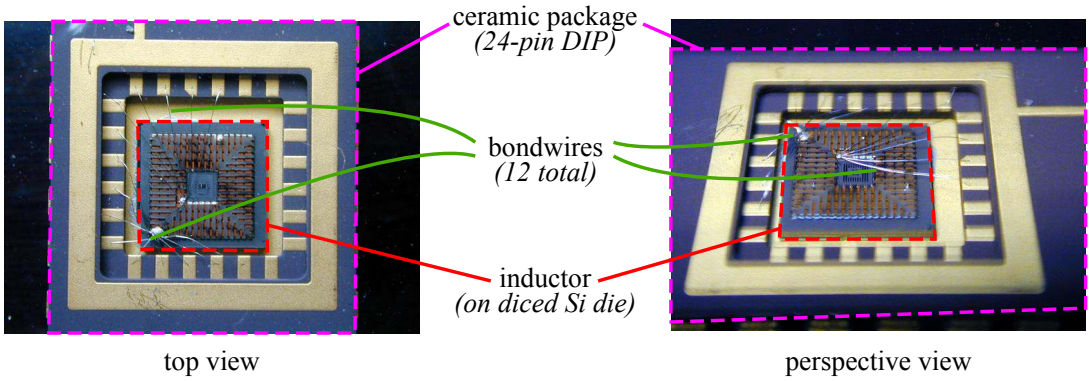

Figure 28: Finished and Packaged Microfabricated Inductor. The all-permalloy inductor, shown as packaged for use in the prototype DC/DC converter. A diced silicon die was epoxied into a ceramic DIP package and wire-bonded to the package leads.

connect the inductor to the package leads.

2.3.2.3 In-System Performance

The experimental setup used to operate the DC/DC converter and record the results is shown in Figure 29. A function generator was used to control both the duty cycle and the frequency of the transistors' switching; the IC's on-board control circuitry was bypassed to allow better control of switching parameters. An electronic load served as the converter's output load. This instrument provided not only output current monitoring but also protection against short-circuit and over-voltage conditions, both of which could damage the microfabricated inductor and/or the switching IC. The switching waveforms and DC output voltage were measured, respectively, with an oscilloscope and a digital multimeter. The DC output voltage was used to manually tune the switching duty cycle to obtain the intended DC voltage. The oscilloscope waveforms provided not only qualitative confirmation of the circuit's proper operation, but also quantitative indications of the inductor's effective inductance. A current probe was used to measure the output current with high temporal resolution.

The converter was operated at switching frequencies of 1 MHz and 5 MHz, with an input voltage of 3.3 V, an output voltage of 2.0 V, and currents up to as 2.5 A. The captured waveforms are shown in Figure 30. The highly-linear current waveforms indicate that the inductor provided a usefully large, highly constant inductance to the circuit. The small amount of rounding evident in the corners of the switching voltage waveforms indicate that the switching IC was functioning properly, and

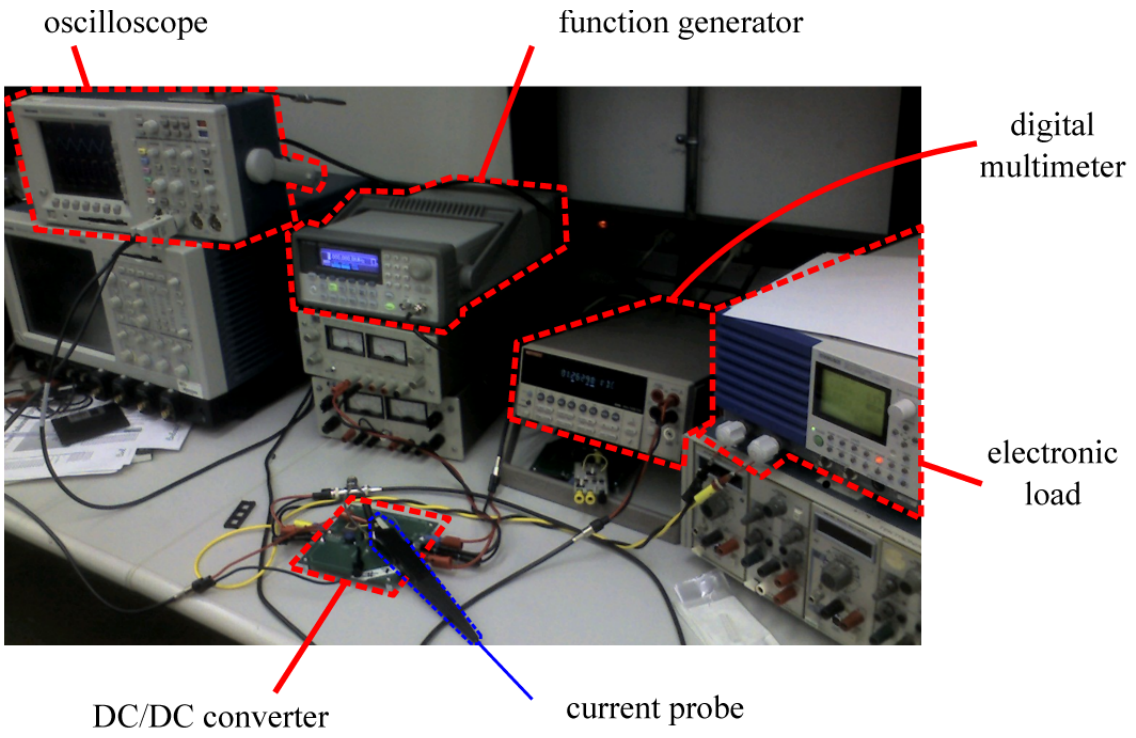

Figure 29: Experimental Setup for System-Level Characterization of Microfabricated Inductor. All equipment utilized in the characterization is highlighted.

that the microfabricated inductor was not presenting any catastrophically large parasitics. At both

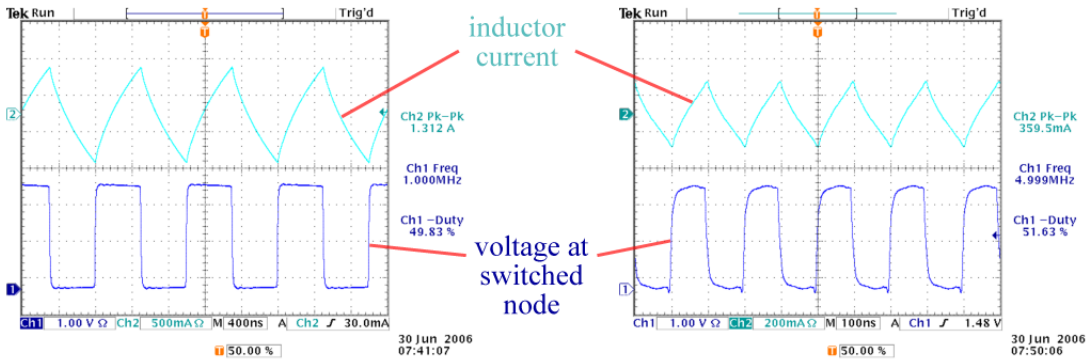

1 MHz switching frequency

5 MHz switching frequency

Figure 30: Switching node voltage and inductor current waveforms obtained from the DC/DC converter

switching frequencies – 1 MHz and 5 MHz – the load current was swept up to 2 A. The efficiency

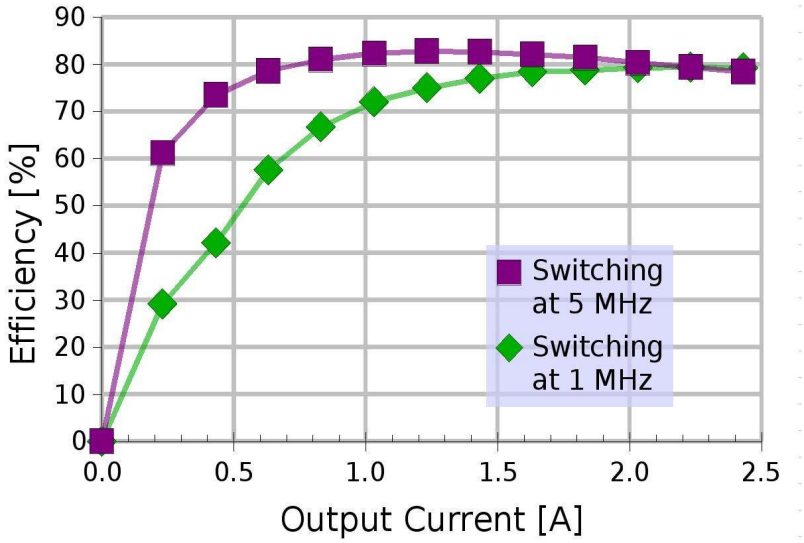

Figure 31: Measured efficiency of the DC/DC converter over frequency, and at 1 MHz and 5 MHz switching frequencies

at each load current value was calculated as ν in the relation

$$\nu = \frac{V_{in} I_{in}}{V_{out} I_{out}} \tag{22}$$

where V_{in} and V_{out} are the input and output voltages, I_{in} and I_{out} are the input and output currents, and the DC character of the input and output signals greatly simplify the calculation. The results are shown in Figure 31.

2.4 Conclusions

2.4.1 Eligibility of Microfabricated Inductors

Chief among the results demonstrated in this portion of the thesis is that a thoughtfully designed and carefully implemented microfabricated inductor can deliver very attractive performance in an integrated power converter application. As power converters move towards increased integration levels, the demand will certainly arise for filter passives that can support the move, from a fabrication process point of view as well as from a performance point of view. This inductor, and the more general approach of utilizing thick electrodeposited metal films, shows great promise in both regards.

2.4.2 Rewards of Maximizing Flexibility

Another important dynamic demonstrated in this work is the value of processing flexibility *per se*. The ability to explore the simplified device constructions (*cf.* subsection 2.2.2) without incurring significant processing costs, either in time or materials, proved very fruitful. Not only was the highest Q surprisingly achieved by the simplified half-permalloy device, but comparative analysis of the devices' performance gave strong insights into the fundamental operation and nonidealities of all devices. Further, it is very plausible that, depending on the application, one of the simplified variants may well prove a better fit than the full construction that was used in the DC/DC converter.

While it might have been possible to foresee some of the design variants during the initial design of the inductor, a general bias towards the type of design provisions that enabled these variants would have been just as effective and much less burdensome. A deeper exploration of the nature of such design provisions is beyond the scope of this work, but two key attributes can be extracted from this experience. First is the modularity of the fabrication process, with inter-module boundaries, such as the planarized surface between the first permalloy layer and the copper spiral, kept as generic and standardized as possible. Second is the alignment of the process module boundaries with functional boundaries of the device itself, as was seen in the division of the magnetic cores among two or three well-separated processes.

2.4.3 Critical Importance of Materials

The final major principle that is emphasized by the results obtained is the pivotal role that materials play in microfabrication. The thick metal layers, both for the magnetic core and the copper spiral, were undoubtedly key in achieving the prototype converter's high efficiency. These thick layers require not only robust and repeatable electrodeposition recipes, but also thick, high-resolution photoresists; without either element, the layers are not achievable. As well, the polymers that serve as structural and insulating elements, along with their putative PECVD-deposited replacements, also play vital roles in the device's overall functionality.

By the same stroke that these materials enable such good performance, their addressable nonidealities also become dominant roadblocks to improving the device's performance. As thick resist technology continues to develop, finer lateral laminations and thicker overall layers will become

possible, resulting in devices with higher inductances and lower losses. Improved magnetic alloy characteristics, such as lowered conductivity, thicker skin depth, or reduced hysteresis losses, will directly improve the device's performance.

CHAPTER III

SEQUENTIAL MULTILAYER PLATING

3.1 Overview

A method for producing high surface area structures has been developed, based on electrodeposition and sacrificial wet etching. This chapter details the extension of this method, through experiential learning as well as the development of custom equipment.

This process serves as the key fabrication technique for both capacitors (Chapter 4) and inductors (Chapter 5). This process can produce extremely high surface area structures of well-controlled geometry. This surface area can, with appropriate subsequent processing, be readily exploited to produce capacitors with high capacitance density.

This process can also produce highly laminated structures, comprising a large number of parallel, electrically isolated layers. If the layers are made of a magnetic alloy, such as permalloy, these laminated structures can be used as magnetic cores, and the electrical discontinuity between layers will serve to inhibit eddy currents. This reduces losses in the core and allows the use of magnetic alloys, which deliver high saturation flux densities and low hysteresis losses, in high-frequency power converters.

3.1.1 Approaches to High Surface Area

In many microscale applications, surface area plays an important role. For many transduction phenomena, rates or magnitude scale with surface area. [63] As well, the skin effect requires that high-frequency electromagnetic fields tend to occur only at the surfaces of conductors, so power-handling capacity should increase with surface area [75]. Consequently, microfabrication methods for creating structures with high surface area have intrinsic value. There are a large number of proven, widely-adopted approaches to achieving high surface area in various applications that demand it.

3.1.1.1 *Electrochemical etching*

One of the most robust means for creating high surface area is the use of electrochemical etching to alter the surface of an initially smooth structure, creating topologies with higher surface area. In aluminum electrolytic capacitors, for instance, an electrolytic etch is used to increase the surface area of a rolled aluminum foil. Examples of this are shown in Figure 32, along with the typical increase in surface area. The increase in surface area of the aluminum foil directly translates into a

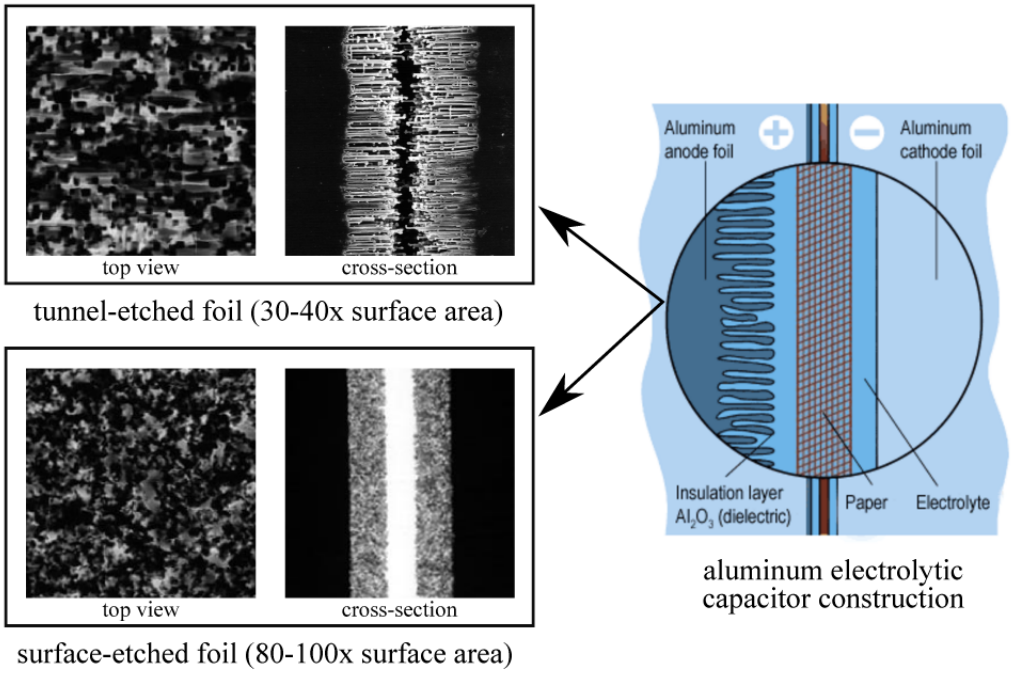

Figure 32: Etched foils for electrolytic capacitors (courtesy Nichicon Corp.)

congruent increase in capacitance.

A similar approach is also effective with silicon, resulting in the formation of what is known as porous silicon. In this process, the polished, optically flat top surface of a silicon wafer is transformed into an extremely high-surface area structure. An electrolytic etch, often in a solution of hydrofluoric acid, is used to initiate pores in the silicon surface and rapidly deepen the pores. [11] A canonical etching cell and a range of pore morphologies are shown in Figure 33.

Porous silicon has a wide range of applications. [17] Its high surface area combined with silicon's intrinsic sensitivity to adsorbed species make it particularly suitable for MEMS sensor applications. Figure 34 shows one such sensor, a biochemical sensor, based on porous silicon. [29]

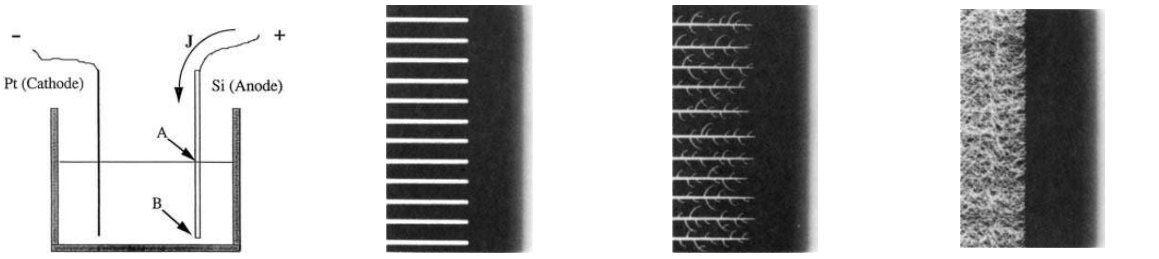

anodization cell for producing porous silicon

lightly *n*-doped wafer, backside illumination (cross-section)

heavily *n*-doped wafer, dilute HF (cross-section)

lightly *p*-doped wafer, concentrated HF (cross-section)

Figure 33: Porous silicon fabrication and cross-sections. Electrochemical dynamics play an important role in the process of creating porous silicon. [11]

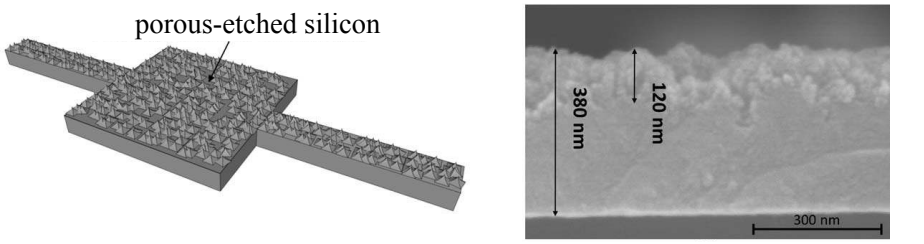

suspended-mass resonator

cross-section of porous-etched silicon

Figure 34: Resonant MEMS sensor based on porous silicon. The extremely high surface area and intrinsic sensitivity of porous silicon enhances the device's sensitivity. [29]

3.1.1.2 Deep RIE Patterning

Deep RIE etching, with appropriately designed etch masks, can be used to produce high surface area topologies in silicon wafers. The anisotropy of the deep RIE process allows the formation of very high aspect ratio structures. These high-surface area silicon structures are attractive as the foundation for silicon-integrated capacitors. This is not only because of the superb compatibility of silicon with further processing steps, particularly for the deposition of high-quality dielectric films, but also the great practical value that would come from the integration of high-density capacitors into conventional CMOS technologies.

Variants of this approach to achieving high surface area in silicon and subsequently achieving high capacitances can be found in [31], [6], (Figure 35) and [44] (Figure 36).

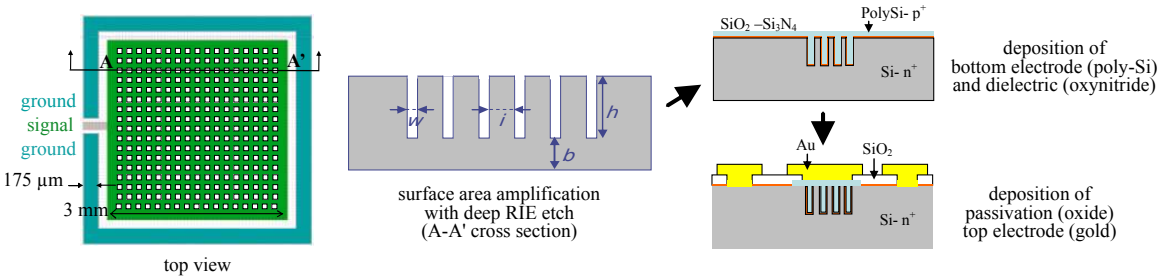

Figure 35: Deep RIE-based silicon capacitors [6]

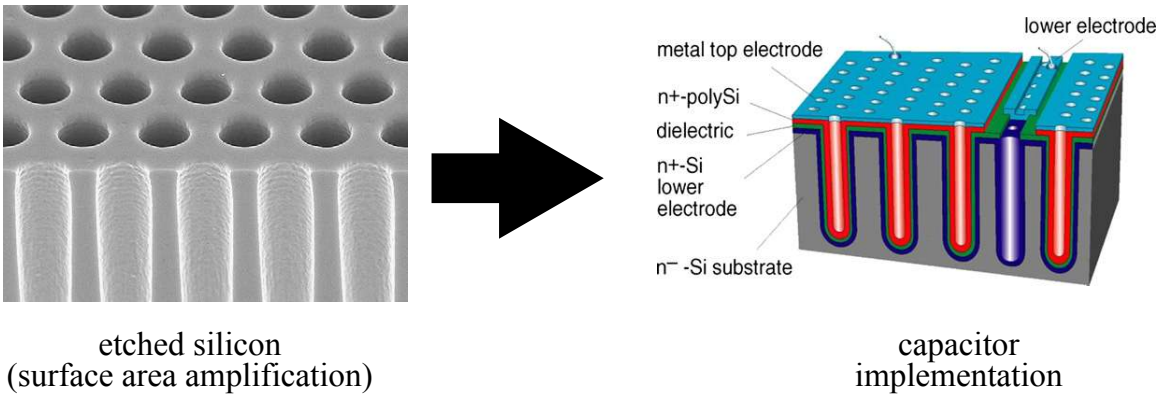

Figure 36: Deep RIE-based silicon capacitors [44]

3.1.1.3 Silicon-Templated Fabrication

Another emerging approach to forming high surface-area structures relies, like the deep RIE-based capacitors above, on the micromachining of silicon. However, in these techniques ([59]), the micromachined silicon is used as a mold through which a metal is plated.

One such application is shown in Figure 37. Note that photoelectrochemical etching, rather than deep RIE, is used to form the high aspect ratio topology in the silicon substrate. For geometrically simple structures where controlled etch depth is not important, such as the through-wafer vias used in [59], photoelectrochemical etching can provide a superior aspect ratio, and therefore superior surface area.

3.1.1.4 Thick photoresists

As the technology that enables thick photoresists progresses, the aspect ratio that can be achieved in the photoresist layer continues to improve. This results in substantially thick structures with high surface areas that can be achieved with a relatively simple, single-step process.

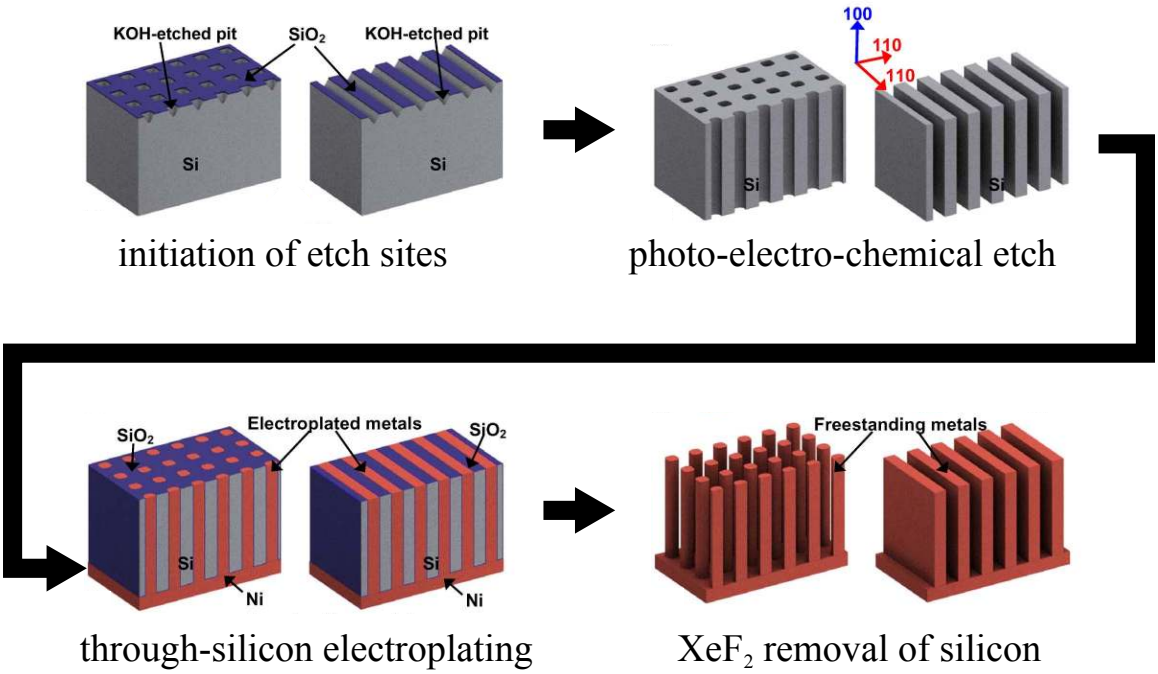

Figure 37: High-surface area metal structure formed by plating through silicon. Photoelectrochemical etching is used to achieve high aspect ratio holes in a silicon wafer, and electroplating is used to transfer this geometry to a metal structure.[59]

In [5], this approach is used to create high aspect ratio structures. These structures are then pyrolyzed to create surfaces composed mostly of carbon. These virtually carbonized electrodes offer apt performance in the role of electrochemical double layer capacitor (EDLC) anodes as well as battery anodes.

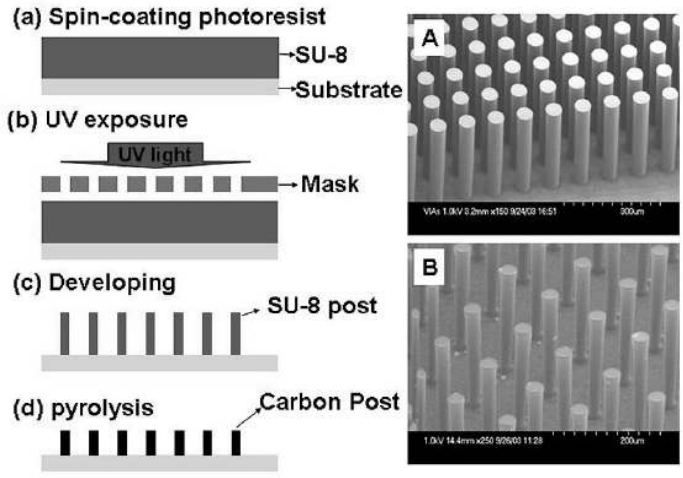

Figure 38: Photolithography-based high-surface area fabrication [5]

3.1.1.5 Nanoscale Approaches

Discoveries and innovations from the field of nanotechnology can also be applied to produce high surface area structures. One example of this is [13], in which carbon nanotubes are used as additive components to form a structure with high surface area. The completed structure is used as the anode of an EDLC. (Figure 39) Another example is the development of silicon nanocones [70], where high aspect ratio tapered structures, called nanocones, are formed directly at the silicon surface. (Figure 40)

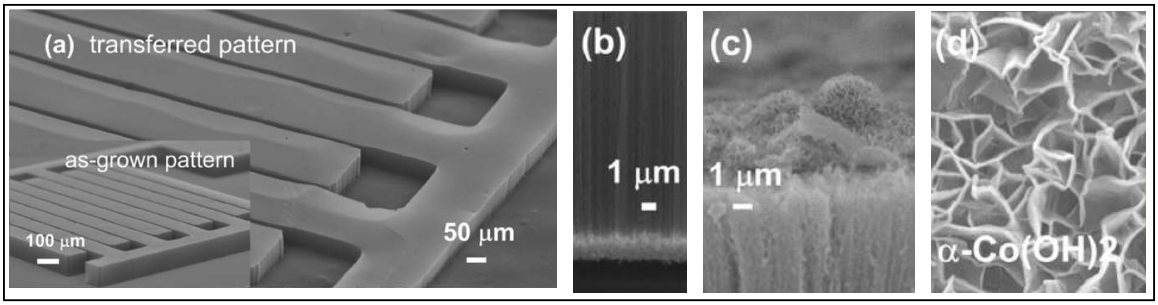

Figure 39: Carbon nanotube-based capacitor. Carbon nanotubes are deposited in an interdigitated pattern. Then, the nanotubes are decorated with cobalt hydroxide flakes to further increase the surface area of the conductive structure. [13]

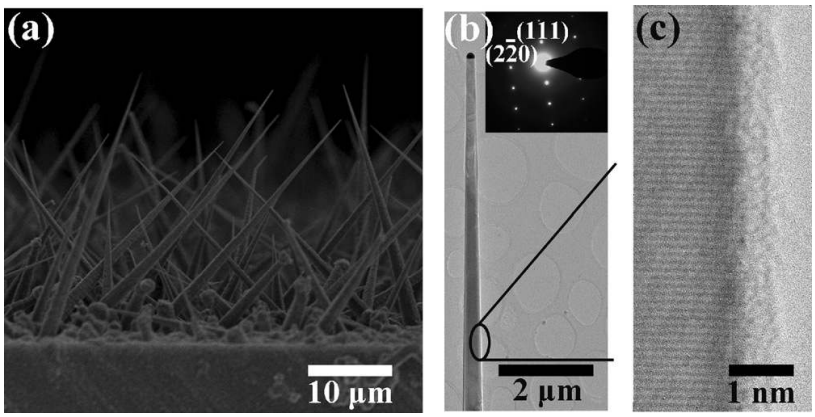

Figure 40: Silicon nanocones. [70]

3.1.1.6 Virus-Templated Assembly

One additional vein of research into high surface area structures borrows advances from biology. The tobacco mosaic virus (TMV) has a cylindrical shape approximately 300 nm in length and 40 nm in diameter, making it an excellent implement for nanoscale structure creation.

As well, the TMV has been genetically modified to control the chemicals present on its outer surfaces in ways to facilitate its application. At one end of its body, the TMV exudes a protein residue that causes the virus to adhere strongly to gold surfaces. Meanwhile, the TMV's entire outer case is highly conducive to nucleation of metal coatings, and is thusly compatible with electroless nickel processes. Typically, TMVs are induced to attach to one or more surfaces of a microstructure, after which they are coated with electrodeposited metal. In [42], [22], and [15], silicon is electrodeposited onto the metal-coated viruses to form a high-capacity battery electrode. This process is illustrated in Figure 41.

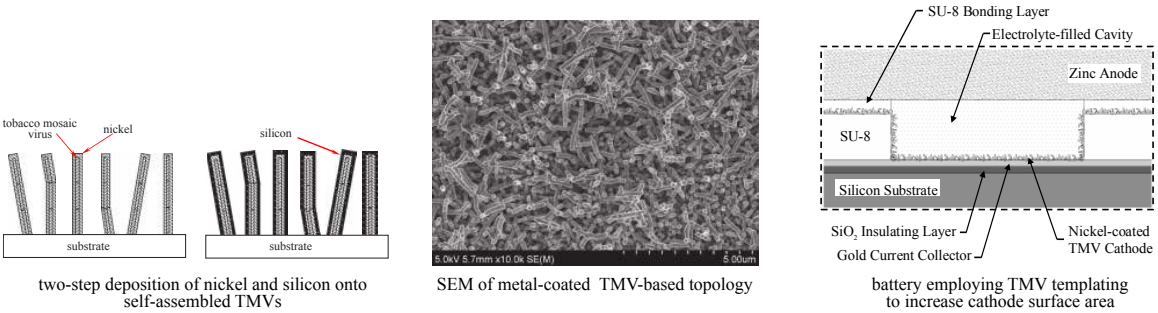

Figure 41: Tobacco mosaic virus (TMV)-based high surface area structures. The TMV has been genetically modified to attach, by its end only, to the substrate, as well as to have outer walls capable of nucleating electroless metal deposition. [15] and [22]

3.1.2 Sacrificial Multilayer Plating

The remainder of this thesis is based on a recently-developed approach to high surface area. The essential component of this approach is the micromachining process illustrated in Figure 42.

In this process, a large number of relatively thin layers of two different metals are electroplated through a plating mask, in an alternating sequence. This produces a metal structure whose height is equal to the sum of the layer heights, and whose composition is laterally uniform but vertically alternating according to the composition and thickness of the layers. This metal structure is then subjected to an etch process that selectively removes only one of the two metals, leaving behind a large number of fully-separate layers of the other metal.

It should be pointed out that this process, along with its great potential, presents two *prima facie* hazards. The first is that the structural layers – the layers that are not removed by the sacrificial etch – lose their mechanical support as the sacrificial material is removed. Generally, this means that

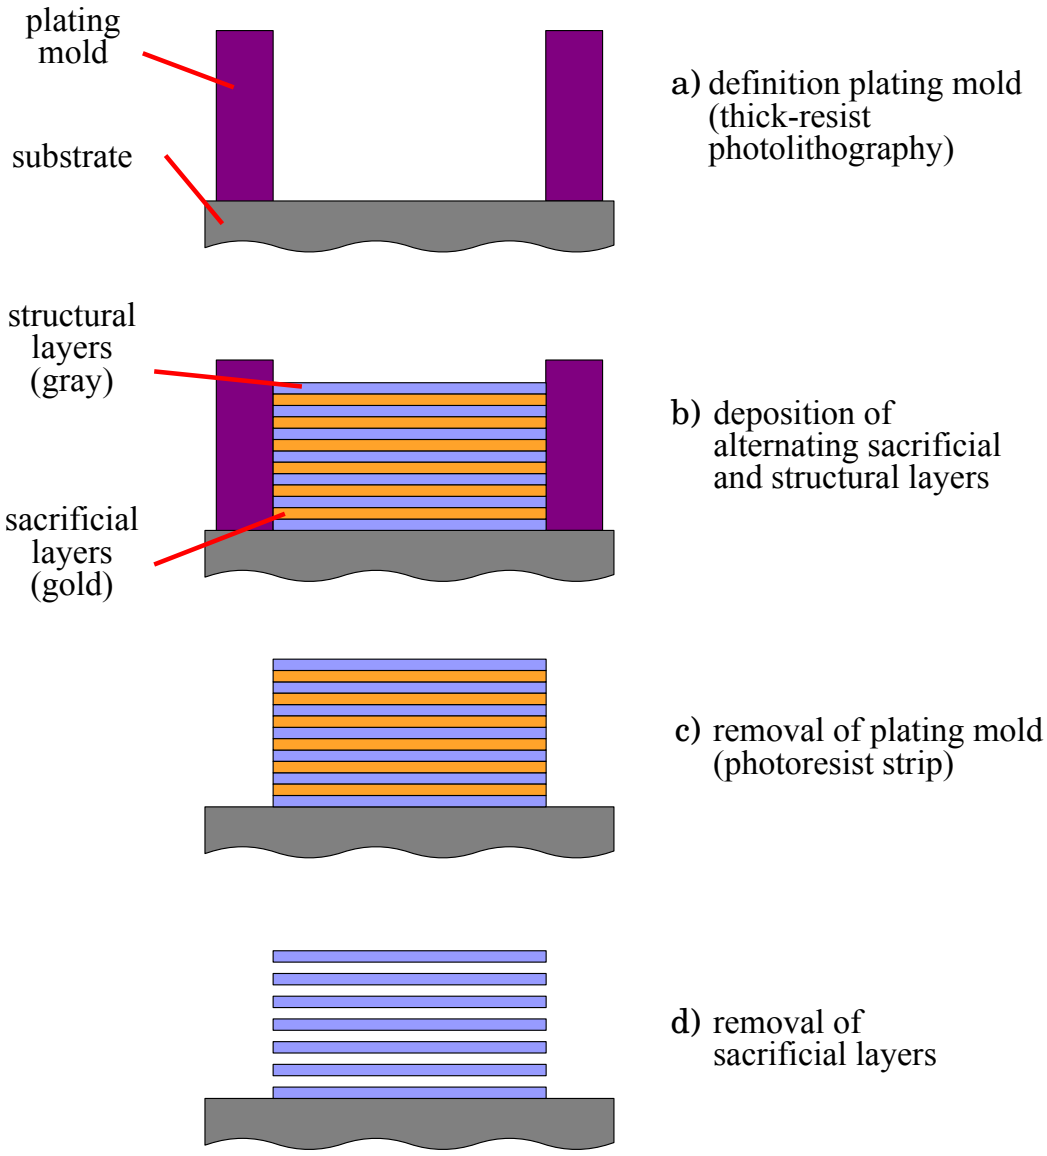

Figure 42: Generic sacrificial multilayer plating process.. Many layers of metal, including some intervening sacrificial metal layers, are deposited. Then, the sacrificial metal layers are removed without affecting any other layers (*i.e.* the structural layers). What remains is a vertically-laminated, high surface area structure formed by the surviving structural layers.

the structural layers are liable to fully detach from the substrate, unless otherwise anchored. The second basic hazard that must be considered is the extreme fragility and compliance of the structural layers after release. This presents an array of secondary needs, including mechanical protection of the finished structure and prevention of stiction during and after the sacrificial release.

3.1.2.1 Process Invention

The sacrificial multilayer fabrication process was invented in 2003 by Dr. Jin-Woo Park, a graduate student under Dr. Mark Allen at the Georgia Institute of Technology. It was developed as a means to implement ultra-fine lamination of magnetic cores, to produce inductors capable of operating at very high switching speeds. As discussed in Chapter 2, lamination is a well-proven technique to reduce eddy current losses in magnetic cores. The process used in Dr. Park's research is shown in

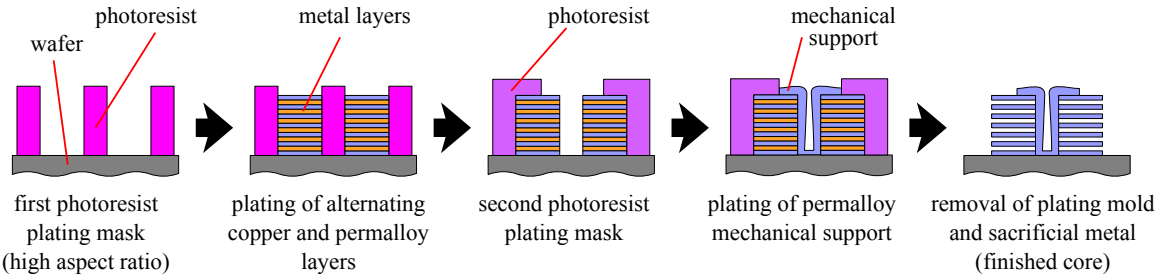

Figure 43: Multilayer process for laminated magnetic cores (Park)[46]. A gullwing-like structure is deposited prior to release to provide mechanical support for the finished layers.

Figure 43. The layers are anchored, prior to release, by a gullwing-like mechanical support that is electroplated on top of the finished stack. The application goal of eddy current suppression placed considerable constraints on the design of the mechanical support. As this support provides electrical contact along with its mechanical attachment, an improperly-placed support could provide a low-impedance path for eddy currents, despite the effort placed into achieving the ultra-thin laminations.

The high compliance of the plated layers was addressed by the dimensions chosen for the total width of the core. Limiting horizontal extent of the layers to 1-3 mm proved sufficient to avoid these problems.

3.1.2.2 First-Generation Plating Robot

Urvi Shah, another graduate student studying under Dr. Mark G. Allen at the Georgia Institute of Technology, advanced the state of the art of sacrificial multilayer plating in 2008. Ms. Shah not only

achieved layer counts in excess of 160 and layer thicknesses under 1 μm , but she also constructed a fully-automated robot to execute the multilayer plating process.

However, it was found that this first-generation robot was not compatible with long processing times and the frequent use that was demanded by the research in this thesis. A higher degree of repeatability was also required than could be provided by the first-generation robot. This was especially true for the electronics that drive the electrolytic deposition process, where finer control over plating current and precise control over process timing was needed to dependably create laminations far below 1 μm in thickness.

3.1.2.3 Motivation for the Second-Generation Plating Robot

Accordingly, a second-generation was developed. This new robot was based broadly on the first-generation system, and certainly incorporated the practical knowledge gained from the construction and operation of the latter. However, to achieve this work's targeted levels of robustness and repeatability, nearly every subsystem of the robot needed to be rearchitected and reimplemented, as did most of the mechanical components. The *x*-stage mechanism was the only system component that was directly inherited from the first-generation

In the new system, all software and electronics were implemented using microcontroller-based PCBs, rather than the protoboard-based, LabView-driven approach used in the first-generation robot. As well, radical new designs for the plating tanks as well as the *z*-stage mechanism were introduced. A closed fluidic system was also added to the second-generation robot not only to facilitate the handling of the large volume of processing liquids (*viz* plating baths), but also to maintain the quality of these liquids both during and between plating runs.

The details of the design and implementation of the second-generation robot are presented in this chapter, as well as in Appendix C.

3.2 Multilayer electroplating robot

Implicit in the sacrificial multilayer fabrication technique is the use of multiple electrolytic deposition reactions. Each deposition reaction entails its own bath chemistry as well as electrode materials. Along with the needs for precise control over plating times and currents and very long total processing times, this clearly creates the need for an auotmated electromechanical plating system.

3.2.1 Architecture

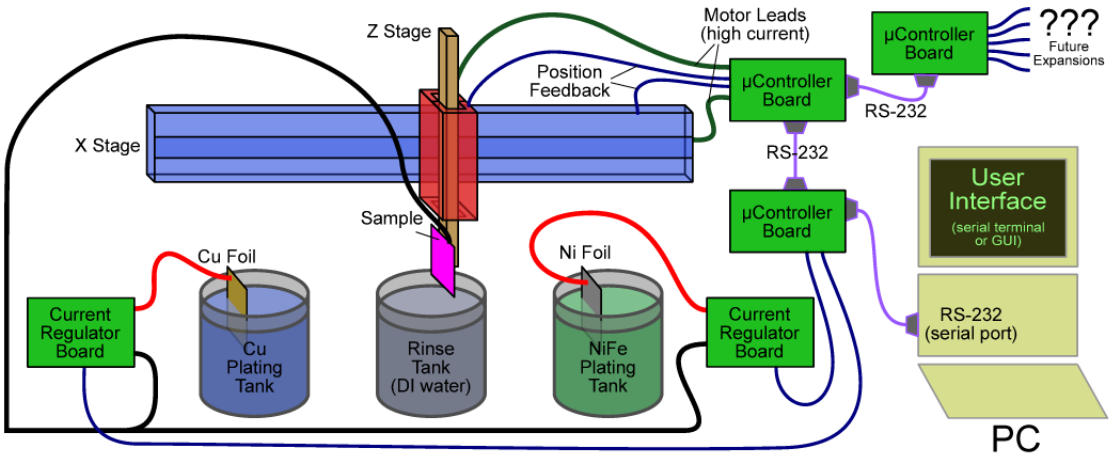

Figure 44: Automated Multilayer Electroplating System Architecture

Figure 44 shows the architecture followed for the multilayer plating robot. There are three main elements to the system: the mechanical assembly, the fluidic system, and the control electronics.

The mechanical assembly consists of the two mechanical stages (x and z), the four process tanks, and all brackets and frames supporting the stages, tanks, wafers, and anodes.

The fluidic system provides for the movement and storage of the plating baths, as well as the management of nitrogen gas flows, used to deoxygenate the baths during plating as well as storage. The baths are kept in large carboys, and moved between the carboys and the process tanks by peristaltic pumps. Nitrogen gas is bubbled through all process baths primarily to prevent oxidation of aqueous iron ions, but also to provide gentle agitation.

The control electronics actuate the mechanical assembly, control the plating processes, and execute the overall fabrication sequence. The control electronics consist of four separate circuit boards, each responsible for a separate task, but all operating in concert. The top-level process control board, known as the "sequencer," is located in its own enclosure and was, in general, separated as completely as possible from the rest of the control electronics. This was done to allow for easy substitution of other sequencer implementations, perhaps according to the preferences of individual robot users. To minimize the chances of damage to the control electronics from the harsh chemicals used in the plating process, all control electronics were packaged in closed boxes and connected to the rest of the system by lengthy cables.

3.2.2 Electronics

Each of the four control boards, being responsible for a single task, contained a single microcontroller: the Cypress "Programmable System on a Chip (PSoC)" CY8C29466. The decision to endow each task with its own processor was based on the limitations of the selected microcontroller. These limitations included not only processing and memory capacity, but also the on-chip digital and analog hardware blocks, such as PWM generators or A/D converters. The primary burden introduced by this multiprocessor architecture was communication.

3.3 Process Extension

The sacrificial multilayer plating process had previously achieved successful fabrication of magnetic cores. [56] [46] However, the full impact of this technology is believed to be as yet unrealized, consisting not only in further refinement of its application to magnetic cores, but also in its application to new types of devices.

Achieving these goals requires the extension of the sacrificial multilayer plating process along several dimensions. These are discussed in this section.

3.3.1 Thicker structures

Increasing the overall thickness of these multilayer structures is an important aspect of expanding their range of applications. For magnetic cores, power handling capacity scales with the core's cross-sectional area, which in turn scales with core thickness. As well, surface area will scale with structure thickness. Similarly, applications which rely on surface phenomena, such as capacitors or heat exchangers, will benefit from thicker structures.

3.3.1.1 Thicker molds

One straightforward, yet challenging, requirement for producing thicker multilayer electroplated structures is correspondingly thick plating molds. Previous multilayer plating work used molds up to approximately 100 μm , and roughly an order of magnitude increase in thickness was targeted.

NR-21 molds

Futurrex NR-21 was used to make the plating molds for most of the structures presented in this

thesis. However, this material proved troublesome at thicknesses over 150 μm . Substantial complications were found in both the soft-baking and the exposure of the material at these thicknesses.

However, at more conventional thicknesses, particularly 50 μm and 100 μm , NR-21 offered a compelling baseline option. Its appeal came not only from its relatively short and simple lithography process, but also from its robust release, post-plating. Acetone was found to be an excellent solvent for removing NR-21. Although the acetone did not fully dissolve the photoresist, it was very effective at removing the photoresist's adhesion to the sample.

However, on occasion, during acetone release steps, large particles of NR-21, diameter 250 μm or larger, were found to re-adhere to the plated metal after being initially detached. Subsequently, removal of these large particles proved impossible, even with the help of the vendor's recommended stripper, RR-41, heated to 50 C.

This recurring problem was solved by beginning the post-plating strip with the wafer submerged in trichloroethylene (TCE) and slowly adding acetone into the TCE. TCE was not found to attack the photoresist, and so it was used to buffer the acetone's action while providing adequate fluid convection to remove the photoresist debris as it was detached. This TCE/acetone photoresist removal bath was always started as 100% TCE, and acetone was added until the volume roughly quadrupled, *i.e.* until it the TCE composition dropped to around 25%. Adding this amount of acetone at a constant rate over a period of 12-15 minutes produced the best results.

Other than NR-21's thickness limitations, it was also judged not suitable for very long multilayer plating runs because it showed some evidence of chemical vulnerability to the electroplating baths. This vulnerability manifested as two distinct phenomena. First, the photoresist, at the end of plating runs longer than 10 hours, tended to show slight discoloration of all portions that were submerged into the plating baths. This was almost always seen, but seemed to worsen with the aging of the baths. Second, the submerged portions of the photoresist would wrinkle and eventually delaminate from the wafer. This, too, seemed to be exacerbated by the aging of the plating baths, but, even with new baths, it limited the effective run time of any NR-21 plating run to roughly 24 hours.

SU-8 plating molds

SU-8 was investigated as an alternative plating mold material. SU-8 is well-known for its ability to produce thick, high-aspect ratio structures, as well as for its chemical durability. Clearly, these

two strengths directly address the shortcomings of NR-21 in this application.

SU-8 performed very well as a plating mold, even for plating runs up to 36 hours. Additionally, it was found that the edge profiles of the structural layers were much smoother when deposited into an SU-8 mold than when deposited into an NR-21 mold.

The principal difficulty with SU-8 was its removal after the plating process. It was found that boiling the wafers in deionized water for 8-10 hours was sufficient to release most of the SU-8 from the wafer. However, mechanical assistance was usually necessary to bring the SU-8 removal process to an acceptable degree of completion.

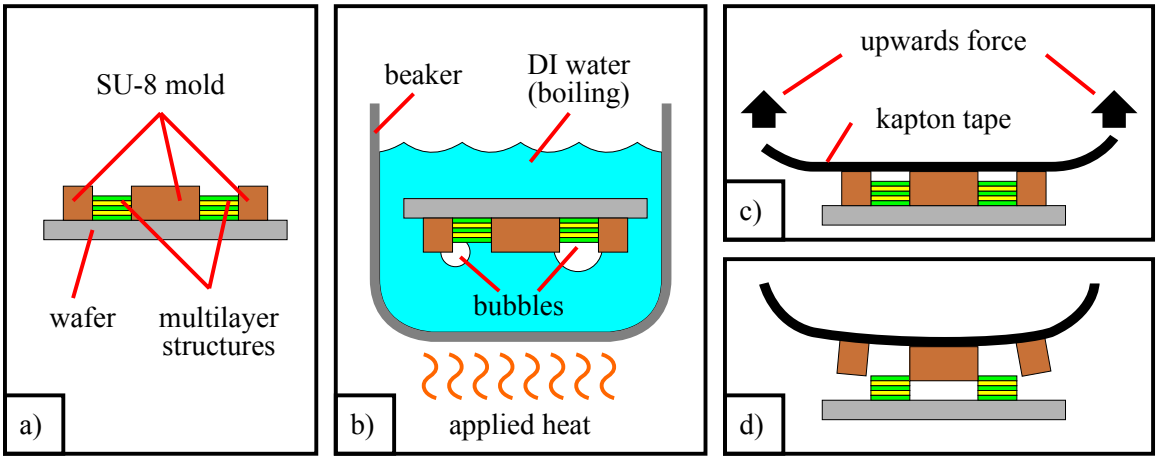

Figure 45: Hydrophilic removal of SU-8. After several hours of submersion in boiling deionized water, SU-8 can be removed from the wafer with gentle mechanical assistance (frame d). Note that the formation of bubbles on the wafer (held inverted in the boiling water) likely accelerates the penetration of water into the SU-8 and the initiation of delamination between the SU-8 and the wafer.

The SU-8 removal process used for this research is shown in Figure 45. The mechanisms responsible for this removal were not well understood, but it seems that the formation of bubbles during the deionized boil played an important role. Presumably, the vapor pressure inside the bubbles not only drove the diffusion of water molecules through the volume of the SU-8, but also initiated separation between the sidewalls of the SU-8 mold and the plated structures, allowing further penetration of water down to the wafer surface. For this reason, the wafers were oriented upside-down during the boiling step, as shown in Figure 45.

3.3.1.2 Large-scale, long-run plating

Beyond producing microfabricated plating molds of increased thickness, it was found that a number of enhancements to the multilayer plating process were required to make use of these molds. Simply depositing enough metal to fill a 100 μm -thick mold, at the typical rate of 10 μm per hour, would require 10 hours of run time. Filling a 100 μm -thick mold with this multilayer plating process, however, incurs the extra time required for steps that must occur between each deposited layers, such as rinsing and the physical transfer of the wafer between tanks. For 1 μm layers, with a corresponding per-layer plating time of 6 minutes, these interstitial steps increase the run time by roughly 50%. Further, the time consumed by these interstitial steps, as well as any resultant cross-contamination between tanks, will, for a fixed mold thickness, scale inversely with the layer thickness.

Therefore, extension of the multilayer plating process to create very thick structures comprising very thin layers will require not only very long run times, but also measures for dealing with the undesirable side effects of the inter-layer processing steps.

Tank geometry

The geometry of the plating tanks was adapted to fit these needs in two ways. First, the volume of the tanks was made as large as practically possible. This was done to generally promote invariant bath composition, but, more specifically, it was a measure to prevent depletion of certain ion species within the baths. For instance, in the nickel-iron bath used, while the nickel atoms that deposit onto the sample are replaced by atoms dissolving off of the nickel foil, there is no source of iron atoms (ions) to replace those that deposit. The travel of the plating robot's x stage imposed one limit on the volume of the baths, while the available height for the tanks imposed another limit.

Two other major factors influencing the tank geometry were: the need to allow for nitrogen aeration from the bottom of the tank, and the desire to minimize contact between the plating bath and the ambient air. Both of these needs were driven by the pronounced tendency of the iron ions in the nickel-iron bath to oxidize, but they also served the broader goal of maintaining the bath's condition throughout the long runs. These factors favored tanks with high vertical aspect ratios, which would allow for maximum volume while minimizing the contact area with the ambient environment.

The final tank design, with emphasis on the required characteristics discussed, is shown in

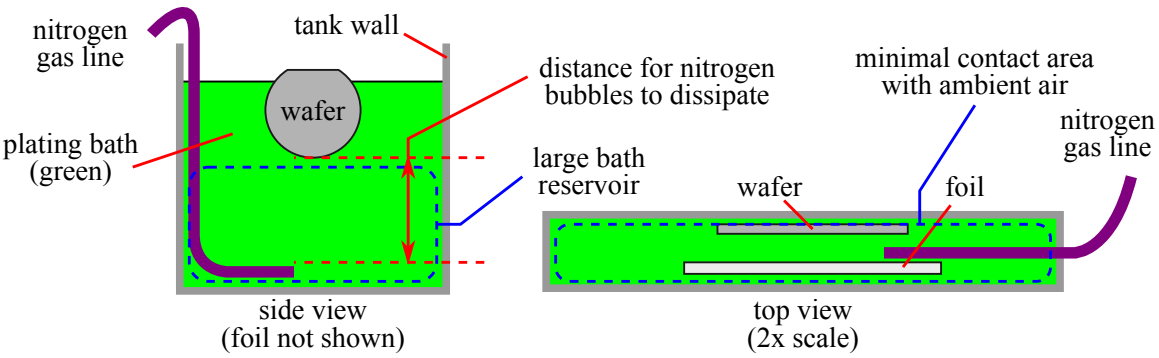

Figure 46: Guiding concerns in plating tank design. The demands of executing long plating runs with the automated multilayer electroplating system translated into a set of desirable features in the plating tanks. One lateral dimension must be minimized to accommodate limited travel range, while the total volume of the tank must both be maximized and made accessible to nitrogen sparging.

Figure 46. It was found that 5 liter polyethylene rectangular-shaped jugs (part number 73062, US Plastic Corp., Lima, Ohio, http://usplastic.com) , with one wall removed, fit these design goals very well.

Bath maintenance

As presented in the previous section, constant-flow nitrogen was used to prevent oxidation of the plating baths during plating runs. While this imposed a requirement on the tank design, this was also part of a larger need to support long runs: the need to maintain the baths' conditions, not only during long runs, but also over the weeks and months between runs.

This need was a consequence of the fact that large volumes of baths were needed during runs. In light of the cost of the raw materials (metal salts) used to make the baths, as well as the manual labor likewise involved, it was judged impractical to frequently create new baths.

In practical terms, this need was met by a number of key provisions. First, the baths were stored in sealed carboys between runs. Second, the baths were passed through 0.6 μm filter paper every time they were transferred from the plating tanks back into the carboys. Third, a small amount of nitrogen gas was bubbled through the carboys during storage, to keep the baths free of dissolved oxygen. Finally, the baths' overall condition and volumes were monitored and adjusted during storage; specifically, deionized water was regularly added to the baths to compensate for any observed evaporation.

Rinse water acidification

One largely unanticipated need in the extension of the multilayer plating process to longer run times was the maintenance of the rinse tanks' contents. A small amount of orange precipitate, presumably iron oxide, was observed in the nickel-iron rinse tank after long runs. It was found that the presence of this orange precipitate correlated with the presence of defects on the plated structures.

Two measures were put in place to prevent the formation of this orange precipitate in the rinse tank. Each measure by itself was insufficient, but when deployed together, they proved to be a reliable solution. First, the deionized rinse water was de-oxygenated by a vigorous flow of nitrogen gas for 2-3 hours before starting the run. Also, a small amount of boric acid (HBO_3) – approximately 25 g – was dissolved into the deionized rinse water during the nitrogen bubbling.

3.3.2 Thin layers

The main motivation to pursue thicker multilayer structures was to achieve higher surface area within the same substrate area. This same motivation also dictates the pursuit of thinner layers; with thinner layers, and an invariant limitation on total structure height, a larger number of layers and, consequently, a greater total surface area, could be achieved. A number of approaches to creating thinner layers within the sacrificial multilayer plating process were taken.

There were two issues that directly limited the minimum achievable layer height: layer integrity and thickness uniformity. Layer integrity refers to the lateral continuity of the plated layer. For this process to work, in general, each layer needs to fully intervene between the adjacent layers. Incomplete sacrificial layers would create contact between neighboring structural layers, not only likely reducing the total surface area and exacerbating any mass transport issues, but also directly hindering application-specific performance. An example of the latter effect would be electrical shorts between layers reducing eddy current suppression.

Thickness uniformity of the layers in this process is particularly important for a few reasons, and thus presents a chief limitation. Firstly, any thickness nonuniformity would be propagated upward and accumulated through the layer stack. Further, any correlation in thickness nonuniformity between layers would only speed up the accumulation. Secondly, given the minimum thickness at

which a continuous layer develops, any nonuniformity in thickness entails the presence of unanticipated structural material. In certain applications, this can be detrimental. For instance, in eddy current suppression, increased layer thickness generally amounts to increased eddy current losses.

3.3.2.1 Electrical

One direction of approach to thinner layers was through the control of the electrolytic plating currents. The highly-flexible plating current controller implemented in the plating robot allowed a great deal of flexibility in this regard. In addition to precise control over the current amplitude traveling between the plating foil and the wafer, it also allowed for three temporal variations:

- a turn-on delay, allowing for the Helmholtz double plane at the layer deposition surface to fully develop before plating commences;
- pulse plating, allowing insertion of brief periods during which no plating current flows, to accommodate slow mass transport dynamics and recovery of the initial Helmholtz double plane; and
- reverse pulse plating, allowing the inclusion of similarly brief periods during which a reversed-direction plating current flows, with the intention of producing smoother plated surfaces.

All three of these features were investigated to help minimize layer nonuniformity and reduce the minimum continuous layer thickness. The initial zero-current period was found to be very helpful in achieving uniform initial deposition and lowering the time required to achieve a continuous layer. Pulse plating produced no discernible improvement in the layer uniformity, and its putative amelioration of mass transport issues was outweighed by the increase in plating time it mandated. For unrelated reasons, such as the apparent chemical vulnerability of the photoresist to the plating baths, the total time available to deposit a given multilayer stack was sharply limited. Reverse pulse plating was found to produce very high levels of particulate contamination at the plating surface, leading to extreme roughness and nonuniformity as subsequent layers deposited atop these particles. It is believed that these particles were the result of the reverse pulse current effectively removing material from the plated surface (*i.e.* reversing the plating process), but the apparent tendency of

the removed material – either as fully dissolved ions or as micron-scale particles – to linger near the plating surface, rendered this approach unsuitable.

3.3.2.2 Smooth copper

A chemical-based approach to achieving thinner layers was also taken. Since copper is a very popular material for electrodeposition, there is a wide array of chemical additives available for achieving uniform, smooth deposition. An attempt was made to leverage this readily-available technology by using a decorative copper plating bath for the sacrificial layers, in place of the simpler sulfate-based bath used in previous work. The "Clean Earth™ Mirror Copper" plating bath, part number 45.209 from Grobet USA (Carlstadt, New Jersey, http://http://www.grobetusa.com) was found to produce very smooth, highly-uniform copper layers.

There were some minor disadvantages encountered in leveraging this commercial technology, beyond its higher cost. First, the bath was found to be less stable than the sulfate bath, frequently producing black solid precipitate that not only contaminated the plating tanks, but also often completely clogged the bath filters. Second, likely due to a chemical interaction between the chemical composition of the smooth copper layers and the nickel-iron bath, discontinuous nickel-iron layers were also sporadically encountered.

3.3.3 Etch optimization

3.3.3.1 Fixturing

One problem that became apparent early in the process extension effort was the importance of mass transport during the etch process. It was found that appropriate fixturing was an effective aid in addressing these mass transport challenges. Figure 47 shows a typical symptom that was consistently found across a large number of samples. The apparent etch rate near the substrate, at the bottom of the multilayer stack, was found to be substantially less than at higher elevations within the stack. In these cases, the etch was performed with the substrate laying flat, and right side up, on the bottom of the blue etch container.

It was postulated that this was due to the pooling of relatively heavy etch products in the voids between the sidewalls of the plated structures. Blue etch is a saturated solution of copper sulfate in ammonium hydroxide, so it is plausible that increases in the concentration of dissolved copper and

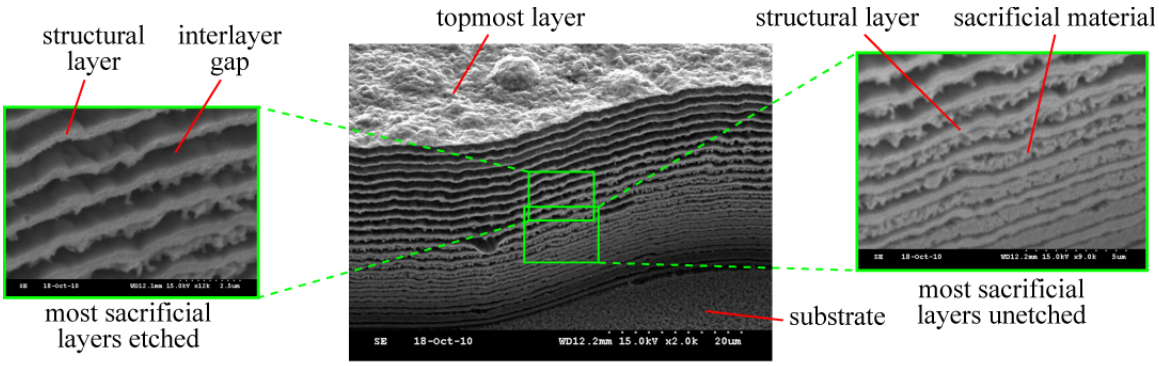

Figure 47: Etch rate dependence on height. This picture was taken after 4 minutes in blue etch. The apparent decrease in etch progression at lower layers is consistent with a convection-dependent mass transport limitation

the consumption of ammonium ions in the complexing of newly-dissolved copper ions would bring the solution to the point of precipitation. Indeed, small amounts of solid precipitate were frequently observed on the substrate and the bottom of the blue etch container. It is also plausible that the blue etch solution at the precipitation point would possess a higher density than the bulk of the blue etch solution, and therefore tend to settle towards the bottom of the container. The hypothesized situation is shown in Figure 48.

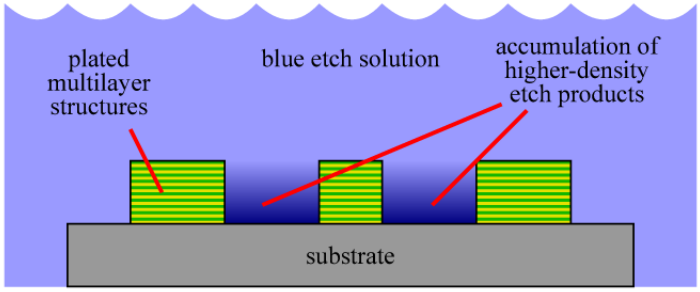

Figure 48: Hypothesized pooling of blue etch products. It was hypothesized that the blue etch products, likely saturated with copper, were heavier than the bulk of the blue etch solution and therefore prone to accumulating near the substrate surface, producing slower etch rates near the substrate.

The solution to this inevitable phenomenon was to alter the orientation of the substrate. The fixture shown in Figure 49 was developed for this purpose. It provided for the fixation of the samples in an upside down orientation, so that the downward flow of the heavier etch products would carry them away from the sample, enabling faster and more uniform etching.

Further, this fixture provided a mesh layer onto which samples could fall, in cases where they

were fully released from the substrate. The mesh layer provided mechanical support of the sample while still allowing the etch reactants to naturally convect away from the sample. With the mesh layer, it was also possible to safely remove the fully-released samples, consisting of extremely delicate layers, from the blue etch without incurring the extreme damage that conventional manipulation implements, such as tweezers, would have incurred.

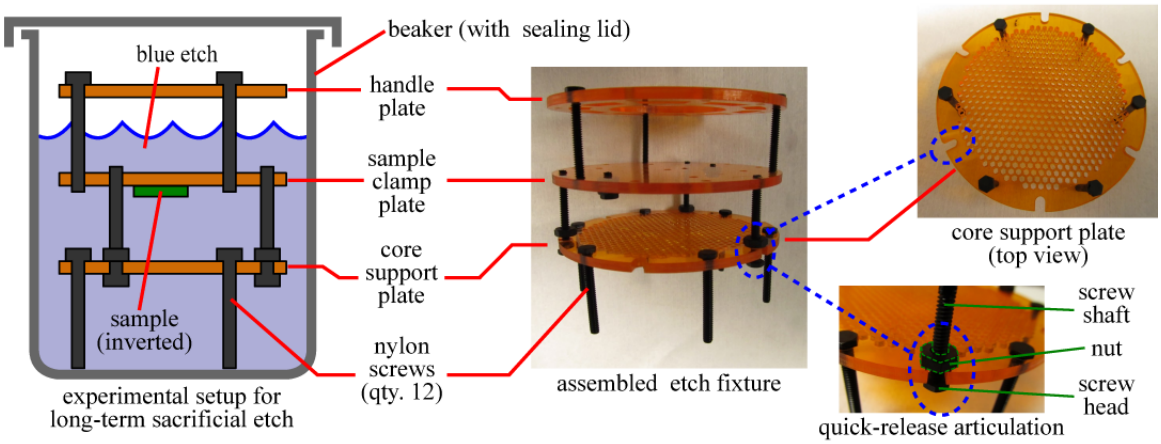

Figure 49: Custom fixture for sacrificial etching. The mesh plate (core support plate) proved a critical feature.

This fixture was found, in initial testing, to alleviate the vertically uneven etching seen in Figure 47. However, as the release process for the laminated magnetic cores – the principal vehicle for the development of this research – involved removing the multilayer structures from the substrate *prior* to the sacrificial etch, the upside-down sample fixturing was not used extensively. Nonetheless, the bottom mesh plate was used in the release of all cores, including the devices presented in Chapter 5 and Chapter 6.

Despite the great utility of this fixture, some potential improvements are obvious upon inspection. First, the fixture was formed entirely from laser-cut poly methyl-methacrylate (PMMA). While PMMA showed no signs of vulnerability to begin submerged in blue etch for many days continuously, PMMA was found to be extremely vulnerable to both acetone and methyl ethly ketone (MEK), and also moderately incompatible with isopropanol. This incompatibility with the organic solvents of interest in this research complicated the development of the all-liquid release and fixation process shown in Figure 51.

Another straightforward improvement to this fixture would be the mesh plate itself. Not only

did the laser trimming process certainly introduce an unknown array of partially ablated and heat-damaged PMMA molecules in the vicinity of the numerous cuts involved in the mesh's fabrication, but it also imposed limitations on the critical dimensions of the mesh. A fine mesh woven from a strongly resistant polymer such as teflon (PTFE) would not only avoid the contamination concern associated with laser machining, but would also allow optimal the choice of optical mechanical dimensions for the mesh. Commercial meshes can be woven much finer than the minimum hole size achievable with laser cutting – 50 μm for the former *vs* 1 mm for the latter – and would be cheaper to produce.

3.3.3.2 Blue etch protection

For extended-length etches intended to produce very high-aspect ratio sacrificial voids, due either to thin layers or large lateral dimensions, it was found that the condition of the blue etch itself was a crucial parameter. Blue etch that had been improperly stored, contaminated, or heavily used was found to prevent full release of layers. In contrast, freshly-concocted blue etch not only achieved consistent success, but also exhibited, upon anecdotal observation, much faster etch rates. Keeping the blue etch in as pristine of a condition as possible was achieved mostly by standard means, including basic contamination prevention measures, frequent inspections, and liberal removal of visible debris.

However, one less-obvious aspect of preserving the potency of the blue etch was learned in this research. The sole liquid ingredient in blue etch, ammonium hydroxide, consists of ammonia (NH_3) dissolved in water. The ammonia molecules clearly play a crucial, if not fully understood, role in the etch process, so preventing their depletion is an essential aspect of maintaining the blue etch solution. As ammonia is a gas at standard conditions, it readily escapes from its dissolved state in the blue etch solution. This is confirmed by the noticeably large amount of concentrated ammonia gas present at the mouth of any beaker containing blue etch (or ammonium hydroxide). Therefore, the the unchecked escape of ammonia gas must be prevented. This was done by always keeping the blue etch in beakers with sealed lids, not only for storage between etches, but also during the etches themselves.

3.3.3.3 *No-Dry Release Process*

Previous sacrificial multilayer plating research had always involved drying the multilayer structures after release. However, it was found that the drying process presented considerable difficulty in the fabrication of laterally large structures with thin layers. First and foremost, air drying was found to frequently induce the downward collapse of the released layers. This is shown clearly in Figure 50, and was also visible at the sidewalls of the released structures. As this downward collapse

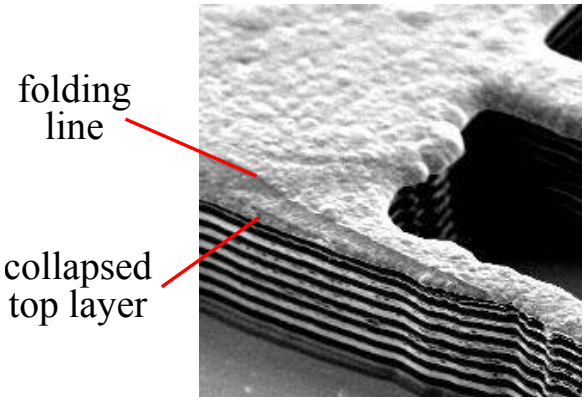

Figure 50: Collapse of top layer during drying. A portion of the top layer has collapsed downward, presumably due to excessive airflow.

effectively removed the sacrificial voids that were so carefully created between layers, it would be detrimental in most applications of these multilayers. In the investigated applications of capacitors and laminated magnetic cores, this was certainly true; in the former case, collapsed voids impede the deposition of dielectric and the infiltration of the liquid electrode; in the latter case, contact between layers creates a low-impedance path for eddy currents.

Consequently to the discovery of this problem, the airflow during drying was held to very low levels or eliminated completely. However, in these cases, stiction, discussed in subsection 1.3.1, became a major difficulty, ultimately producing the same catastrophic symptom, collapse of the interlayer gaps. A progression of solvents after the etch, starting with the basic deionized water rinse, continuing to isopropanol, and ending with methanol alleviated these stiction-related issues, but did not completely solve them.

For the laminated-core application, the essential requirement for the interlayer gaps was not the preservation of the exact gap geometry, but rather the imposition of electrical discontinuity between

layers. While an intact gap would achieve this, so would the deposition of an electrical insulator between layers. In fact, relying on continuous layers of an insulating material, instead of relying on unfilled sacrificial gaps, to maintain electrical separation between layers was judged to provide superior mechanical robustness of the finished structures. Deposition of an insulator in the interlayer gaps entailed the transport of precursor gases and/or liquids into the gaps. This presented a mass transport problem in general. But for the case of liquid-phase precursors, if the precursor solvent had low adhesion to the structural layers, the difficulty in re-wetting a dried structure would compound the mass transport difficulties.

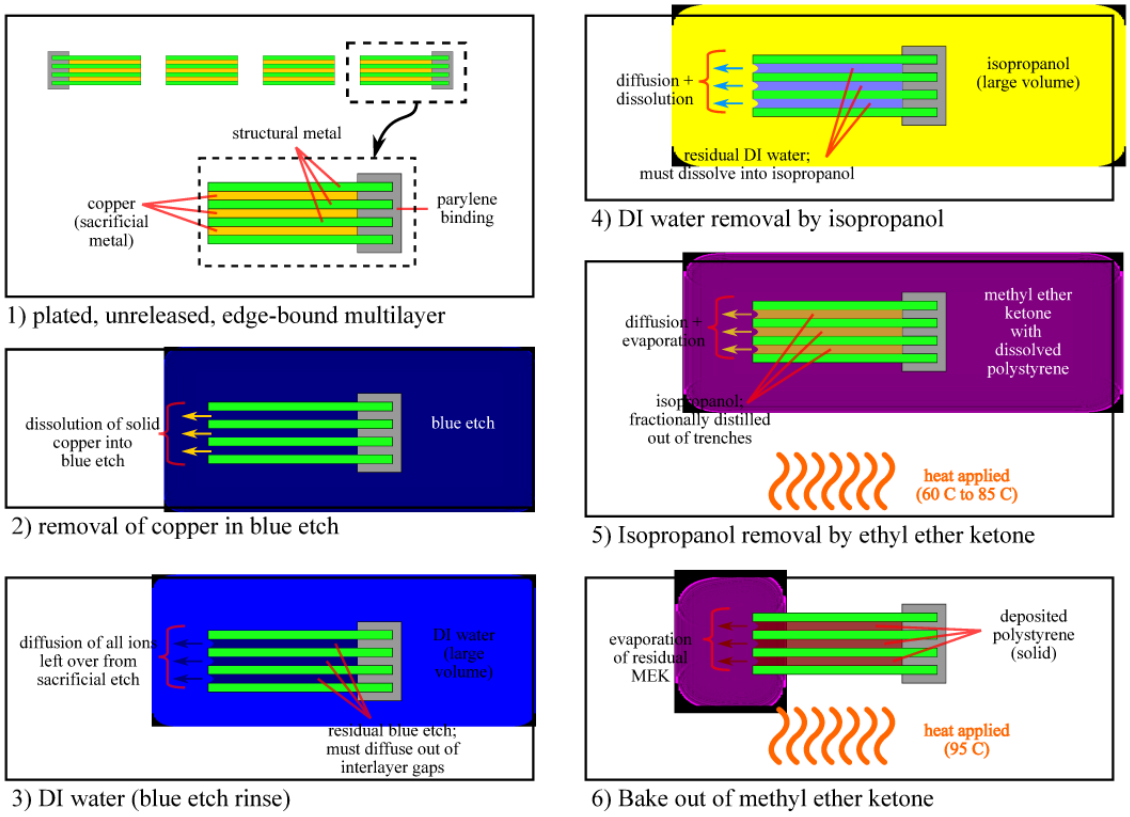

Figure 51: All-Liquid Sacrificial Post-Release Process. This sequence of steps not only avoids drying-related stiction risk, but also allows deposition of an insulator (polystyrene) between released layers.

Accordingly, great advantage was seen in the creation of a post-etch process that did not require drying of the sample. This would simultaneously bypass the difficulties of stiction and re-wetting after drying, while allowing the use of liquid-phase insulator deposition. While vapor-phase deposition and liquid-phase deposition have their advantages relative to one another, a liquid-phase

deposition process was found to be adequate for the purpose of electrical insulation. Polystyrene (PS), in the form of styrofoam, was dissolved in methyl ethly ketone (MEK), to a concentration that was roughly half of the cloud point. It was found that a wide variety of materials, including stainless steel and nickel, could be dip-coated with this PS/MEK solution, and that the coatings provided electrical insulation. The high vapor pressure of MEK at room temperature produced near-instantaneous formation of a highly-contiguous polystyrene layer.

A no-dry process that both accomplished rinsing after blue etch, as well as the deposition of polystyrene between layers, was successfully developed, and is shown in Figure 51. As the very first step of the process, the plated structures are removed from the substrate. As the substrates were oxidized silicon wafers, this substrate release was accomplished by an extended soak in hydrofluoric acid. Next, a brief (4-6 minute) sacrificial etch is performed and a liftoff-patterned parylene deposition is performed. The brief sacrificial etch creates topology on the structure's sidewalls that allows the deposited parylene to articulate with the structural layers, providing them with mechanical support. Parylene was selected not only for its robust deposition process, but also for its apparent imperviousness to long-term exposure to blue etch.

Next, the main etch is performed, with the sample resting on the mesh plate of the fixture shown in Figure 49 typically requiring 2 to 3 days to completely traverse lateral distances of 300 μm . After the main etch, the sample, still held by the fixture, is left in room-temperature, unagitated DI water for approximately 5 hours. Presumably, since blue etch is an aqueous solution, practically all of its components are able to diffuse out of the gaps and into the bulk of the DI in this time. Following the DI rinse, the sample is transferred from the fixture to a small piece of Texwipe for mechanical support.

It is then moved from a DI rinse into a beaker of isopropanol, and left for approximately 3 hours. Isopropanol and DI water are miscible, so, given sufficient time and the large volume of the beaker relative to the interlayer gaps, the composition of the liquid within the gaps should be practically 100% isopropanol. Immediately after this, the sample is transferred from the isopropanol beaker to a beaker containing a PS/MEK solution at approximately half of the cloud point concentration. It would be theoretically plausible to rely on the miscibility of MEK and isopropanol, in combination with their relative volumes, to achieve full infiltration of the MEK into the gaps.

However, it was found that the MEK evaporated at a rapid rate, making the the completion of this interdiffusion dubious. Instead, fractional distillation was used to ensure the isopropanol's removal from the interlayer gaps. This fractional distillation was carried out using a hotplate at approximately 80 C. This temperature, while it was clearly sufficient to remove the IPA, producing visible currents flowing outward from the structure, was also sufficient to evaporate most of the MEK within 2 hours. After the majority of the MEK has been evaporated, presumably signaling that all of the isopropanol has also been removed, the hotplate temperature was set to 95 C to complete full evaporation of the MEK.

3.4 Conclusions

3.4.1 Alignment with electrochemical theory and electroplating practice

Electrochemical investigations have been conducted on a vast array of materials, including metals, semiconductors, nanomaterials, and polymers. Certainly, the ability to use these materials in electroplated multilayer structures could enhance many aspects of their performance as well as facilitate their fabrication. However, the relatively specialized body of knowledge and equipment involved in classical electrochemistry would require considerable effort to assimilate, as the processing detailed in this chapter is focused on relatively large-scale electrochemical reactions.

Similarly, an enormous amount of practical electroplating knowledge has been compiled, especially in industrial applications. This includes plating equipment, knowledge about tank and electrode preparation, and bath maintenance techniques, such as Hull cell analysis. Unfortunately, again, this work is, relative to industrial electroplating, focused on extremely small-scale production, making it difficult to bring know-how from the former to bear upon the latter.

In summary, there are very few interfaces by which well-developed knowledge, either from theoretical electrochemistry or from electroplating practice, may be clearly incorporated into this work. However, the decorative copper plating bath described in this chapter, as used to create smooth structural layers, illustrates one such interface: off-the-shelf, pre-formulated plating baths. This is a somewhat coarse interface, in that the bath's properties could not be finely engineered, and that options were not present to troubleshoot occasional problems with the bath. However, its impact on this research was undeniably substantial.

Therefore, further efforts to establish more interfaces by knowledge from these adjacent fields can be brought to bear on this research could have potentially large payoffs. Bath analysis and conditioning tools and techniques are one obvious dimension of development along these lines. At the least, being able to detect and correct out-of-specification baths would avoid wasted fabrication runs. Another possibly large benefit could be derived from comprehensively surveying industrially significant plating and finishing processes. For an electrochemical process to achieve economic viability under high-volume, industrial conditions, the fundamental process must be extremely robust. Consequently, such processes should be readily reproducible and therefore easily imported into a research lab environment.

3.4.2 Difficulty of imaging

In micromachining, imaging plays a pivotal role in verification of the fabrication processes, as well as in troubleshooting of problems and correlating characterization results with structural features. However, in this work, as the layer thicknesses decreased below 1 μm , imaging of the structures became increasingly difficult. This factor, alone, greatly slowed the overall progress of the research.

The most basic type of imaging germane to this research was electron micrography of the outer sidewalls of the multilayer structures. This was the main source of information about the fabrication results, but for layers thinner than 700 nm, even this technique was unreliable. More informative techniques, such as cross-sectioning and energy dispersion spectroscopy (EDS) were unsuccessful with thin layers. Staple techniques of alloy characterization and magnetic characterization, such as transmission electron micrography (TEM) and SPLEEM, were simply not compatible with these multilayer structures. These would have required either cross-sectioning or singulation and manipulation of individual layers.

Conventional spectroscopy techniques, such as atomic absorption (AA) spectroscopy or X-ray fluorescence (XRF) spectroscopy, are very useful for characterizing electroplated layers. But, as they do not implement spatial resolution, they could only be used to analyze the average composition of the entire plated structure, and not to analyze its physical features. Through-depth spectroscopy techniques, such as secondary ion mass spectroscopy (SIMS) or ion milling-assisted X-ray photoelectron spectroscopy (XPS), could potentially be used to provide fine details of the layer stack.

However, analyzing the full depth of a typical multilayer stack (50-100 μm) with one of these techniques would take a prohibitively long amount of time under available conditions. Further, to provide comprehensive information about a given multilayer stack, through-depth spectroscopy would have to be performed at multiple points across the stack's lateral surface, compounding the time requirements.

Any optimization that could break through even one of these imaging obstacles could deliver profound impact on this research. Alternatively, the application of any other readily available through-depth imaging technologies, such as C-mode scanning acoustic microscopy (CSAM) or tomographic techniques, would be equally valuable. Unfortunately, the aforementioned advanced imaging techniques are, just like the research into the multilayer structures itself, extremely complex, expensive, and time-consuming.

3.4.3 Dependence on mass transport

Many problems found in this research, from nonuniform thickness of the deposited layers to long sacrificial release times to refilling of interlayer gaps, can be viewed as mass transport problems. The very strength of the sacrificial multilayer process – the extremely high aspect ratio of the interlayer gaps – entails that physical access to the layers' surfaces happen through long, narrow paths.

Being able to faithfully model these mass transport conditions could provide important guidance to research efforts. For fabrication steps which depend on slow mass transport dynamics, reliable estimates of the required timescale could eliminate iterations during process development. For instance, the long sacrificial etch times required might have been predictable from a combined modeling of the blue etch reaction along with the diffusion of the reactants into and out of the interlayer gaps. More importantly, construing any critical mass transport limitations in conventional terms would potentially allow application of mass transport solutions from other research disciplines, especially chemical engineering.

CHAPTER IV

MICROFABRICATED CAPACITORS

4.1 Background

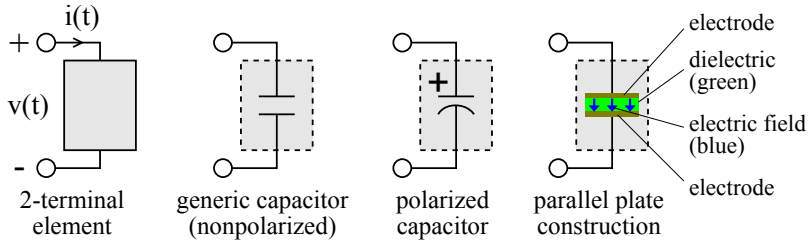

Figure 52: Each of the above can represent real capacitors: *(from left to right)* 2-terminal component; schematic symbol for a generic capacitor; schematic symbol for a polarized capacitor; cross-section of a parallel-plate capacitor

Capacitors, one of the four fundamental lumped, two-terminal passive component types, are essential in all of electronics. Figure 52 shows various canonical representations of capacitors, from a nondescript 2-terminal element through the frequently-employed parallel plate construction. In all cases, the behavior is described by the equation

$$i(t) = C\frac{d}{dt}v(t).$$

This relation entails elegant behavior that is very useful in linear signal processing, but it also describes an energy storage phenomenon. The latter is seen more clearly by considering the charge q contained in the capacitor,

$$i(t) = \frac{d}{dt}q(t),$$

yielding

$$q(t) = Cv(t).$$

Much like the energy storage behavior of inductors, capacitors' energy storage behavior is utilized extensively in switched-mode power supplies (SMPSs) to smooth the flow of electrical energy at various points in the circuit. This is certainly true at SMPS outputs, which are typically low-impedance nodes; here a sufficiently large capacitor can be added to provide near-constant voltage.

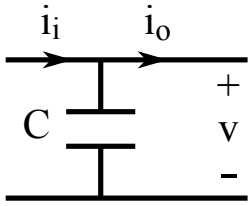

$$\begin{aligned} v &= \frac{1}{C}q \\ \Delta v &= \frac{1}{C}\Delta q \\ &= \frac{1}{C}\left[\int i\,dt\right] \\ &= \frac{1}{C}\left[\int (i_i - i_o)\;dt\right] \end{aligned}$$

Figure 53: Capacitor used for voltage filtering. As long as i_i and i_o have equal DC magnitudes, v maintains a constant DC value. Further, C can be chosen large enough to make Δv as small as desired

4.1.1 Basic Theory

Capacitors typically store electrical energy in an electrostatic field developed across a solid dielectric between their two electrodes. There are counterexamples to this, such as ruthenium-based electrochemical pseudocapacitors and air-core capacitors, but they are relatively rare in practical electronic systems [16]. The capacitor's variable of state can be either charge $q(t)$, voltage $v(t)$, or the stored energy $w(t)$:

$$w(t) = \frac{1}{2}Cv(t)^2 = \frac{v(t)\,q(t)}{2} = \frac{1}{2}\frac{q(t)^2}{C}$$

Maxwell's laws dictate that an electric field will develop between any two conductors with unequal voltages; the capacitance C, usually treated as a constant, is determined by the structure of this field. The invariance of C relies on the invariance of the capacitor's geometry as well as the invariance of the permittivity of any materials containing the capacitor's field.

From a fully-known electrostatic field configuration, the capacitive quantities v, q, w, and C can be extracted, given that both electrodes are conductive. w is a volume integral of local electrical energy storage:

$$w = \int_V \frac{1}{2}\vec{D}\cdot\vec{E}\,dV,$$

q is found from a surface integral over any surface S enveloping either electrode:

$$q = \int_S \vec{D}\cdot d\vec{A},$$

and v is found by a path integral

$$v = \int_M^N \vec{E} \cdot \vec{ds}$$

where M is any point on one electrode, N is any point on the other electrode. In the previous relations, D and E assume their typical values of, respectively, electric flux density and electric field intensity, in Maxwell's equations, with the typical linear constitutive relation

$$\vec{D} = \varepsilon \vec{E}$$

where ε is permittivity.

4.1.2 Key Behavioral Aspects

Having such conceptually simple behavior, practical capacitors are characterized in terms of typical nonidealities.

4.1.2.1 Nonconstant Capacitance

The assumption of constant capacitance relies on the electrical linearity (*i.e.* constant permittivity) of the field-storage medium. Many materials exhibit changes in dielectric behavior as temperature and/or voltage change. This is especially critical in signal processing applications, such as communication signal filtering, where circuit or system performance can depend sharply on capacitors maintaing precise, linear characteristics over a range of operating conditions.

However, in SMPS filtering applications, maintenance of some minimum capacitance is more important than maintenance of a constant capacitance. This is because larger capacitance will, in general, provide improved attenuation of voltage and current ripple, and can thus be readily tolerated. In contrast, SMPS filter capacitors are typically optimized both to minimize physical size as well as parasitics such as series resistance and inductance. Not surprisingly, this tradeoff between capacitance on the one side and physical size and parasitics on the other is present across all capacitor types. [38]

4.1.2.2 Voltage Breakdown

Arguably the most critical nonideality encountered in capacitors is dielectric breakdown. Any insulating medium, even a vacuum, will allow the passage of large currents in the presence of a strong

enough electric field. This passage of current generally comes with permanent damage to the dielectric material, and the field at which it occurs is known as the breakdown field strength. In a capacitor, where the dielectric geometry is fixed, this field strength is mapped to a corresponding voltage across the device's terminals.

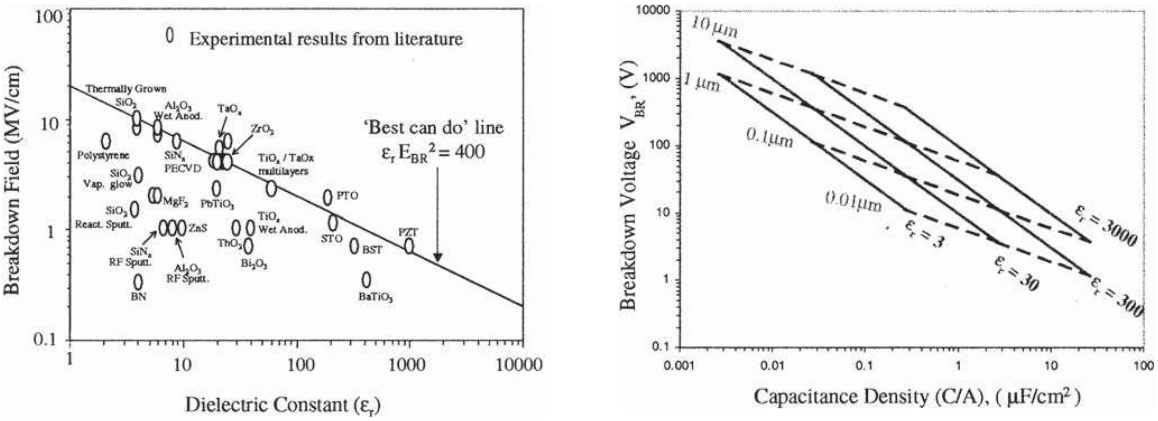

Figure 54: Breakdown characteristics of common dielectric materials. [30]

A capacitor's breakdown voltage depends strongly on the dielectric material's intrinsic properties as well as on the dielectric thickness. As well, certain structural defects in the dielectric layer, such as thinner-than-nominal areas, will result in lower breakdown voltage.

4.1.2.3 Dissipative Parasitics

Like any electrical component, capacitors are subject to dissipative parasitics. These can occur either in series or in parallel with the capacitance itself. Series resistance is commonly known as electrical series resistance (ESR), while parallel resistance is known as leakage resistance. Leakage resistance is determined primarily by the dielectric material's intrinsic properties along with the dielectric's thickness and any structural flaws. At higher frequencies, the dielectric material itself may experience dissipation (*i.e.* complex-valued permittivity), and this additional loss is also modeled as leakage resistance. Series resistance is a function of the electrode material's properties and geometry. At high frequencies, series resistance can also present elevated resistance, due to the skin effect-induced current crowding within the electrodes.

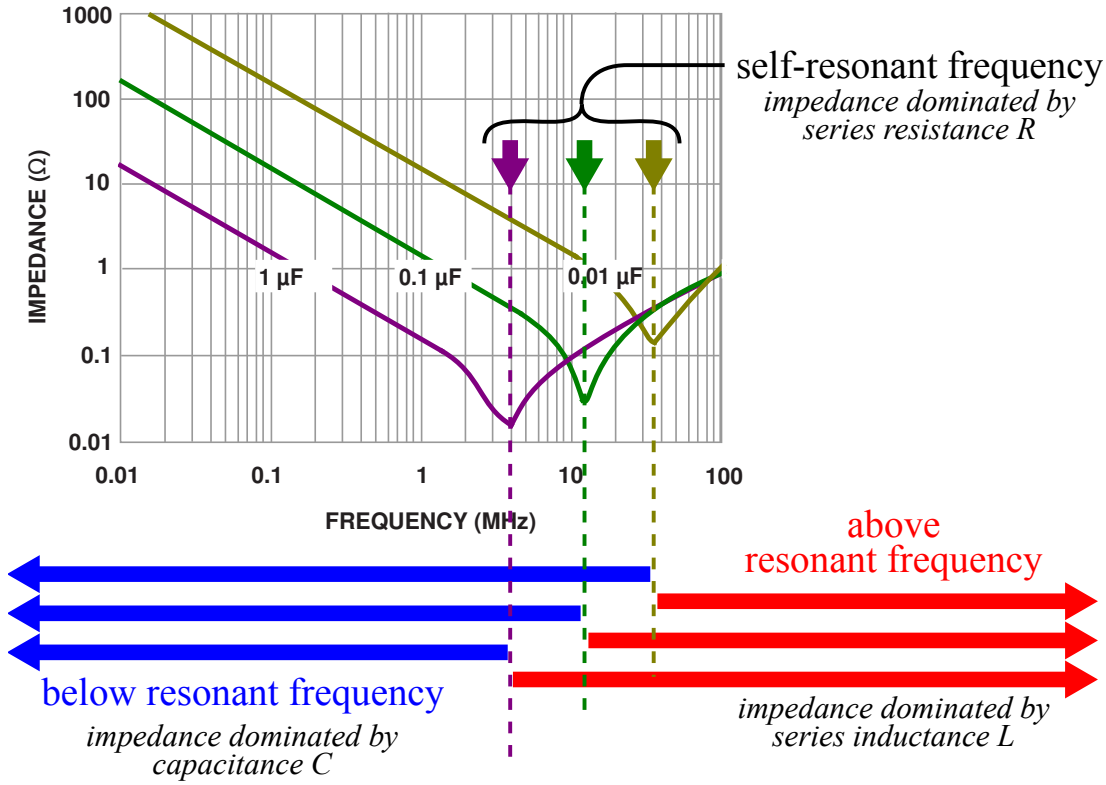

Figure 55: Typical Capacitor Impedance vs. Frequency [67]. Impedances of three capacitors, values 1 μF , 0.1 μF , and 0.001 μF , are shown. Each capacitor goes through self-resonance at its own frequency.

4.1.2.4 Reactive Parasitics

Capacitors also exhibit series inductance (ESL), due to the finite expanse of their electrodes, as well as any leads exterior to the device. The interaction between a capacitor's capacitance and the prevalent parasitics is frequently modeled as a resonant RLC circuit. Figure 55 shows a typical capacitor's impedance as a function of frequency.

At low frequencies, the device's capacitance presents the dominant impedance. However, as frequency increases, the impedance due to the capacitance decreases, and the impedance due to the ESL increases. At some frequency, called the device's self-resonant frequency, the impedances of the capacitance and the series inductance add to zero. At this frequency, the impedance of the device is dominated by the ESR. Beyond the self-resonant frequency, the ESL presents the dominant impedance.

4.2 Capacitor Implementations

4.2.1 Conventional Constructions

There are a number of commercially important capacitor technologies. Aluminum electrolytic and multilayered ceramic (MLCC) varieties are the most commonly used discrete capacitor types. [64]

4.2.1.1 Multilayer Ceramic

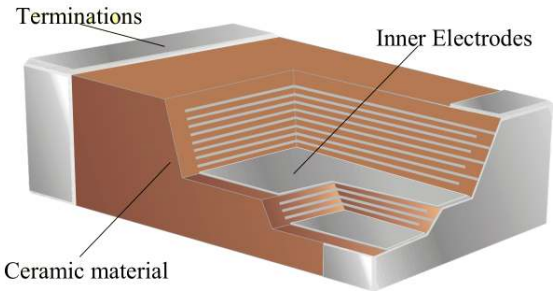

Figure 56: Multilayer Ceramic Capacitor Construction (courtesy Walsin Technology Corp.)

Ceramic capacitors are fabricated by stacking alternating layers of conducting material and dielectric material. The layers are cut and stacked prior to a single firing process, which simultaneously cures both electrode and dielectric ceramic materials. After the ceramic firing, opposing ends of the multilayer structure are metallized to create the device's two terminals. A typical MLCC construction is shown in Figure 56.

MLCC capacitors offer excellent performance and density for a number of reasons. First, many of the highest-performing dielectric materials are ceramics. Second, the multilayer construction allows the use of very thin layers, as thin as the pre-fired ceramic materal sheets may be created, providing for high capacitance density. Finally, ceramic materials are extremely durable, and highly compatible with even the most demanding packaging and PCB assembly flows.

The principal disadvantage of MLCCs is the high thermal budget required for their fabrication. Ceramic firing typically happens well above 1000 C, precluding the integration of MLCC fabrication with most electronics fabrication processes.

4.2.1.2 Electrochemical Double Layer

Electrochemical double layer capacitors (EDLCs) are a relatively new breed of capacitor that offer very high energy storage densities. First described in a 1957 patent issued to Howard Becker [28], EDLCs consist of a solid electrode immersed in a conductive electrolyte. [16].

There is no directly-constructed dielectric layer; rather, the electrolyte naturally develops a very thin nonconducting region in the vicinity of the electrode. This nonconducting region is known as the electrochemical double layer, and it is formed as charge carriers in the electrolyte diffuse into, and adhere to, the conductive electrode. As more charge carriers deposit onto the electrode, an electrical potential develops between the electrode and the bulk of the electrolyte until an equilibrium is reached. The nonconducting region is known as the electrochemical double layer (EDL), and its charge storage capability is the source of an EDLC's capacitance.

The EDL is not able to withstand voltages far above 2 V, as most electrolytes become subject to hydrolysis at these potentials. However, the EDL can provide high per-area capacitance, on the order of 15-50 μF /cm^2. At the same time, since the EDL is spontaneously formed at the liquid-solid interface, without the precondition of having any particular dielectric material deposited, it is extremely conformal to the electrode. In the presence of electrodes with large surface area, such as those made from pressed carbon powder, the EDL's high effective surface area outweighs its relatively low voltage limitation and allows EDLCs to achieve superlative capacitance and energy storage ($\frac{1}{2}CV^2$) attributes.[16]

4.2.1.3 Electrolytic

Electrolytic capacitors offer high density and low cost. Electrolytic capacitors are formed through the electrolytic growth a thin layer of oxide on a metal electrode. This electrolytic oxidation, or anodization, gives a dielectric layer with extremely precise thickness control as well as high quality and conformality. As well, the metal electrode can be pre-treated to provide extremely high surface area.

There are two principal varieties of electrolytic capacitors, each identified by metal that forms the anode: aluminum and tantalum. Aluminum electrolytic capacitors utilize aluminum foil as the anode. Prior to anodization, the foil's surface area is amplified by a brief wet etch that selectively

dissolves grain boundaries. The anodes of tantalum electrolytic capacitors are typically masses of very small tantalum particles that have been pressed into a rigid, porous structure [37].

After the oxide layer is grown, the second capacitor electrode, the cathode, is deposited over the anode. To achieve the desired conformality, a liquid-phase cathode, *i.e.* a conducting liquid, is typically used. Vacuum can be used to help remove any gas trapped in the anode's surface area-enhancing features (pores for tantalum anodes, trenches for aluminum anodes). Practical contact is made to the liquid cathode is accomplished by bringing a mass of inert metal into contact with it. An electrochemical double layer will form at the cathode's liquid-solid interface, nominally creating a parasitic capacitance within the cathode. However, this capacitance does not affect the overall device's behavior; not only is the parasitic capacitance much larger than that of the electrolytic oxide layer, but the electrochemical double layer is easily punctured and its capacitance thereby shunted around. However, it is still important to choose an appropriate material for the solid cathode electrode, to facilitate EDL puncture and minimize the resistance at the puncture points.

Even the most conductive liquids give far higher bulk resistivity than metals, leaving electrolytic capacitors with a correspondingly higher series resistance. This is universally undesirable and is mitigated by locating the solid portion of the cathode as close as possible to the anode, thus minimizing the mean path that charge carriers must travel. To minimize reliability issues due to evaporation of the liquid portion of the cathode, low-vapor-pressure organic solvents such as poly ethylene glycol (PEG) are preferred in practice.

For tantalum electrolytic capacitors, manganese dioxide (MnO_2), a solid semiconductor, can be used as a cathode material. MnO_2 is still significantly less conductive than a metal, but it is much more conductive than liquid electrodes. MnO_2 is formed by immersing the anodized anode in an aqueous manganese nitrate ($MnNO_3$) solution, and then removing the anode and heating it to 140 C. As the water is boiled off, a thin layer of solid $MnNO_3$ is left on the tantalum anode. The $MnNO_3$ is pyrolyzed to form MnO_2. [49] This procedure must generally be repeated several times to achieve high-quality MnO_2 layers.

4.2.1.4 Thin Film

Thin film capacitors are utilized heavily in IC design to provide on-chip capacitance. Many types of circuits, both digital and analog, require capacitors, In cases where the required capacitance is small enough or on-chip integration is valuable enough, the capacitor can be fabricated along with the MOSFETs.

There are two typical thin-film constructions: gate oxide and metal-insulator-metal. In gate oxide capacitors, the thin layer of silicon dioxide used to form MOSFET gates is used to create capacitors. In metal-insulator-metal (MIM) capacitors, a very thin layer of insulating material is deposited on top of a metal layer, and a second metal layer is deposited on top of the insulator, producing a metal-insulator-metal "sandwich." MIM capacitors offer higher density and a wider variety of dielectric materials than gate oxide capacitors, but at the cost of diminished stability over temperature and process variations as well as the extra processing steps. [34].

4.2.1.5 High-performance capacitors

Capacitors' performance is generally proscribed by the dielectric material. Variation over temperature and voltage, as well as leakage dissipation, are principally due to the dielectric. In applications where the best performance is needed, polymer and mica are frequently preferred dielectric materials. In cases where extremely low losses and minimal variations are demanded, air or vacuum can serve as the dielectric. To achieve very high breakdown voltages, such as seen in power grids, certain oils can be used as a dielectric material. The major disadvantage of all of these capacitor types are their relatively high cost and low capacitance density. Particularly for oil and vacuum capacitors, packaging costs can be very high.

Certain polymers, such as polypropylene, exhibit an important behavior known as self-healing. In a generic capacitor, a breach in the dielectric layer will obviously lead to increased leakage current, as well as a potentially catastrophic amount of heat dissipated in a very small volume. In a self-healing capacitor, such a breach will, due to the high-temperature interaction of the electrode and dielectric materials, devolve into a nonconducting region. In this sense, any points of failure in the capacitor's dielectric are "healed."

4.2.2 Microfabricated Capacitors

A large number of approaches exist to the microfabrication of capacitors. These are all based on combination of the formation of high surface area structures and the deposition of a suitable dielectric material, and a subsequent second electrode, on top of the high surfare area structures. Several such constructions are reviewed in subsection 3.1.1, with particular attention paid to the formation of the initial high surface area structure; examples include [6]. and [31].

The utility of microfabricated capacitors is augmented by the wide array of dielectric materials that can be deposited onto high surface area structures. Whereas electrolytic constructions only allow for one type of dielectric material (*e.g.* aluminum oxide or tantalum oxide), microfabricated capacitors may employ any material, either to tailor its performance, or to extend its performance as new materials are developed. A number of extremely high-*k* dielectric materials, such as barium strontium titanate (BST), are available for this role [53].

4.3 Microfabricated Capacitor

4.3.1 Overview

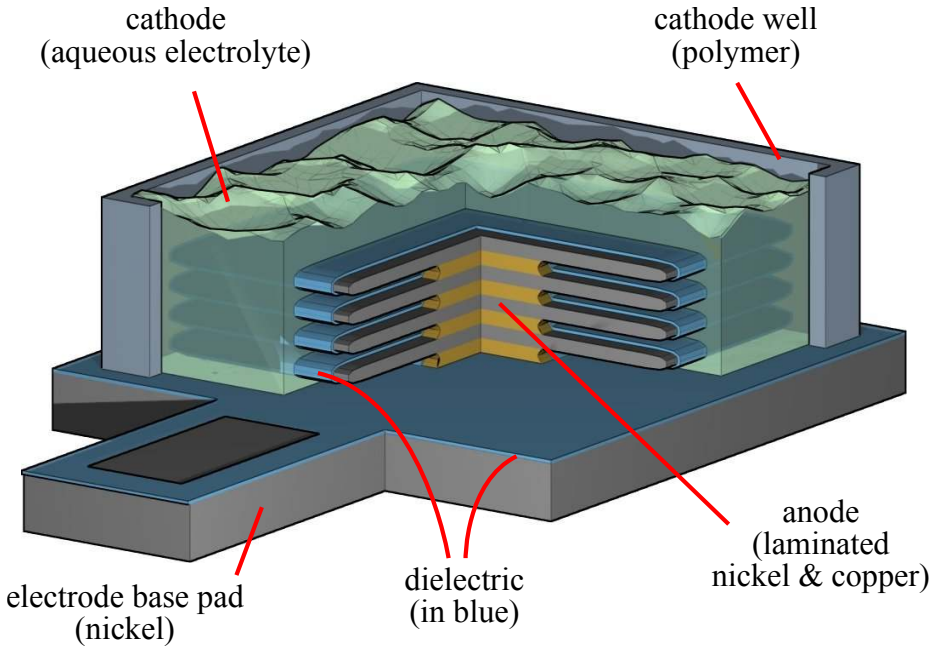

Figure 57: Microfabricated capacitor schematic

The sacrificial multilayer plating process described in Chapter 3 offers a number of advantages for capacitor implementation. Most importantly, it gives extremely high specific surface area, promising very high capacitance density. But it also gives access to a very wide device design space, allowing basically free choice of electrode thickness and dielectric thickness, as constrained only by available volume. At the same time, this is all achieved with a very low thermal budget – under 200 C – and full CMOS backend compatibility.

The capacitor construction demonstrated here is somewhat similar to that of electrolytic capacitors: a single, contiguous metal structure with high surface area is used as an anode, and a conducting liquid is used as a highly conformal cathode. However, in this new construction, the capacitor's dielectric layer is deposited onto the anode, rather than electrolytically formed at its surface. Figure 57 shows a schematic of the microfabricated capacitor.

4.3.2 Design

4.3.2.1 Lateral Geometry

One of the dominant difficulties anticipated with this capacitor's development was the fragility of the anode's layers. This was addressed by designing the anode as an array of doubly-supported cantilevers, or air bridges. The coalescence of the air bridge arrays and their supports into a monolithic structure is shown in Figure 58. Implicit in this combination is the ability to recover both the air

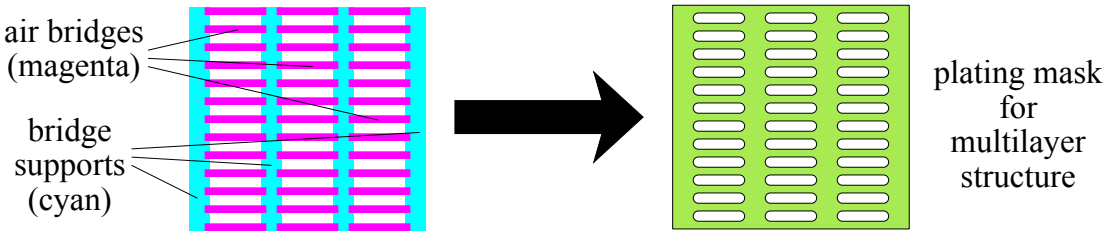

Figure 58: Plating Mask Design for Air Bridges and Supports. Two distinct features – air bridges and supports – are created by a single plating mask.

bridges and their supports from the monolithic structure. This is accomplished via a timed etch of the sacrificial layers, and is illustrated in Figure 59.

A variety of air bridge arrays were designed. In all cases, the air bridges were 15 μm wide, and the columns supporting them were 100 μm wide. All anode structures had outer dimensions

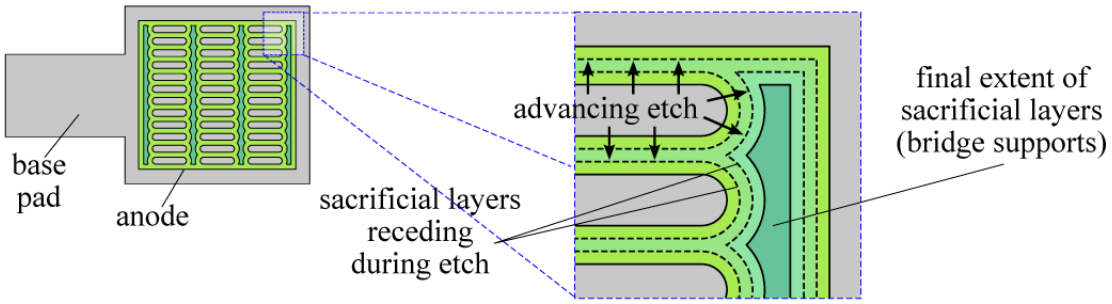

Figure 59: Formation of Air Bridges and Supports During Etch. The timed wet etch is assumed to act isotropically on the plated structure.

of approximately 1.5 mm x 1.5mm. The length of the air bridges, as well as the distance between them, was varied. Air bridge lengths of 130 μm , 360 μm , and 590 μm were used, and air bridge separations of 15 μm , 45 μm , and 75 μm (corresponding to pitches of 30 μm , 60 μm , and 90 μm , respectively) were used. The array configurations that were characterized are summarized, along with their relevant geometrical data, in table Table 2

4.3.2.2 Dielectric Layer

In all cases, atomic layer deposition (ALD) was used to add the dielectric layer onto the released anodes. The most critical aspect of the dielectric layer design was to ensure a breakdown voltage of at least 5V, to allow significant margin for characterization. ALD-deposited alumina has a breakdown strength of 2 to 4 MV/cm, translating to a minimum dielectric thickness of 25 nm.

4.3.2.3 Expected Capacitance

The microfabricated capacitors were analyzed with the help of the parallel-plate approximation, which gives the total capacitance C as

$$C = \frac{\varepsilon_0 \varepsilon_R A}{g}$$

where A is the total area, g is the dielectric thickness, ε_0 is the permittivity of free space, and ε_R is the relative permittivity of the dielectric. g was estimated from the ALD process parameters, and ε_R was taken to be 6, the typical value for ALD-deposited alumina. Analytical determination of the anode surface area A requires modeling of the progress of the timed etch. To simplify this calculation, a constant-rate, isotropic etch was assumed. This allows use of the results presented in Appendix A.

The following table summarizes the calculated areas and perimeters of the fabricated air bridge array structures, both before and after the 10-minute sacrificial etch.

Table 2: Air Bridge Array Geometries

Design	Air bridge length	Inter-bridge gap	Array size (rows × columns)	Etch depth $(-1 \times s)$	Area $A_G(s)$	Perimeter $P_G(s)$
X	360 μm	75 μm	16×3	0 μm	0.857 mm^2	47.63 mm
				10 μm	0.459 mm^2	12.4 mm
Y	130 μm	45 μm	24×6	0 μm	1.311 mm^2	56.27 mm
				10 μm	0.804 mm^2	21.8 mm

In an etched multilayer anode, the total surface area A is given by:

$$A = n_K h_K P_K + n_C h_C P_C + A_K + 2n_C (A_K - A_C)$$

where n_K, h_K, A_K, and P_K are the number, height, area, and perimeter, respectively, of the structural layers, and n_C, h_C, A_C, and P_C are the number, height, area, and perimeter, respectively, of the sacrificial layers. Since the structural layers are unaffected by the sacrificial etch, their area and perimeter can be taken as that of the unetched geometry (etch depth = 0 μm). This allows the following substitutions into the total area equation above:

$$\begin{aligned} A_K &= A_G(s)|_{s=0} \\ P_K &= P_G(s)|_{s=0} \\ A_C &= A_G(s)|_{s=-10\mu m} \\ P_C &= P_G(s)|_{s=-10\mu m} \end{aligned}$$

where $A_G(s)$ and $P_G(s)$ are the area and perimeter of the plated multilayer structure. Further, in all multilayer structures produced, the structural and sacrificial layers were of equal height h:

$$h_K = h_C = h$$

Combining these, the total anode area can be expressed in terms of the quantities given in Table 2:

$$A = h\left[n_K P_G(0) + n_C P_G(-10\mu m)\right] + A_G(0)(2n_C + 1) - A_G(-10\mu m)(2n_c)$$

giving the expected capacitance as

$$C = \frac{\varepsilon_0 \varepsilon_R A}{g} \tag{23}$$

$$= \frac{\varepsilon_0 \varepsilon_R}{g} \left(\left[n_K P_G(0) + n_C P_G(-10\mu m) \right] + A_G(0)(2n_C + 1) - A_G(-10\mu m)(2n_c) \right) \tag{24}$$

Evaluating this equation for the characterized capacitors, which had 10 structural layers and 9 sacrificial layers, all of height 4 μm yields, and a 100 nm alumina dielectric layer gives the expected areas and capacitances shown in Table 3. The dielectric layer was a composite of 100 nm alumina ($\varepsilon_R = 7$) and a parylene coat of 250 nm ($\varepsilon_R = 3.5$) on the top surface of the anode and outer edges of the layers.

Table 3: Anode Areas and Expected Capacitances

Geometry	Anode area	Expected capacitance
X	10.44 mm^2	3.7 nF
Y	13.85 mm^2	4.6 nF

4.3.3 Fabrication

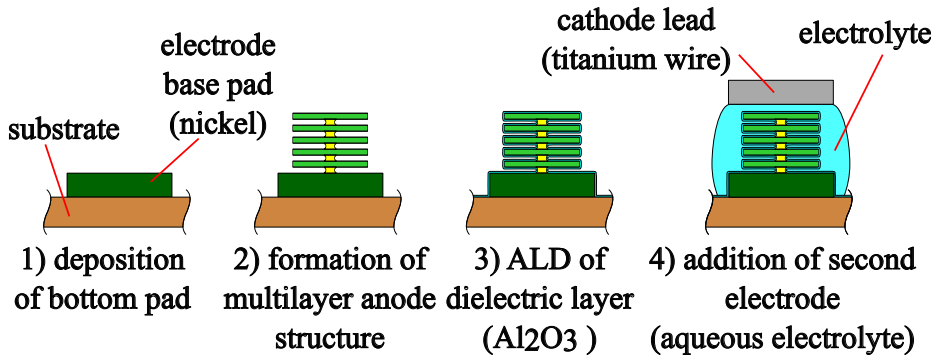

Figure 60: Capacitor Fabrication Flow. A multilayer nickel structure is formed on top of a nickel basepad, and then partially released. Onto this multilayer structure, a layer of alumina is deposited using atomic layer deposition (ALD). Finally, a liquid-phase eletrode is added and vacuum-infiltrated into the multilayer structure.

The fabrication flow for the microfabricated capacitors is shown in Figure 60. The sacrificial multilayer plating process was used to form a multilayer structure, with nickel as the structural material and copper as the sacrificial material. A timed etch was used to partially release the multilayer structure. Trimethylammonium (TMA) and water were used as precursors for the ALD deposition of alumina, the capacitor's dielectric layer. An aqueous solution of nickel sulfate, boric acid,

and saccharin was used as the liquid electrode. Weak vacuum levels, no lower than 1 Torr, were used to infiltrate the liquid electrode into the multilayer structure. This process flow was relatively straightforward, but a few challenges still required solutions.

4.3.3.1 CTE Mismatch

The first of these was delamination of the plated nickel structures from the silicon wafer. This was presumably due to the difference in coefficient of thermal expansion between silicon and nickel, as exacerbated by the 200 to 250 C ALD process. A minor case of delamination is shown in

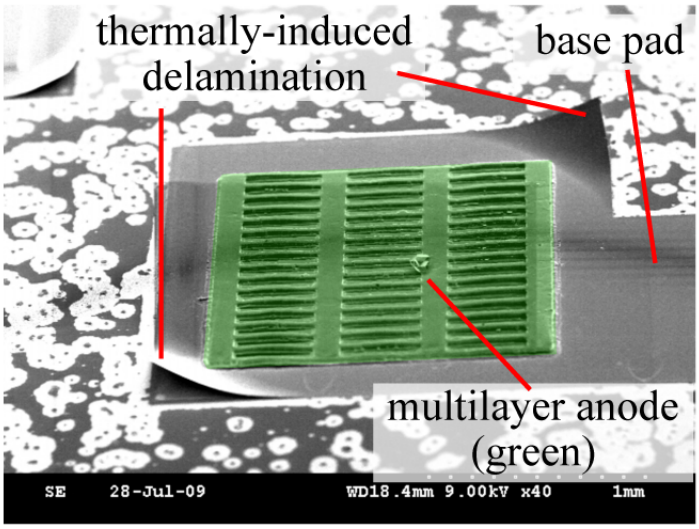

Figure 61: Delamination of Capacitor Base Electrode(false-color SEM micrograph).The thin base electrode showed a strong tendency to delaminate during ALD deposition. Aggravating factors included the sharp corners of the electrode, and the poorly controlled thermal ramp into and out of the ALD equipment.

Figure 61, where only the corners of the base pad have lifted. In more severe cases, the entire plated structure and base pad have been detached from the wafer, often taking a portion of the wafer's bulk (crystalline silicon) with them.

However, even minor failures, like the one pictured, are enough to render a capacitor useless, due to the anode metal being allowed to directly contact the cathode. Clearly, in Figure 61, the cathode liquid will be able to penetrate underneath the lifted corners of the basepad. Even if this lifting occurred prior to the ALD deposition, and there was an alumina coating on the exposed underside of the basepad, this alumina coating would be extremely vulnerable to cracking as the relatively flimsy basepad nickel layer was incidentally exercised during subsequent fabrication steps, such as addition

of the cathode well or cathode lead attachment. The recommended temperature for the ALD process was 250 C, but it was found that this produced very high incidence rates of delamination, affecting over 50% of the samples on a given wafer.

A twofold solution was adopted to mitigate this problem. First, to minimize steady-state CTE stress, the ALD process temperature was lowered to the minimum allowed value of 200 C. As well, to minimize the effects of thermal transients, the ALD chamber was cooled to 100 C during loading and unloading of the wafers. A 10 C/min thermal ramp was used during the heating of the ALD chamber to 200 C, and during the cooldown from 200 C back to 100 C. These solutions were found to be sufficient to lower the delamination yield losses to lower than 10%. Further, the parylene overcoat, discussed in the next section, also would have helped minimize the effects of minor delaminations.

4.3.3.2 ALD Nonidealities

Another fabrication issue that demanded considerable effort to manage was nonideal behavior of the ALD process. In principle, ALD is capable of producing extremely conformal layers of tightly controlled thickness. However, in practice, the ALD machine used in this work produced noticeably uneven coatings, as shown in Figure 62 and Figure 63.

Figure 62 shows the process's tendency to deposit rapidly at the outer edge of the layers, probably due as much to simple mass transport issues as to the beading of water vapor at the sharp layer edges. Also shown is the tendency of this relatively rapid deposition rate to bridge the gaps between layers. The presence of thicker dielectric or even gap bridgements was not a critical issue for the capacitor's performance. However, this almost certainly served to inhibit mass transport into the trenches, both during subsequent ALD cycles as well as during liquid cathode infiltration.

The incomplete ALD layer coverage shown in Figure 63 probably presents a more serious problem. Although this image shows an air bridge in which the sacrificial copper was not fully removed – a pathological case which was putatively not present in the characterized devices – the same interface conditions would exist within the air bridge supports. The interface conditions would consist of an uneven boundary of residual copper, which would potentially contain trapped etch residue as well, along which ALD deposition was apparently inhibited. This structural defect would probably

result in increased leakage current, if the liquid cathode could infiltrate to the regions of nickel that were not covered by the dielectric.

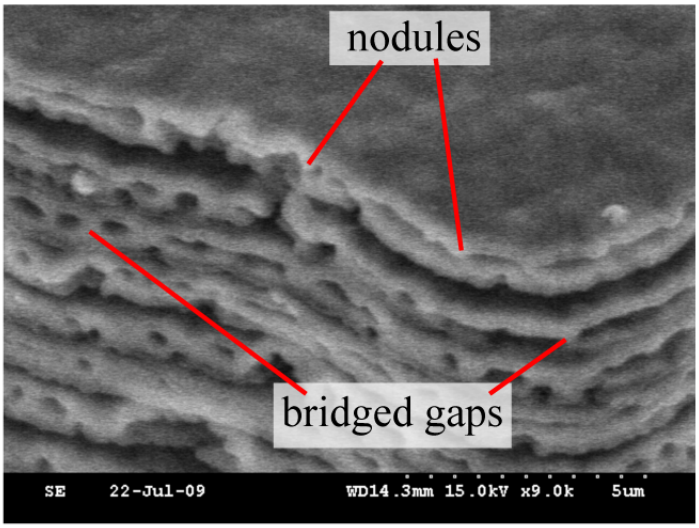

Figure 62: Excessive Dielectric Deposition on Capacitor Electrode. The bulk of the dielectric film is not visible this micrograph, but nodules formed at the edge of the nickel layers are visible.

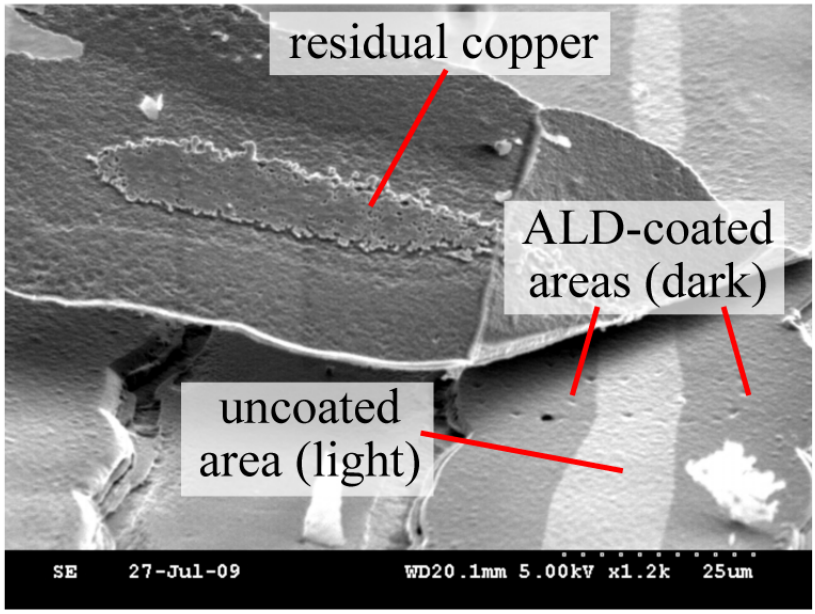

Figure 63: Exposed Interior of Capacitor Anode Gap. This micrograph indicates several features within the capacitor anode gaps. Residual copper can be seen, as well as contrast due to partial ALD dielectric coverage.

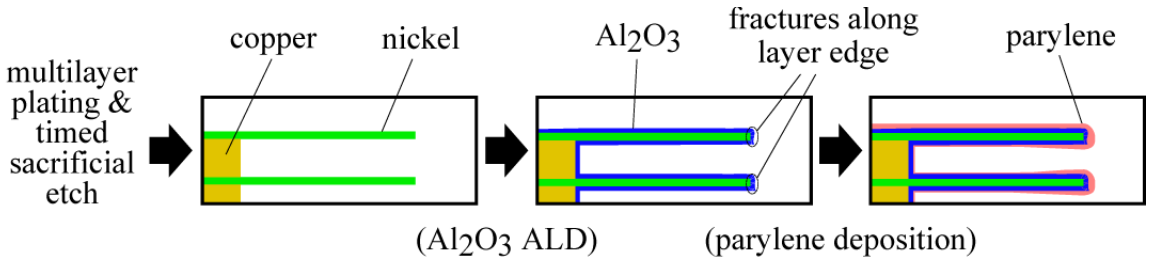

Figure 64: Parylene for Protecting Layer Edges. The role of parylene in repairing pinholes at the layers' edges is shown. The parylene deposition does not need to penetrate the gaps to serve this role.

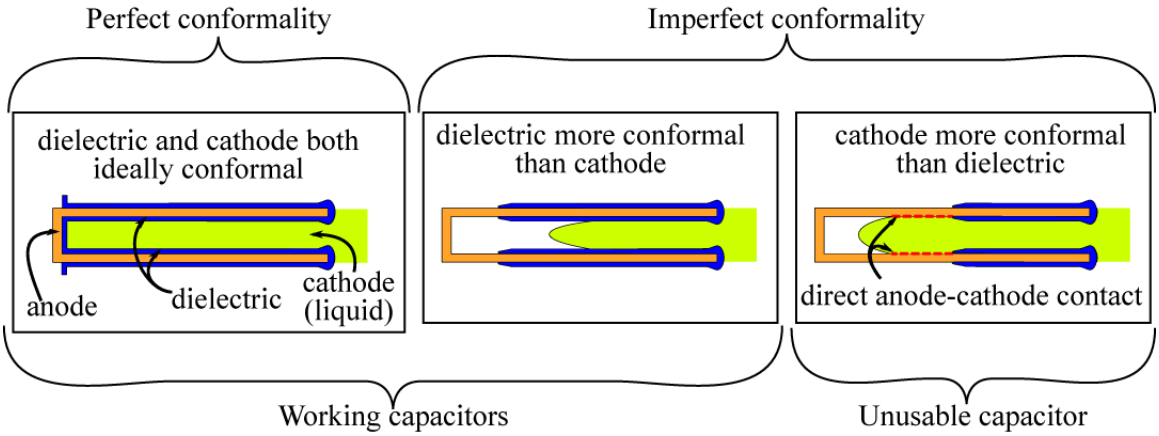

Figure 65: Potential Scenarios for Liquid Electrode Infiltration. Three qualitative regimes are possible. The liquid electrode need not penetrate the full depth of the gaps to create a functional capacitor. However, if the dielectric deposition is not more conformal than the electrode, a large short circuit path will be created between the capacitors' electrodes.

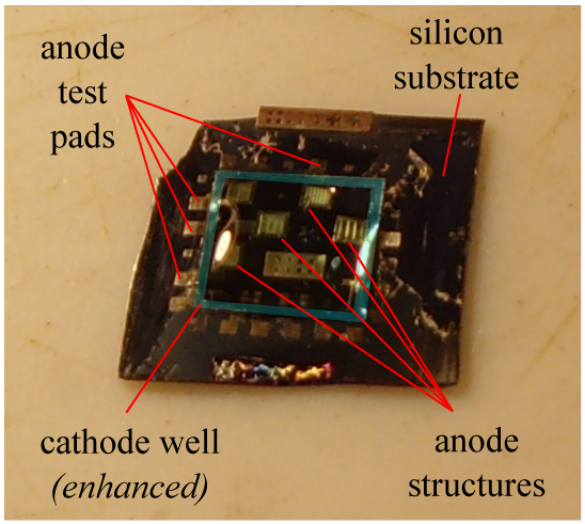

Figure 66: Photograph of Microfabricated Capacitor. There are several capacitors attached to this silicon die. A single cathode is shared by all devices shown. The surface gloss of the aqueous cathode can be seen.

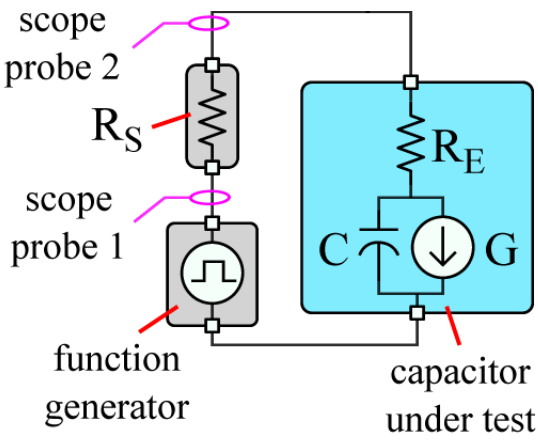

Figure 67: Capacitor Characterization Schematic

4.3.4 Characterization

4.3.4.1 Theory

The model assumed for the microfabricated capacitor, as well as the test circuit used, is shown in Figure 67. A leakage current source (G) was assumed, instead of a leakage resistance, because it was presumed that most leakage current was due to electrochemical reactions between the anode and the electrolyte, as enabled by pinholes in the dielectric layer. This assumption was supported by the observation during testing of small, slowly-developing bubbles at the edges of the anode – where the highest pinhole density could be expected. As the excitation voltage increases, these electrochemical reactions will certainly enter a mass transport-limited regime, and behave as voltage-independent leakage sources.

There are three quantities to be determined: capacitance C, series resistance R_E and leakage current G. G was measured by observing the DC voltage across R_S and dividing the voltage by R_S. C and R_E were calculated from time constant measurements, as extracted from the captured waveforms. Two time constant measurements, taken at different values for R_S, were sufficient to determine both C and R_E. The governing relations are

$$\begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix} = \begin{bmatrix} (R_{S1} + R_E)\,C \\ (R_{S2} + R_E)\,C \end{bmatrix},$$

which can be solved as:

$$\frac{\tau_2}{\tau_1} = \frac{R_{S2} + R_E}{R_{S1} + R_E} \tag{25}$$

$$(R_{S1} + R_E)\,\tau_2 = \tau_1\,(R_{S2} + R_E) \tag{26}$$

$$R_E = \frac{R_{S1}\tau_2 - R_{S2}\tau_1}{\tau_1 - \tau_2} \tag{27}$$

and

$$C = \frac{\tau_1}{R_{S1} + R_E} \tag{28}$$

$$= \frac{\tau_1}{R_{S1} + \frac{R_{S1}\tau_2 - R_{S2}\tau_1}{\tau_1 - \tau_2}} \tag{29}$$

$$= \frac{\tau_1 - \tau_2}{R_{S1} - R_{S2}} \tag{30}$$

From the captured waveforms, the RC time constant τ could be directly extracted, as shown in Figure 69.

4.3.4.2 Waveform Analysis

Figure 68 shows typical waveforms obtained during characterization. Capacitive charging and discharging is clearly illustrated. The exponential decay region of the waveforms, or rather substantial subsets thereof, were easily identified.

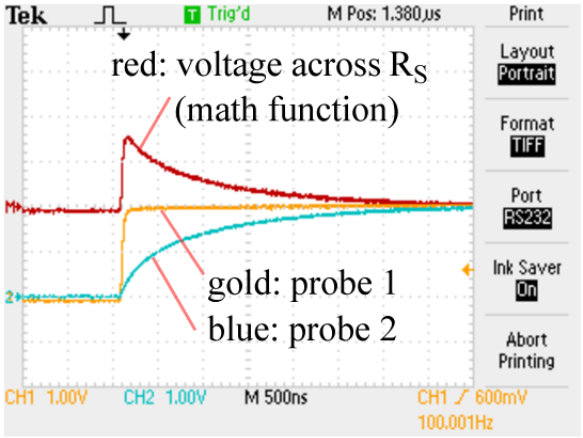

Figure 68: Microfabricated Capacitor Test Waveforms. The oscilloscope's built-in math function is used to indirectly measure the capacitor current.

Figure 69 illustrates the use of linear regression to extract the time constant from a waveform. The voltage across the sense resistor R_S is used to observe the time constant. The current through

an RC circuit $i(t)$ during charging and discharging of the capacitor is given by

$$i(t) = I_0 e^{\frac{-t}{\tau}} \tag{31}$$

where I_0 is a constant value, and τ is the RC time constant. The observed voltage across this resistor, $v_R(t)$, obtained using the oscilloscope's waveform arithmetic functionality, is a time series of values:

$$v_R(t_n) = R_S\, i(t_n)$$

where n is an integer that indexes the waveform's constituent datapoints, and t_n are the time values at the corresponding datapoints.

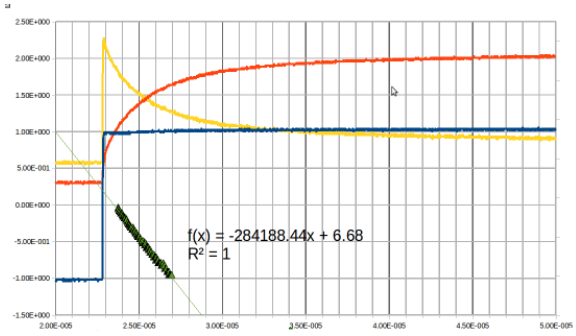

Figure 69: Capacitor Characterization Waveform Analysis. The exponential decay regime of the waveforms was manually identified, and analyzed with a semilog-scale linear regression.

A linear fit can be used to extract τ from a set of $v_R(t_n)$ datapoints. First, the analytical equation for $i(t)$ is manipulated into a form compatible with linear regression:

$$i(t) = I_0 e^{\frac{-t}{\tau}} \tag{32}$$

$$\ln[i(t)] = \ln\left[I_0 e^{\frac{-t}{\tau}}\right] \tag{33}$$

$$\ln i(t) = \frac{-t}{\tau} + \ln I_O \tag{34}$$

$$\ln i(t) = \frac{-t}{\tau} + \ln I_O \tag{35}$$

Next, the equation is converted into discrete time (t_n substituted for t) and current is rewritten in terms of the actual values for $v_R(t_n)$ to obtain a standard-form linear regression equation:

$$\ln\left[\frac{v_R(t_n)}{R_S}\right] = \frac{-t_n}{\tau} + \ln I_O \tag{36}$$

$$\ln v_R(t_n) = \frac{-t_n}{\tau} + \ln I_O R_S \tag{37}$$

$$y_n = m x_n + b \tag{38}$$

where

$$y_n = \ln v_R(t_n) \tag{39}$$

$$x_n = t_n \tag{40}$$

$$m = \frac{-1}{\tau} \tag{41}$$

$$b = \ln I_0 R_S \tag{42}$$

Successful application of linear regression using this formulation will permit, finally, an expression for τ:

$$\tau = -\frac{1}{m} \tag{43}$$

4.3.4.3 Results

Using the above analysis, the data from the characterized samples indicated the device characteristics shown in Table 4.

Table 4: Capacitor Characterization Results

Device Geometry	R_{S1} [Ω]	R_{S2} [Ω]	τ_{S1} [ns]	τ_{S2} [ns]	Capacitance C [nF]	Series Resistance R_E [Ω]	Parallel Leakage G [μA]
X (prior to vacuum infiltration of cathode)	380	760	575	1046	1.24	84.3	2.11
X	380	760	2119	3519	3.68	195	3.42
Y	380	760	575	1046	3.33	435	2.89

4.3.5 Discussion

The results in table Table 4 agree reasonably well with the expected values for capacitance (3.7 nF for Device X, 4.6 nF for Device Y). One of the dominant sources of uncertainty in estimating the capacitance is the conformality of the 250 nm parylene overcoat layer. The analytical model assumed that the parylene only deposited on the top surface of the anode and along the outer edges of the layers. In reality, this assumption underestimates the conformality of parylene's conformality, thereby underestimating the thickness of the actual dielectric layer. This poorly-controlled and

poorly-modeled conformality could explain the difference seen between Device Y's expected capacitance and its actual capacitance. In Device Y, the considerably shorter air bridges (130 μm vs Device X's 360 μm) likely exhibited less sagging and stiction-related deformation, thus facilitating deeper penetration of gas-phase parylene. Adjusting the expected penetration depth s of the parylene from the implicit value of 0 to a modest value of -5 μm lowers Device Y's expected capacitance to 3.5 nF, much closer to the observed value.

Also worth noting is the large increase in all three device attributes (C, R_S, and G) during vacuum infiltration in Device X. Certainly, capacitance and leakage current can be expected to increase as a greater area of contact is achieved between the liquid cathode and the dielectric layer. However, the series resistance R_S should also exhibit some increase as the electrolyte enters further into the anode's trenches, since current carriers have to travel significantly further to access the capacitance found at the far reaches of the trenches. Thus, the increase in these quantities is in perfect agreement with the qualitatively expected behavior that would come with a successful vacuum infiltration step.

These results demonstrate the capability of the sacrificial multilayer plating process to produce very high surface area structures. Clearly, these are not high-performance capacitors, but they do establish a plausible path towards a family of devices with excellent capacitance density (as high as 45 nF/mm^2 for 100-layer structures) and a high degree of flexibility, both in design and fabrication.

4.4 Conclusions

4.4.1 Utility of Sacrificial Multilayer Plating

The results presented in this chapter show, most importantly, what they were intended to: that the sacrificial multilayer process described in Chapter 3 is a powerful way to achieve high surface area throughout relatively large microfabricated structures. The high capacitance densities achieved, as well as the dramatic increase in density seen with vacuum infiltration, prove that the anode structures produced by sacrificial multilayer plating have high surface areas and high aspect ratio features.

Further, it can be seen that the nonidealities of the dielectric layer deposition and of the liquid cathode infiltration process present dominant limitations to the achievable capacitor performance. The requirement to include a parylene deposition step to overcome dielectric layer issues at the layer edges led to a considerable decrease in observed capacitance.

4.4.2 Inestimable Need for Robust Packaging

A secondary theme in this work was the primacy of packaging. The overall design of the device was, of course, thought out from the beginning of the design process. But, a more robust approach to the device assembly process – everything that happened after the sacrificial release etch of the anode structure – would have allowed much easier experimentation, and a considerable increase in the scope of the results.

This capacitor was not intended to be subjected to vigorous, long-term use. Creating the cathode wells, handling the devices, maintaining the liquid electrode, and attachment of the test leads were all executed with the proscribed intention of device-level characterization. Still, these steps proved to be a source of an unexpectedly large set of technical challenges.

4.4.3 Indispensability of ALD

Another dynamic well illustrated by this work is the great utility of atomic layer deposition. Although it is a relatively slow deposition process that is subject to its own set of nonidealities, it was a key enabler of this work. Its well-documented ability to deposit a highly conformal layer of tightly-controlled thickness served as a solid foundation of known behavior upon which the features of the microfabricated anode could be evaluated. Alternative deposition approaches, such as CVD or PECVD, would have introduced a prohibitive amount of development and characterization work, if they would have been capable of adequate conformality at all.

As well, ALD deposition of a wide range of dielectric materials, including barium strontium titanate (BST), is widely reported. With improved ALD deposition as well as access to the full set of ALD-compatible materials, it is estimated that the capacitance density of these devices could be increased by at least two orders of magnitude.

4.4.4 Future Improvements

There is much room to improve the capacitors, beyond the dimensions of packaging and ALD dielectric materials already discussed. One of the key barriers to the practical use of this capacitor is its high ESR. Use of a solid-phase cathode would be very helpful in reducing the ESR. Numerous conductive polymers are available for this role, as well as manganese dioxide (MnO_2), which is

widely used as a second electrode in solid tantalum capacitors [37].

Another major improvement in this device could come from the incorporation of an electrolytically formed dielectric. Electrolytic dielectrics offer the dual advantage of good control over thickness, as well as low pinhole density, due to the self-passivating nature of electrolytic oxide formation.. ALD deposition of aluminum has been reported, and this could serve as an ideal precursor to a higher-quality alumina dielectric layer.

CHAPTER V

HIGHLY LAMINATED-CORE INDUCTORS

5.1 Background

This chapter presents the application of the extensions to the sacrificial multilayer plating process, presented in Chapter 3, to produce inductors with magnetic alloy cores that are capable of supporting high switching frequencies. Electrodeposited magnetic alloys, such as permalloy, offer many benefits as magnetic core materials. They provide higher saturation flux densities and lower coercivities than ferrite materials [66]. Moreover, whereas ferrites generally require high-temperature firing, magnetic alloys are deposited at low temperatures, permitting their integration onto silicon wafers.

However, magnetic alloys, due to their good electrical conductivity, are prone to the presence of excessive eddy currents. At high switching frequencies, these eddy currents introduce high losses in magnetic materials. As well, eddy currents in a magnetic core will prevent utilization of the core's full volume of magnetic material, as described by the concept of skin effect.

A viable method for defusing the development of eddy currents in a magnetic core would alleviate these often prohibitive effects. Lamination of a core structure is a simple and well-proven technique for suppressing eddy currents. The sacrificial multilayer process can be used to create magnetic cores with extremely high degrees of lamination. This goal is pursued with the express purpose of reducing eddy current losses. The design and implementation of such highly laminated cores is presented.

Additionally, thorough attention is paid to measuring the core loss of these inductors, as loss measurements are key to not only establishing their utility in high-efficiency power converters, but also to confirming the expected suppression in eddy current losses. Toward this end, a new test system for measuring core losses at high frequency and high magnetization levels was constructed. Finally, advanced numerical curve fitting techniques are used to elucidate the proportion of measured core losses that can be conclusively attributed to eddy current effects.

5.2 *Design*

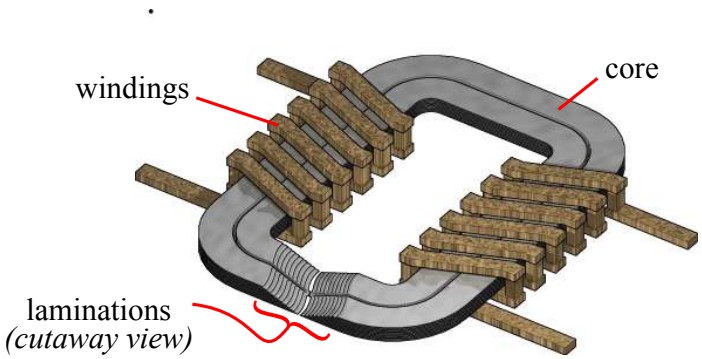

Figure 70: Conceptual Representation of Highly Laminated-Core Transformer

The generic design of the highly laminated inductor is shown in Figure 70. Three are three main components to the design of this inductor: magnetic core design, winding design, and packaging design. The core design was based largely on the knowledge gained in the extending of the sacrificial multilayer process, as presented in Chapter 3. The winding design entailed navigating an array of potential approaches. The packaging design required simultaneously addressing a number of needs, including winding integration, mechanical protection of the laminated core, and easy interconnection into test systems.

5.2.1 Core Design

5.2.1.1 Magnetic Design

Assumptions

In analyzing and designing the magnetic operation of the inductor's cores, a few simplifying assumptions were made. First, ideal symmetry was assumed along the thickness direction. This allowed the core's design to be considered in only the two lateral dimensions.

Second, it was assumed that *all* magnetic flux would arrange into eikonally- related paths within the magnetic core. This concept is illustrated in Figure 71. It enabled a straightforward analytical approach using the eikonal analysis conventions detailed in Appendix A. It also includes the implicit assumption that all magnetic flux is contained within the core (*i.e.* zero leakage flux).

Finally, it is assumed that the magnetic circuit behaves symmetrically along the length of each path. Mostly, this assumption amounts to ignoring any crowding of flux or in-plane eddy currents.

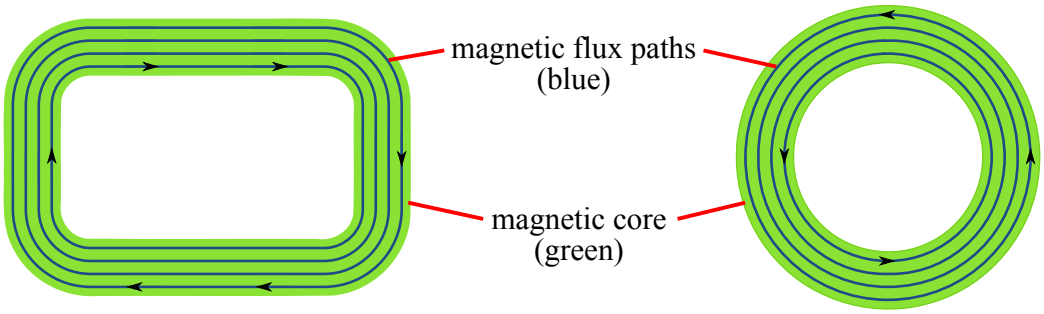

Figure 71: Organization of magnetic flux into eikonally-related paths

This is a highly expedient framework for the analytical modeling carried out below.

Closed Core

It was decided that the cores would be closed cores, meaning they would not incorporate any longitudinal air gaps. While including an air gap would allow better control over the inductor's device-level behavior, particularly controlled inductance and increased saturation current, not including the air gap conferred a number of compelling advantages to this of research.

First and foremost, it was found that these cores would not produce very large inductances – in most cases, not much higher than 1 μH. Compared to conventional cores, it is the extremely low thickness of these microfabricated cores that accounts for their low inductance. Higher inductance is desirable not only for use of these inductors in practically-sized SMPSs, but also to limit the magnitude of required capacitance for the chosen balanced reactivity measurement scheme.

Second, a non-gapped core strengthens the grounds for the assumption that all flux is contained in the core, and that it arranges into eikonally-related paths. Including an air gap would not only raise the core's reluctance, but would also introduce a potentially sharp sensitivity to vertical misalignment. These factors would not only exacerbate field fringing at the air gap, but a poorly controlled and potentially intractable quantity of fringing. Furthermore, any curvature of the cores near the air gap would raise the spectres of self-demagnetization and oblique eddy currents at the boundaries of the air gap, again introducing potentially dominant uncertainty into the core loss measurements.

Finally, a non-gapped core would offer secondary advantages in the process of gaining insight into core losses. This is mostly because greater flux will be generated with a smaller amount of coil current. This will reduce the magnitude of nuisance losses associated with coil current, such as image eddy currents and resistive losses. As well, by removing the requirement for mechanical

means to maintain alignment across a gap, many unknowns, and potential parasitic pathways (*ex.* eddy currents in a silicon substrate) can be removed from the picture.

Eikonal Modeling

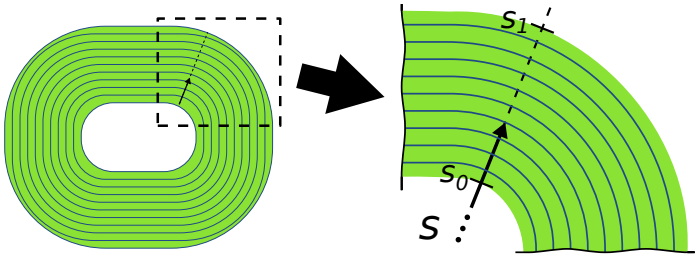

Figure 72: Organization of magnetic flux into eikonally-related paths

Figure 72 shows the orientation of the eikonal coordinate s used in modeling of these cores. s is always taken to be perpendicular to the flux paths, and is equal to s_0 at the inner margin of the core, and equal to s_1 at the outer margin of the core. In the simple case of a ring-shaped core, s is identical to the radial coordinate r in a polar coordinate system.

Most importantly, for each value of s, a path length $l(s)$ is defined. Due to the simple geometric nature of the cores analyzed, and the assumption of flux path geometries, the path length $l(s)$ varies linearly with s. l_0 is defined as the shortest path length for a given core, or, equivalently, $l(s_0)$; likewise, l_1 is the longest path for a given core, occurring by assumption at its outer margin, or where $s = s_1$.

With these simple conventions, the expected reluctance R of a given core can be easily derived. First, since the integration will be done across parallel core segments, the core's behavior is treated in terms of admittance G, such that

$$G = \int dG = \frac{1}{R} \tag{44}$$

The admittance of a given segment, corresponding to a differential increase ds in the s coordinate, is

$$G = \frac{wh}{l}\mu \tag{45}$$

where h is the height of the given segment (left unbounded throughout this analysis), w is the width of the segment, and μ is the material's permeability. Using this to develop dG, a differential

contribution to the total admittance G, gives

$$dG = \frac{h\,ds}{l(s)}\mu_R\mu_0 \tag{46}$$

with the width w taken to be ds, and the permeability rewritten as the product of μ_R and μ_0, relative and free-space permeability, respectively. Carrying out the integration over s gives

$$G = \int dG \tag{47}$$

$$= \int_{s_0}^{s_1} \left[\frac{h\,ds}{l(s)}\mu_R\mu_0\right] \tag{48}$$

$$= h\mu_R\mu_0 \int_{s_0}^{s_1} \frac{1}{l(s)} ds \tag{49}$$

A change of coordinates from s to l, based on their linear relationship

$$l(s_0) = l_1 \tag{50}$$

$$l(s_1) = l_0 \tag{51}$$

$$\frac{dl}{ds} = \frac{\Delta l}{\Delta s} \tag{52}$$

$$= \frac{l_1 - l_0}{s_1 - s_0} \tag{53}$$

gives

$$G = h\mu_R\mu_0 \int_{s_0}^{s_1} \frac{1}{l(s)} ds \tag{54}$$

$$= h\mu_R\mu_0 \int_{l_0}^{l_1} \frac{1}{l} ds \left[\frac{dl}{ds}\right]\left[\frac{dl}{ds}\right]^{-1} \tag{55}$$

$$= h\mu_R\mu_0 \int_{l_0}^{l_1} \frac{1}{l} dl \left[\frac{dl}{ds}\right]^{-1} \tag{56}$$

$$= h\mu_R\mu_0 \left(\ln l\right|_{l_0}^{l_1} \frac{s_1 - s_0}{l_1 - l_0} \tag{57}$$

$$= h\mu_R\mu_0 \frac{s_1 - s_0}{l_1 - l_0} \ln \frac{l_0}{l_1} \tag{58}$$

Given a particular core thickness h and number of turns N_t, the expected inductance $L(h, N_t)$ is then given by

$$L(h, N_t) = \frac{N_t^2}{R} \tag{59}$$

$$= N_t^2 G(h) \tag{60}$$

$$= \left[N_t^2 h\right] \mu_R\mu_0 \frac{s_1 - s_0}{l_1 - l_0} \ln \frac{l_0}{l_1} \tag{61}$$

With the introduction of a core-specific figure of merit A_L, the inductance calculation can be, for practical purposes, greatly simplified:

$$L(h, N_t) = \left(N_t^2 h\right) A_L \tag{62}$$

$$A_L = \mu_R \mu_0 \frac{s_1 - s_0}{l_1 - l_0} \ln \frac{l_0}{l_1} \tag{63}$$

5.2.1.2 Detailed Design

Plating Mask Implementation

The assumption that the magnetic flux would arrange into eikonally-related paths was heavily relied upon in the design of the cores' plating masks. Fundamentally, the core was laminated along the s coordinate (*i.e.* transverse direction), meaning it was divided into a number of geometrically alike, concentric "tracks." This is shown in Figure 73. By assumption, this division will not affect the cores' magnetic behavior, since the magnetic flux is presumed to flow parallel to the trenches formed between tracks.

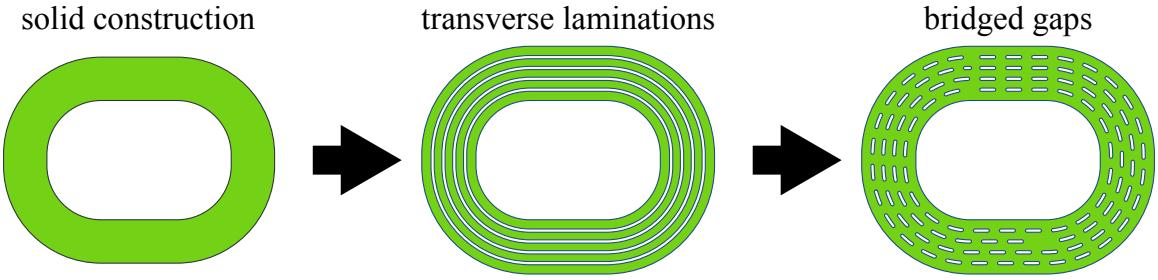

Figure 73: Conceptual Evolution of Plating Mask for Laminated Magnetic Cores

Introducing these transverse divisions was desirable for two reasons. First, it would limit the potential impact of any in-plane eddy currents by cleaving these currents' paths. These were not expected to be a factor, as they are not conventionally considered for flat laminations. However, if the microfabricated laminations developed any out-of-plane curvature, these eddy currents might introduce uncertainty. More importantly, the trenches produced by these transverse divisions facilitated mass transfer during the laminations' release etch.

As a practical matter, it was found that it was advantageous to periodically bridge the trenches between tracks; this, too, is shown in Figure 73. These bridges kept each core layer as a laterally contiguous structure, greatly increasing the core's mechanical integrity.

As well, these bridges play an important role during the multilayer plating process. The formation of an unbridged trench by plating through a mask would necessitate a continuous track of photoresist. This continuous track would be subject to stress-induced deformation and, perhaps, delamination during the plating process. The bridges across the gap amount to breaks in this continuous track of photoresist, eliminating the accumulation of intrinsic stresses across the track's entire length.

Core Variants

A number of core geometries were designed, but the inductance factor A_L was held to one of two constant values across all designs. While many of these geometries were successfully released and qualitatively confirmed as inductors, only ring-shaped cores with A_L of 198 pH/(μm ·turn^2) were characterized.

Table 5: Microfabricated Core Plating Mask Geometries

Ind. factor A_L [pH/μm ·turn^2]	Shape of path at $s = s_0$	s_0 [mm]	s_1 [mm]	Number of tracks	Track width [μm]	Etch hole width [μm]
198	circle	1.5	5.0	14	205	50
198	circle	1.5	4.6	5	540	100
198	circle	1.5	4.3	5	520	50
198	circle	1.5	4.6	5	540	100
130	circle	2.0	4.4	9	220	50
130	circle	2.0	4.1	4	475	50
130	circle	2.0	4.2	4	480	100
130	square	3.0	7.8	9	220	50
130	square	3.0	7.1	4	475	50
130	square	3.0	7.5	4	490	100
130	2:1 rectangle	2.0	6.1	9	470	50
130	2:1 rectangle	2.0	6.5	4	490	100
130	4:1 rectangle	1.5	7.3	10	245	50
130	4:1 rectangle	1.5	6.7	5	480	50

5.2.2 Winding Design

Several options were considered for adding the windings to the cores, to yield inductors. Microfabricated windings might have offered considerable advantages, including higher turn count, reduced handling demands during sacrificial release of the core layers, and direct integration onto other microfabricated systems. However, the development of microfabricated windings would have added considerable scope and unknowns into the primary figures of merit, which were the losses within

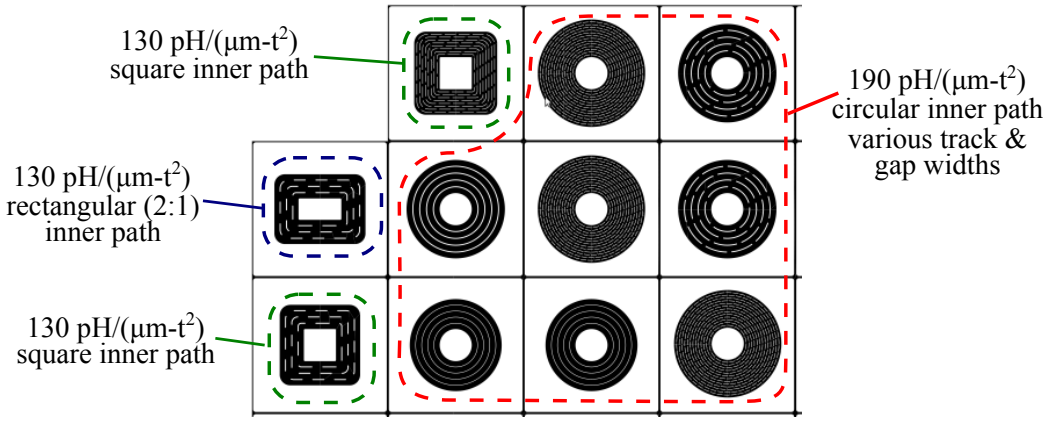

Figure 74: Plating Masks for Magnetic Cores. The rectangular grid was part of the mask used. It served to prevent the lateral accumulation of shear stress in the thick photoresists used.

the core.

Accordingly, manually wound wire was used to create the inductor windings. Modern magnet wire is available in in diameters down to 20 μm – smaller than could have been conveniently attainable with a 3-dimensional winding microfabrication process. Keeping the diameter of the copper conductors on the order of the skin depth of copper at the frequencies of interest (60 μm at 1 MHz to 20 μm at 10 MHz) is an important part of minimizing the winding losses, and obtaining a tractable model of core losses. Also, modern magnet wire is readily available in litz wire configuration, a multi-conductor arrangement which contains multiple strands of individually insulated wire, intertwined in a pattern that distributes conductor proximity effects equally between all conductors. This helps not only minimize losses, but also make the winding losses more repeatable and predictable between devices.

5.2.3 Packaging Design

Much of the top-level device packaging approach chosen for these inductors is embodied in the design of the bobbin assembly. The construction of the bobbin is shown in Figure 75. This bobbin has three critical provisions:

- an annular cavity for holding the magnetic core, including thick outer walls to prevent contact between the litz wire and the top surface of the magnetic core;
- regularly-spaced slots to guide the litz wire, so that a consistent winding arrangement, in

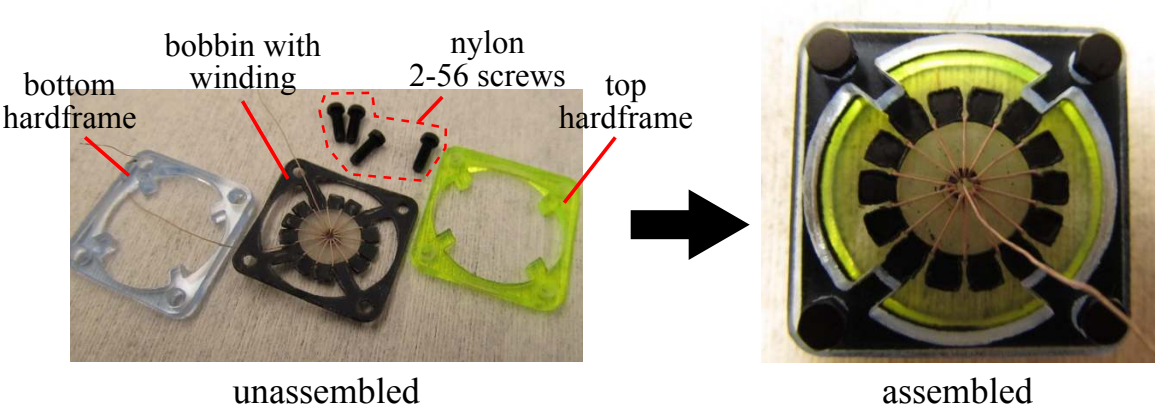

Figure 75: Full bobbin assembly

terms of both turns count and spatial configuration, can be achieved across all characterized samples; and

- a thick, rigid hardframe for manual handling, with highly compliant mechanical connections to the inner portion of the bobbin, to reduce the chance that improper handling would induce any bending of the magnetic core.

The need for each of these provisions was discovered through trial and error, particularly driven by the yield of the inductor assembly and test process. In particular, the hardframe feature evolved over several batches of fabrication and assembly. In addition to allowing for robust manual handling, the hardframe also protects the delicate magnetic core from incidental damage and also articulates with features in the characterization PCBs, again providing repeatable interconnect.

The fabrication of the bobbin assembly consists of two separate assembly steps, as suggested by Figure 75: the fabrication of the bobbin proper, and the assembly of the hardframe onto the bobbin. Four #2-56 screws, one in each corner, are used to hold the entire assembly together. The holes in the bottom hardframe are threaded to receive the #2-56 screws.

The bobbin itself was fabricated with a laser machining process, diagrammed in Figure 76. Two dissimilar polymer materials – ultra high molecular weight (UHMW) polyethylene (PE) and polyester – were laminated together and then laser machined. By varying the laser machining parameters, mainly laser power and sweep speed, it was possible to selectively machine this laminate. The UHMW-PE was extremely resistant to the 10.4 μm CO_2 laser used, while the polyester was

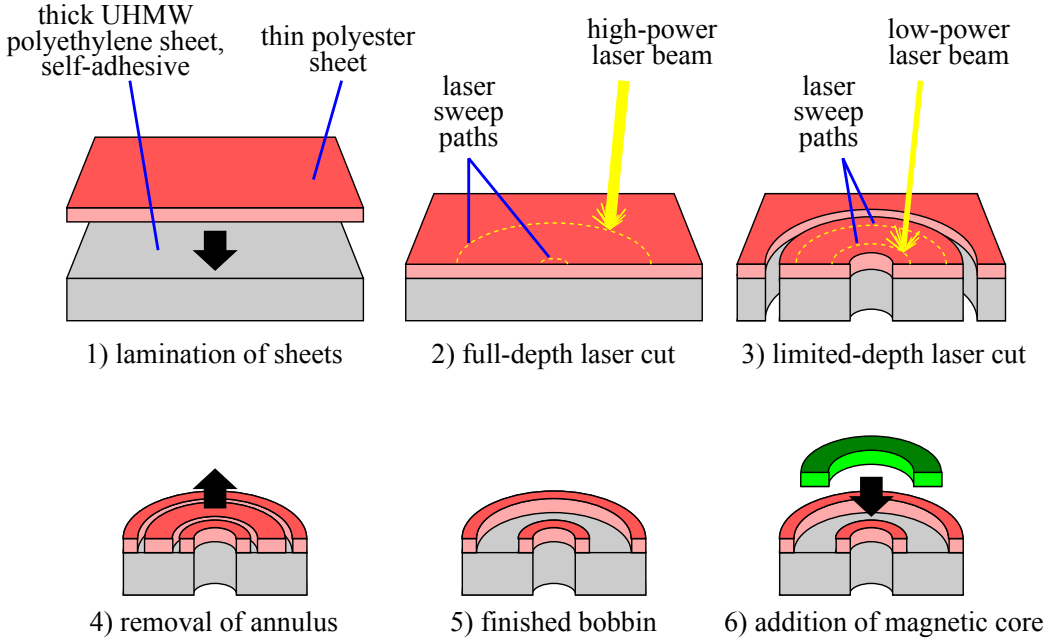

Figure 76: Bobbin fabrication by laser machining (cross-section)

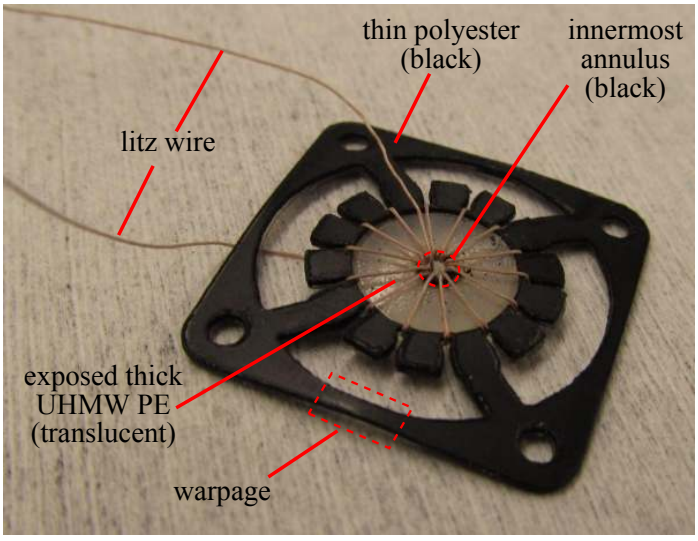

Figure 77: Completed bobbin with winding. The pictured device is also the air-core sample "AC14," which was used to characterize winding inductance and power loss.

readily ablated by the same laser. Beyond their intrinsic difference in laser machinability, the difference in thickness of these layers – 0.25mm for the UHMW-PE layer *vs.* 0.10mm for the polyester layer – accentuated their machinability. A completed bobbin is shown in Figure 77.

5.3 Core Fabrication

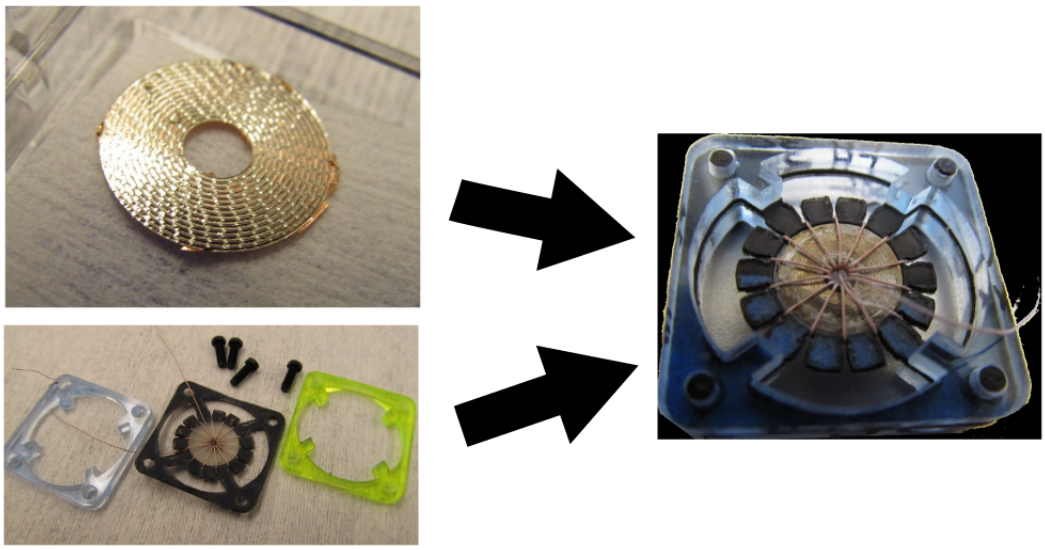

Figure 78: Assembly of Laminated-Core Inductor

The magnetic cores were fabricated according to the sacrificial multilayer plating process described in Chapter 3. Much of the progress presented in Chapter 3 was achieved while pursuing the fabrication of the cores detailed in this chapter. The final set of samples used for core loss characterization is listed in Table 6.

Table 6: Laminated Core Samples

Wafer Name	Number of Layers	Plating Time per Layer	Estimated Layer Thickness	Estimated Core Thickness (Total)
SFH6	200	60 s	0.17 μm	33 μm
SFH5	100	120 s	0.33 μm	33 μm
SFH4	66	360 s	1.0 μm	66 μm
SHT3	25	480 s	1.33 μm	33 μm
SHT2	20	600 s	1.67 μm	33 μm

Figure 79 shows some magnetic cores immediately after plating and photoresist removal, still on the silicon substrate. Figure 80 shows a sample that has received a partial etch, so that the layers are separate and clearly visible. Notably, the layer thickness and interlayer gaps are near the imaging

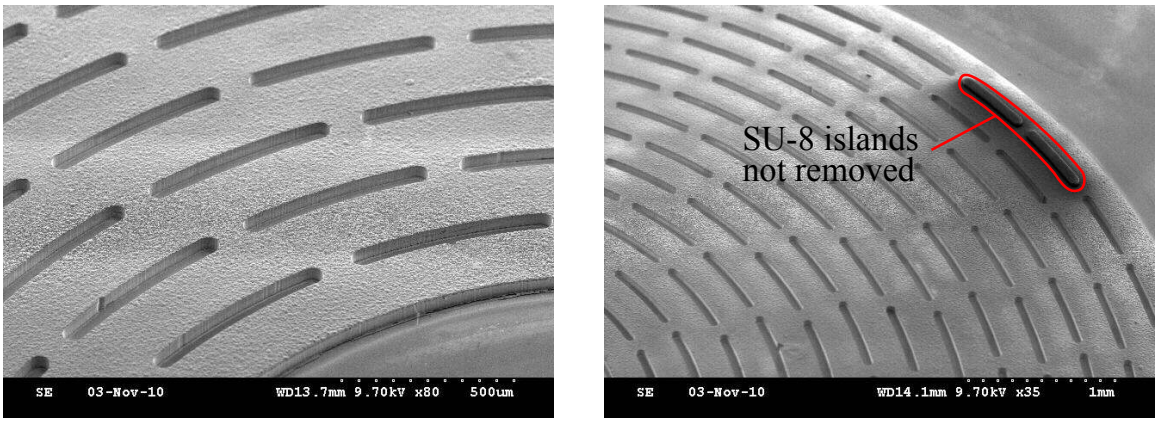

(a) (b)

Figure 79: Laminated Cores on Wafer

resolution limit of the SEM used (Hitachi 3700). Figure 81 shows magnetic cores in the final liquid

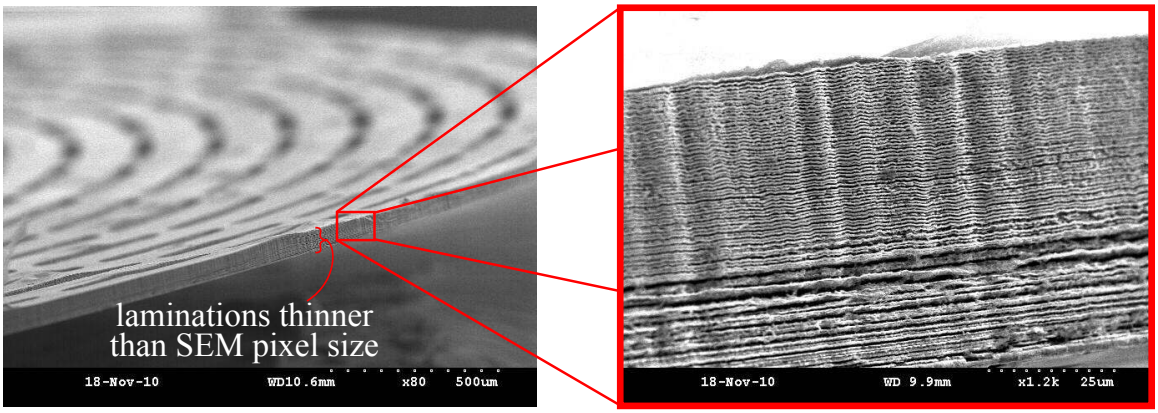

Figure 80: SEM micrograph showing multiscale character of laminated cores

release step, which is the infiltration of polystyrene-bearing MEK into the interlayer gaps.

5.4 Approach to Inductor Characterization

5.4.1 Emphasis on Power Loss Measurement

The highly laminated construction of these inductors' cores, and the associated intention of reducing eddy current losses in the core, serves as the centerpiece of both the devices' purpose as well as their implementation. Accordingly, the characterization of the inductors was focused on exploring the effectiveness of this construction at controlling eddy current losses.

However, given the lack of experimental methods for directly measuring eddy current magnitudes, the characterization could be focused no narrower than on the total core loss, and, in fact, on

magnetic cores in final
methy ethyl ketone (MEK) release

magnetic core disintegrating
during final MEK release

Figure 81: Laminated cores undergoing final release (cf. Figure 51, Chapter 3)

the total power dissipation within the inductor. From these total power loss results, known effects could be removed through numerical compensation, and core loss mechanisms could be separated through numerical curve-fitting. The mathematical theory behind the latter procedure is detailed in Appendix B.

5.4.2 Survey of Loss Characterization Approaches

There are a number of techniques for measuring energy dissipation within a magnetic material. There are three governing assumptions across all such techniques:

- first, that all energy added to the core is carried over the terminals of whatever windings are utilized;
- second, that all energy added to the core is either returned via the windings' terminals, or dissipated as heat within the core; and,
- third, that all energy dissipated as heat within the core is classified as core loss.

The second assumption implies that all all energy added to the core is removed only by electromagnetic induction or by thermal conduction or convection. This in turn entails that any other phenomena capable of transferring energy away from the core, including mechanical actuation and electromagnetic radiation, may be ignored.

5.4.2.1 Calorimetric

Calorimetric loss measurement exploits the assumption that all core loss manifests as heat produced internal to the core. Operating the core at a controlled excitation level for a known amount of time, then, amounts to the production of a generally unknown quantity of heat energy within an enclosed volume. Measuring this unknown quantity of heat, is precisely the problem addressed by the technique of calorimetry.

In a calorimetric measurement, the unknown quantity of released heat energy, Q_{unk} is confined within a system with a well-known heat capacity m_C. This heat produces a rise in temperature ΔT, and these quantities are related as

$$Q_{unk} = m_C \Delta T \tag{64}$$

establishing that if m_C is known and ΔT can be accurately measured, Q_{unk} can be trivially calculated [76].

In a calorimetric core loss measurement application, Q_{unk} would be equal to the total energy lost throughout the entire core volume over t_{on}, the entire time that the core is energized:

$$Q_{unk} = \int_0^{t_{on}} P_{core}(t)\, dt \tag{65}$$

where $P_{core}(t)$ is the total power loss within the entire core volume at time t. Since calorimetric measurements necessarily involve some amount of temporal averaging, it is convenient to regard core loss as a time-invariant quantity, symbolized simply as P_{core}.

$$Q_{unk} = t_{on} P_{core} \tag{66}$$

This is a well-justified assumption since, at frequencies on the order of 1 MHz, magnetic excitation, and presumably core losses, will vary over timescales 1 μs or less, while the run times t_{on} will need to be on the order of 1 second or more to achieve substantial heat accumulation.

The primary difficulty with calorimetric systems is controlling parasitic heat flows. Considering all parasitic heat flows as a single, time-averaged flow P_p, and assuming that these occur not only over the core energization period t_{on} but also the additional time t_{meas} required to measure the system

temperature, the extracted magnitude of Q_{unk} becomes

$$Q_{unk} = t_{on}P_{core} + [t_{on} + t_{meas}]\,P_p \tag{67}$$

$$= t_{on}\left[P_{core} + P_p\left(1 + \frac{t_{meas}}{t_{on}}\right)\right] \tag{68}$$

It can be seen, then, that any uncertainty in P_p will translate into a larger amount of uncertainty in P_{core}. Minimizing the impact of P_p can be accomplished by minimization of its magnitude or by highly accurate characterization. As P_{core} becomes very small, the most minor components of P_p, such as conduction through winding leads, blackbody radiation, and even conduction into moderate vacuum environments can become dominant limitations. As well, highly accurate characterization of P_p can become onerous, involving characterization not only over internal calorimeter temperature T, but also over ambient temperature T_{amb}.

As well, accurate determination of the temperature rise ΔT itself can become difficult if it is not large enough. Again, as P_{core} becomes small, achieving a substantial ΔT will require minimizing the heat capacity m_C. However, since the calorimetric system must involve, at a bare minimum, the core itself, absolute minimization of m_C means having the core material itself dominate m_C. This may introduce an unacceptably high lower limit on m_C by itself. But, the core's heat capacity is given by:

$$m_C = m_{core}C \tag{69}$$

$$= \rho v_{core}C \tag{70}$$

where ρ is the density of the core material, C is its intrinsic heat capacity, m_{core} is the total core mass, and v_{core} is the total core volume; This means that uncertainty in determining any one of the three quantities occurring on the right hand side of Equation 69 can introduce dominant uncertainty into m_C.

Calorimetric core loss measurement is particularly incompatible with the highly laminated cores presented in this thesis. The high aspect ratio trenches formed between laminations can trap residual material, introducing relatively high uncertainty into measurements of the core's mass or volume. This, in turn, limits the accuracy with which the core's heat capacity can be determined.

Another issue with using calorimetric loss measurement in this work is that the high aspect ratio of the laminations themselves creates can create a high-impedance path for heat conduction to the

periphery of the core. This can create significantly uneven thermal profiles within the core, adding a nuisance factor to experimental measurements of core loss. As well, it directly increases the thermal time constant for the system, entailing that a large settle time be incorporated into the temperature measurement time, consequently compounding the impact of parasitic thermal paths. The impact of this can be seen by considering Equation 68, where t_{meas} is the measurement time, and P_p represents any parasitic power leakage.

5.4.2.2 Wattmeters

The assumptions involved with conventional core loss measurement equate the core loss with the net electrical energy delivered to the core, via the attached windings. Therefore, knowing only the current and voltage on the windings, it is possible to determine the energy lost in a magnetic core.

$$P(t) = v(t)\, i(t) \tag{71}$$

$$W(t_1, t_2) = \int_{t_1}^{t_2} P(t)\, dt = \int_{t_1}^{t_2} v(\tau)\, i(\tau)\, d\tau \tag{72}$$

$$P_{avg} = \frac{1}{T} W(t, t+T) = \frac{1}{T} \int_{t}^{t+T} v(\tau)\, i(\tau)\, d\tau \tag{73}$$

Given any two-terminal electrical device, Equation 71 can be used to calculate the net electrical power delivered to the device, and Equation 73 can be used to calculate the averaged power dissipation.

Wattmeters are instruments that carry out precisely this procedure: they measure the device's voltage and current, and calculate the net power. If a magnetic core has only one winding attached, and the losses in the winding are known, its core loss can be measured with a wattmeter.

Hall Effect Wattmetering

Hall effect wattmeters make use of the implicit multiplication effected by Hall effect sensors. Their conceptual operation is shown in Figure 82. [1] The Hall sensor generates a voltage v_H at its output terminals that is proportional to both the current passing through it, i_H, and to its perpendicular magnetic field B:

$$v_H(t) = R_H i_H(t)\, B(t) \tag{74}$$

where R_H is the Hall constant. A Hall effect wattmeter is created by arranging the circuit so that the two time-dependent quantities on the right-hand side of Equation 74 become proportional to

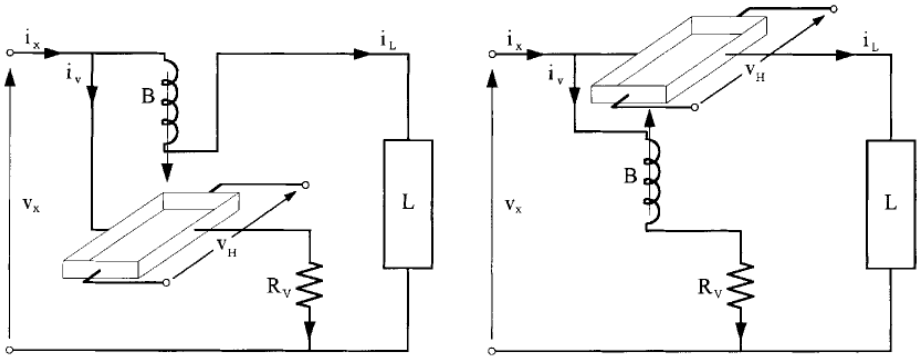

Figure 82: Two Varieties of Hall Effect Wattmeters. In both,the ability to create a magnetizing current i_V that is proportional to V is essential.

$i(t)$ and $v(t)$, respectively, the current through and voltage across the two-terminal device under test (DUT).

Two possible arrangements are shown in Figure 82. In both arrangements, a resistor R_V, with a relatively low resistance, is used to create an energizing current $i_v(t)$ that is proportional to $v(t)$. One of the two currents $i_v(t)$ or $i(t)$ is passed through a coil that produces an appropriately-directed magnetic field $B(t)$, of directly-proportional flux density. The current not passing through the coil is then used to energize the Hall effect sensor (*i.e.* as $i_H(t)$ in Equation 74.

Consequently, in ether Hall effect wattmeter arrangement, the Hall sensor output voltage $v_H(t)$ is directly proportional to the DUT's power dissipation. This voltage can then be integrated, either digitally or with an analog integrator, to give an accurate measurement of average power.

Thermocouple-based Wattmeter

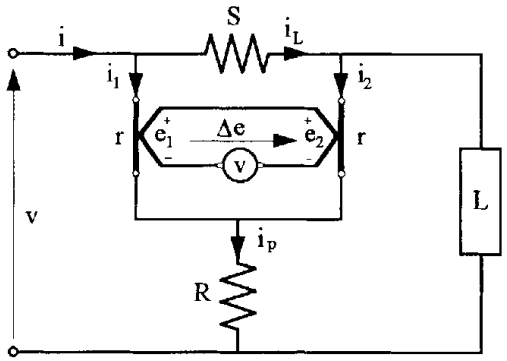

Figure 83: Thermocouple-based Wattmeter

Figure 83 shows a thermocouple- based wattmeter. [1] It is assumed that the two thermocouples

are perfectly matched, and that they generate voltages e_1 and e_2, which are proportional to their internal power dissipation, which is proportional to the square of their currents, i_1 and i_2 respectively. It is also assumed that the bias resistor R is much larger than the thermistor resistances r, which are also assumed to be approximately constant over temperature. In this case, the difference between the thermocouples' currents can be taken as

$$i_1 - i_2 = \frac{S\, i_L}{2R} \tag{75}$$

and the difference between their produced voltages becomes

$$\Delta e = k_1 \left(i_1^2 - i_2^2\right) \tag{76}$$

$$= k_1 \left(i_1 + i_2\right)\left(i_1 - i_2\right) \tag{77}$$

$$= k_1 \left(i_p\right)\left(\frac{S\, i_L}{2R}\right) \tag{78}$$

$$= k_1 k_2 \left(v_L i_L\right) \tag{79}$$

Thus, the voltage difference Δe can then be used as a dynamic output that is exactly proportional to the DUT's consumed power.

Note that this approach relies on two implicit physical assumptions. First, this wattmeter requires a linear relationship between the power dissipated in a thermocouple and its corresponding temperature increase. This requires a constant thermal resistance between the thermocouple and the ambient environment; certain cooling mechanisms create a nonconstant thermal resistance. As well, the thermal time constant of the thermocouples must be much smaller than the inverse of the frequencies being passed through the DUT. This not only limits the maximum electrical frequency that can be accurately measured with a given system, but also creates pressure to minimize the thermal resistance between the thermocouples and the ambient environment.

The latter dynamic is problematic because minimizing the thermal resistance will increase the bandwidth of the system, but it will also simultaneously minimize the temperature rise and, consequently, the thermocouple voltages e_1 and e_2 and the overall sensitivity of the system. This implies that the thermal system design would need to be tuned as the frequencies and magnitudes of the signals under test change.

Sampling Wattmeter

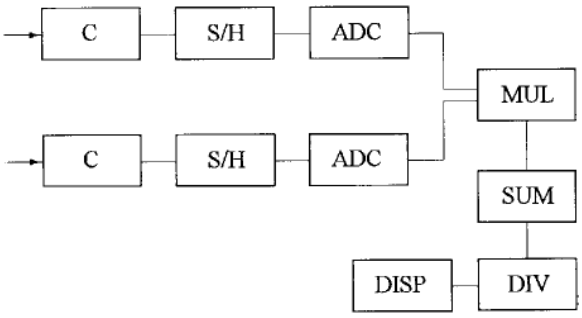

Figure 84: Block diagram of a sampling wattmeter. The operation is straightforward: two analog voltages, one proportional to the current through the DUT, and the other proportional to the voltage across the DUT, are sampled and converted to digital codes, where standard digital arithmetic is used to calculate average power dissipation within the DUT.

Although thermocouple-based and Hall sensor-based wattmeters offer simple construction and considerable robustness, they are limited to the frequency, voltage, and current regimes over which their critical approximations hold. As well, their accuracy is subject to factors that affect the physical behavior of their analog elements.

One way to widen the operating range of these wattmeters is to effectively perform the calculation in Equation 71 with a digital computer. This is the general approach described as sampling wattmetering. The instantaneous values of the DUT's voltage and current values are sampled at particular times.

5.4.3 Balanced-Reactivity Loss Measurement

5.4.3.1 Wattmetering for High-Q Inductors

With sinusoidal excitation – the simplest of cases – wattmetering involves the measurement of three generally independent quantities: the magnitude of v, the magnitude of i, and the temporal relation between v and i. Any relative uncertainty in v and i will be included in the uncertainty for the calculated power. [10].

This can be seen in both measurement of complex impedance as well as measurement of power. Consider a complex impedance given by $\hat{Z}$ in the two equivalent forms

$$\hat{Z} = R + jX \tag{80}$$

$$= |Z| \{\cos\theta + j\sin\theta\} \tag{81}$$

with the familiar conventions

$$R = |Z| \cos\theta \tag{82}$$

$$X = |Z| \sin\theta \tag{83}$$

Assuming that $|Z|$ and θ are the directly measured quantities, the experimental measurements $\left|Z_{exp}\right|$ and θ_{exp} include some random uncertainties e_Z and e_θ, respectively.

$$\left|Z_{exp}\right| = |Z| + e_Z \tag{84}$$

$$\theta_{exp} = \theta + e_\theta \tag{85}$$

The impact of these uncertainties on R_{exp}, the experimentally-extracted value of R, is estimated as a linearization about the true value of R:

$$R_{exp} = R + e_R \tag{86}$$

$$e_R = R - R_{exp} \tag{87}$$

$$= |Z| \cos\theta - \left|Z_{exp}\right| \cos\theta_{exp} \tag{88}$$

$$= |Z| \cos\theta - \Big(|Z| + e_Z\Big) \cos\left(\theta + e_\theta\right) \tag{89}$$

$$\approx \frac{\partial R}{\partial |Z|} e_Z + \frac{\partial R}{\partial \theta} e_\theta \tag{90}$$

$$\approx (\cos\theta)\, e_Z + \Big(-|Z| \sin\theta\Big) e_\theta \tag{91}$$

$$\approx e_Z \cos\theta - e_\theta |Z| \sin\theta \tag{92}$$

In typical impedance analyzers, the accuracy of impedance magnitude $|Z|$ scale with the magnitude itself, while the accuracy of phase angle measurement scales with frequency f:

$$e_Z = |Z|\, \alpha_Z \tag{93}$$

$$e_\theta = k_\theta f \alpha_\theta \tag{94}$$

Converting this to a relative error α_R in R such that

$$e_R = R \alpha_R \tag{95}$$

we proceed

$$\alpha_R = \frac{e_R}{R} = \frac{1}{R}\left\{[\cos\theta]\,|Z|\,\alpha_Z + [-|Z|\sin\theta]\,k_\theta f \alpha_\theta\right\} \tag{96}$$

$$= \frac{|Z|\cos\theta}{R}\alpha_Z - \frac{|Z|\sin\theta}{R}k_\theta f \alpha_\theta \tag{97}$$

$$= \frac{R}{R}\alpha_Z - \frac{X}{R}k_\theta f \alpha_\theta \tag{98}$$

$$= \alpha_Z - k_\theta f Q \alpha_\theta \tag{99}$$

establishing that the relative uncertainty in the phase angle (θ) measurement is magnified by both frequency f and Q in its effect on the uncertainty contained in R_{exp}.

As is typically done in experimental work relying on well-calibrated equipment, [1] the uncertainties e_Z and e_θ are presumed to occupy normal distributions about a mean of zero, and to have zero correlation with each other. In this case, the expected magnitudes of the errors are related as: [10].

$$|\alpha_R|^2 = |\alpha_Z|^2 + |k_\theta f Q \alpha_\theta|^2 \tag{100}$$

and, as Q and f grow very large,

$$|\alpha_R|^2 \approx |k_\theta f Q \alpha_\theta|^2 \tag{101}$$

$$\alpha_R \approx k_\theta f Q \alpha_\theta \tag{102}$$

5.4.3.2 Resonant Inductor Loss Measurement

A direct solution to the problems posed by excessive reactive components in a wattmetering scheme is to nullify the reactive components. This can be accomplished most easily by adding in series with the DUT a component with equal and opposite reactive impedance. With the DUT's reactance balanced in this way, the only impedance presented by the circuit is the real (non-reactive) impedance of the two components. This has the ultimate effect of removing the uncertainty introduced by the reactive components. If the DUT is an inductor, its reactance can be canceled by a capacitor connected in series. This combination creates what is known as an LC circuit, or an RLC circuit if the nonzero resistances of the inductor, capacitor, and interconnect are accounted for. [24]

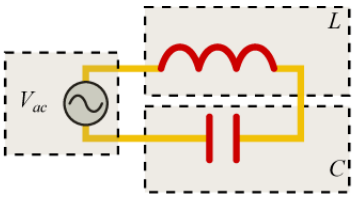

Figure 85: Idealized Schematic of Resonant Loss Measurement Circuit

5.4.3.3 Off-resonance Measurement

To reap the benefits of balanced-reactivity measurement, the LC circuit does not need to be precisely at resonance. Again framing the problem of high-Q inductor loss measurement problem in terms of experimental uncertainty, it is only necessary to reduce the uncertainty contribution of the reactive waveforms to a point where it is no longer a dominant limit. This can be accomplished with the same LC setup as described in the previous section, with the augmentation of time-domain measurements.

5.5 Characterization Results

5.5.1 Characterization System Implementation

A system to implement balanced-reactivity core loss measurement was implemented, based on the series LC approach. The capability of resolving total core loss into components associated with individual loss mechanisms, such as eddy currents or hysteresis, requires measurement of core loss over substantial ranges of both frequency and flux density. Accordingly, the loss measurement system was constructed to characterize the losses of the highly laminated cores over a range of frequencies and excitation levels.

It supports operation at frequencies from 100 kHz to 30 MHz and coil currents up to 3 A. Its design accommodates both on-resonance and off-resonance measurements, as well as extensive data collection.

5.5.1.1 Architecture and Implementation

A high-power RF amplifier, driven by a function generator, was used to excite the LC circuit, and an oscilloscope was used to capture the waveforms at both nodes of the test circuit. The core under test is wrapped with a single coil, forming the inductor L, and the capacitance C is constructed as a network of high-quality, low-resistance surface-mount capacitors. This is shown in Figure 86.

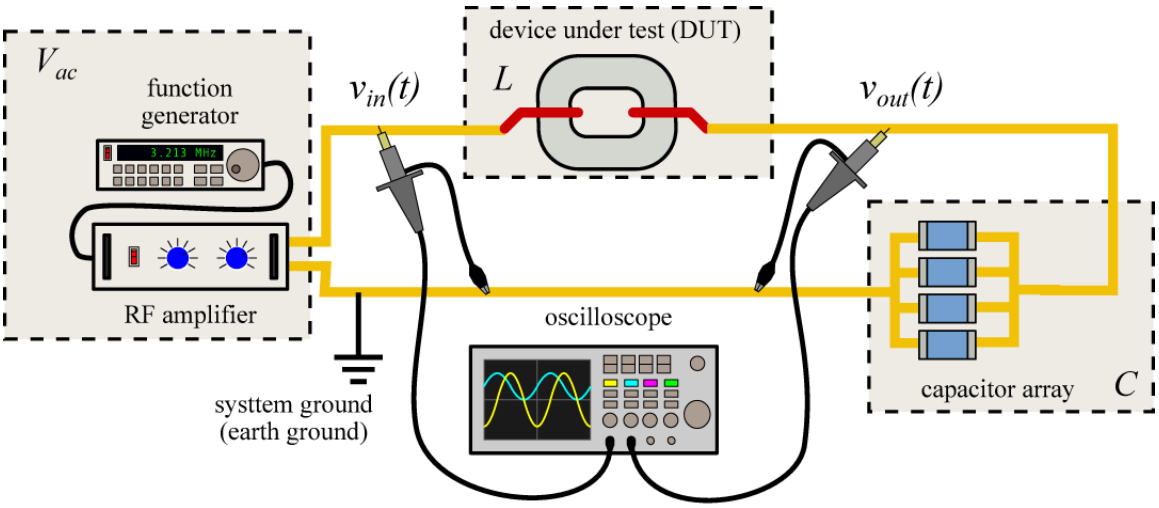

Figure 86: Implementation Schematic of Resonant Loss Measurement System

Although all test equipment utilized was capable of supporting the entire range of frequencies and drive levels desired, it was necessary to implement a number of values for the resonant capacitance. This is due to the fact that for a given DUT, having inductance L, a balanced-reactivity measurement can only be taken near one frequency f_0 for each series capacitance C.

$$f_0 \approx \frac{1}{2\pi\sqrt{LC}} \tag{103}$$

Moreover, accurate determination of L as well as the core loss itself, relies on knowledge of the

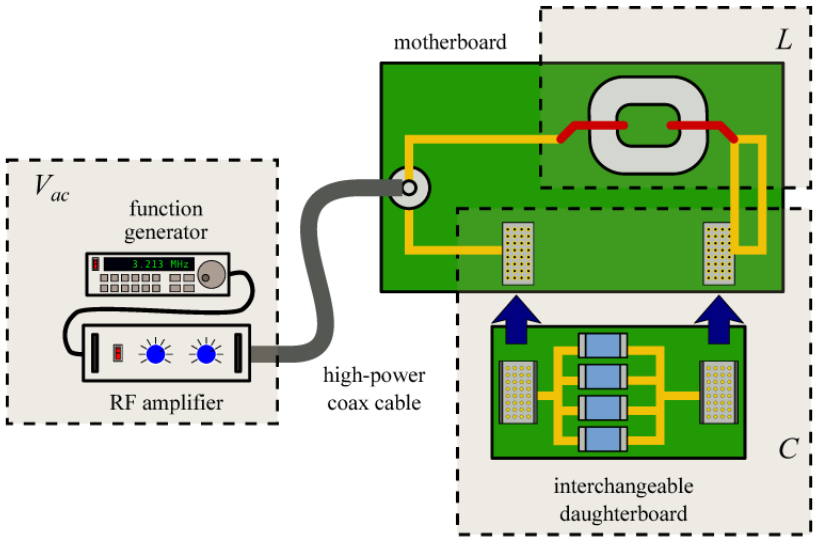

Figure 87: Full loss characterization system schematic

precise impedance of the capacitance bank. Inductive and resistive parasitics, as well as dielectric

dissipation, can significantly affect the capacitance bank impedance at higher frequencies. Also, dielectric nonlinearities can introduce up to 80% variation in capacitance over the operating voltage range, as well as undesirable distortion into the test waveforms. Therefore, it is necessary to measure the capacitance bank's impedance at each utilized frequency and excitation level. This potentially onerous need was addressed in three ways.

1. Constituent capacitors were selected to give maximal linearity with respect to voltage, to allow impedance data taken at any excitation level to be directly usable at all other excitation levels. This was entirely a matter of the dielectric material used in the selected capacitors. Fortunately, the dielectrics that offer the highest linearity also offer the lowest dielectric losses, so it was possible to simultaneously minimize capacitor nonlinearity and dissipative losses. However, low-loss, highly-linear dielectrics also tend to provide lower capacitance density, so some compromise between dielectric ideality and the physical bulk of the capacitance bank was required. Capacitors with PTFE, polypropylene, mica, and air dielectric were considered, and, due largely to their availability in high-density multilayer surface-mount packages, mica capacitors (Cornell Dubilier) were selected.

2. An impedance analyzer was used to directly measure the capacitance bank's impedance over the frequency range of interest. This approach permitted use of capacitance banks with generally arbitrary frequency behavior, and allows the system to accommodate significant parasitics throughout the entire test circuit.

3. A small number of fixed-value capacitance banks were constructed and reused throughout all of the testing. These fixed capacitance banks only needed to be measured once, obviating the need seen in [24], wherein an adjustable capacitor was used, to disconnect the capacitor from the test circuit and measure its capacitance after each single loss measurement.

Another guiding concern in the development of this test system was repeatability. This was addressed through the use of connectors wherever feasible, and the minimization of circuit adjustments that would involve soldering. Excessive soldering and manual connection not only hastens wear on the components and circuit boards, but also introduces significant variability into the circuit's behavior. Connectors were used to connect the amplifier to the test circuit, as well as to connect the

capacitance banks to the test circuit. Although fully connectorized oscilloscope probes were not used, due to cost concerns, test points were provided for use with clip-on scope probes. The test circuit was implemented, accordingly, as a motherboard-daughterboard combination, shown in Figure 87. The motherboard contained the input connector, the DUT, all scope attachment test points, as well as provisions to support an impedance matching transformer. The daughterboards contained only the capacitance banks.

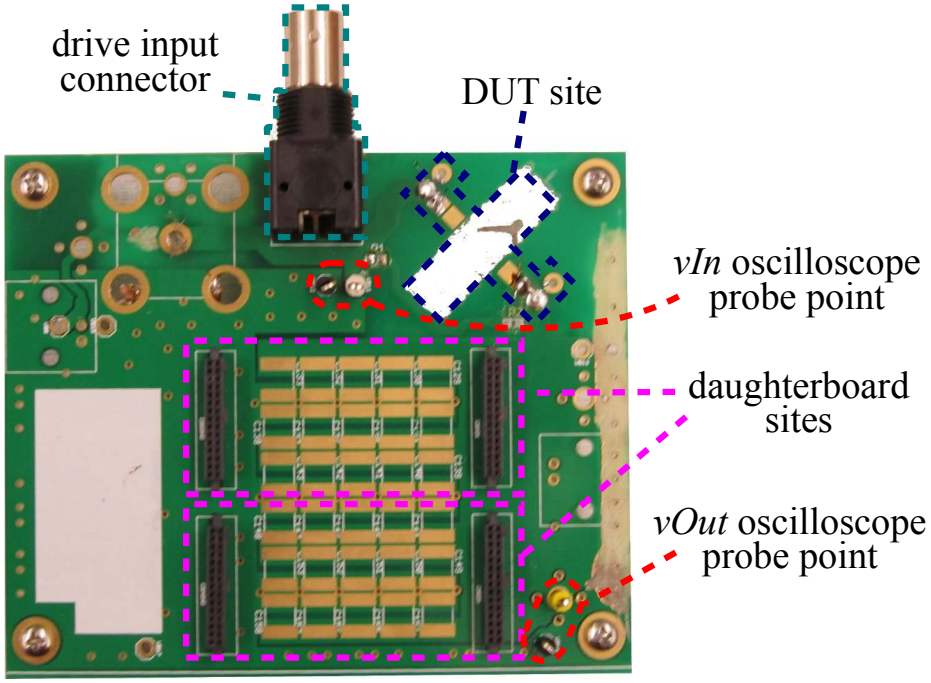

Figure 88: Motherboard for loss characterization system

5.5.1.2 Data Collection

To reduce the possibility of operator error impacting the measurements, the loss measurement system was automated as much as possible. There were two major dimensions to this automation: instrument operation and data management. Automated instrument operation consisted of the development of a software program that operated the function generator and oscilloscope via remote interfaces, and also collected all measurement data over these interfaces. The primary need in this dimension was to avoid the need for manual transcription of values and manual transfer of waveform data, both of which introduce significant risk of clerical errors. But, a workable instrument operation solution also provides a great deal of convenience, allowing the testing to be executed

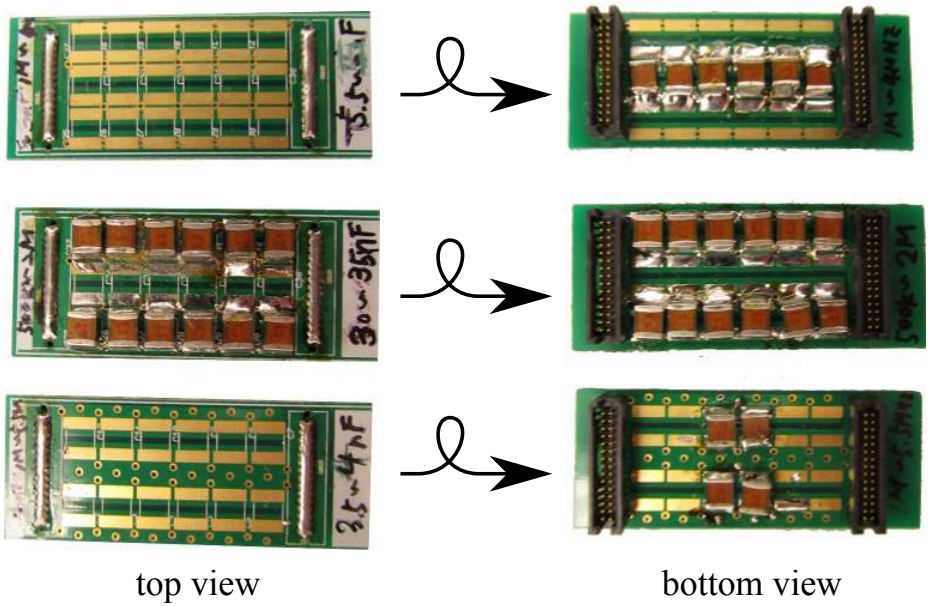

Figure 89: Daughterboards for loss characterization system

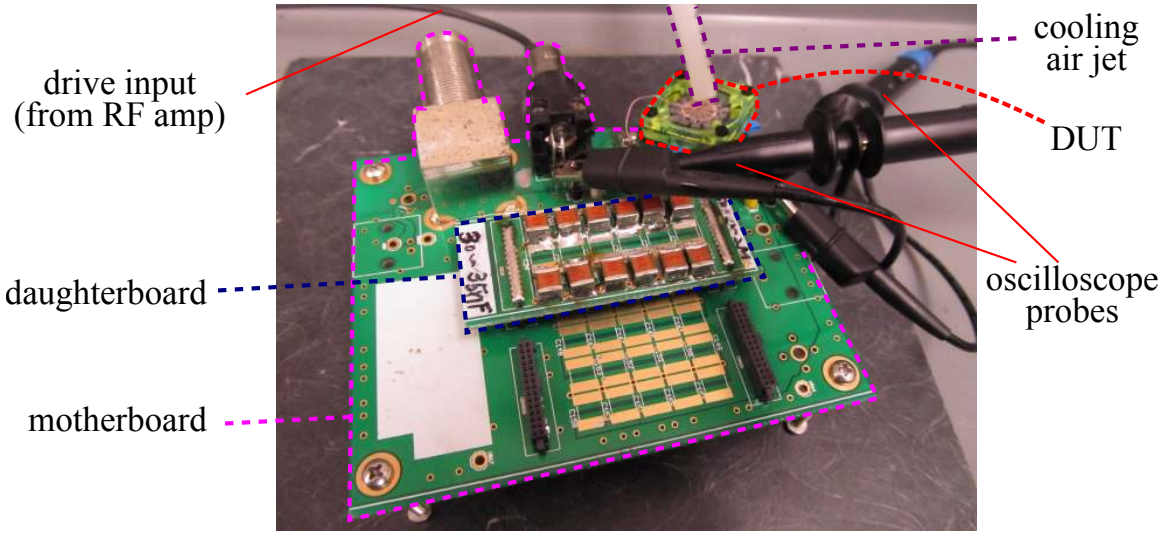

Figure 90: Fully connected loss characterization board

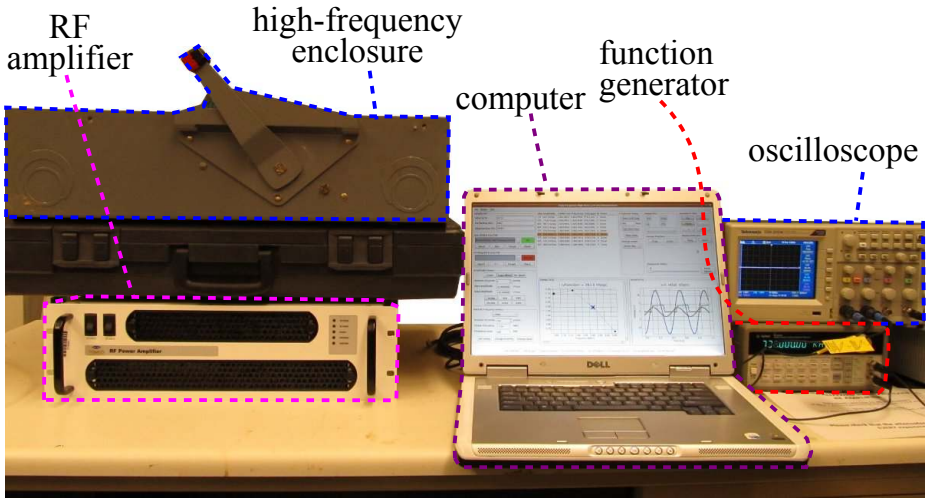

Figure 91: Complete loss characterization system

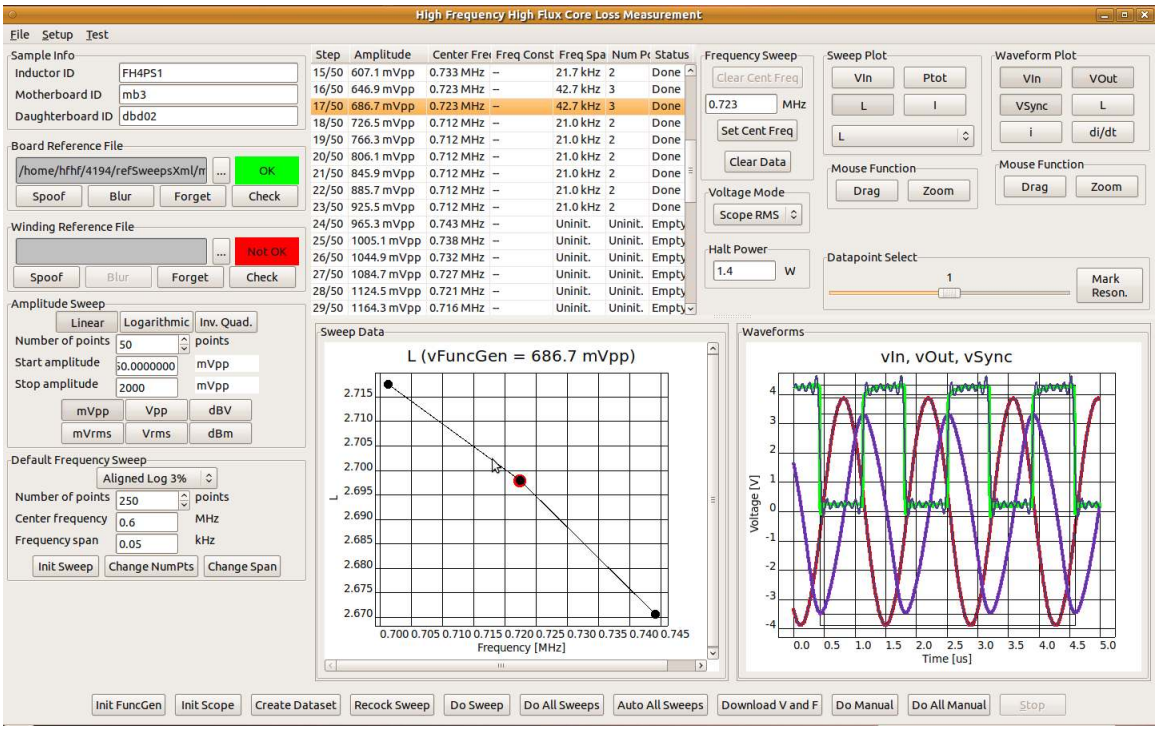

Figure 92: Screenshot of loss characterization system software

much faster.

5.5.2 Characterization System Qualification

5.5.2.1 Qualification Sample

To ensure that the loss characterization system was functioning properly, it was run with an off-the-shelf magnetic core. The magnetic core was a FT3767 ferrite core from Amidon, Inc. Based on its datasheet, the properties in Table 7 could be readily calculated.

Table 7: Qualification Sample Properties

Property	Value	Units
inner diameter	4.75	mm
outer diameter	9.53	mm
thickness	3.18	mm
relative permeability	40	-
average path length	22.4	mm
cross-section area	7.58	mm^2
volume	170	mm^3
inductance with 10 turns		nH

5.5.2.2 Expected Losses

The core's vendor supplied the core loss data for the chosen material, material "67," shown in Figure 93. From this chart, since the loss data forms a straight line when plotted on logarithmic x and y axes, an exponential relationship of the form shown in Equation 104 can be inferred.

$$P(B) = kB^{\beta} \tag{104}$$

$$\frac{P(B)}{P_0} = \alpha \left(\frac{B}{B_0} \right)^{\beta} \tag{105}$$

$$\log \frac{P(B)}{P_0} = \log \alpha + \beta \ \log \frac{B}{B_0} \tag{106}$$

At each of the three frequencies shown approximately 20 datapoints were manually measured from the graph shown in Figure 93 and fit, using simple least squares regression, to the linearized exponential relation Equation 106.

The results are shown in Table 8. To make direct use of these quantities, the relation

$$P(B) = P_0 \times 10^{\wedge} \left\{ \log \alpha + \beta \log \frac{B}{B_0} \right\} \tag{107}$$

which is equivalent to Equation 106, can be used.

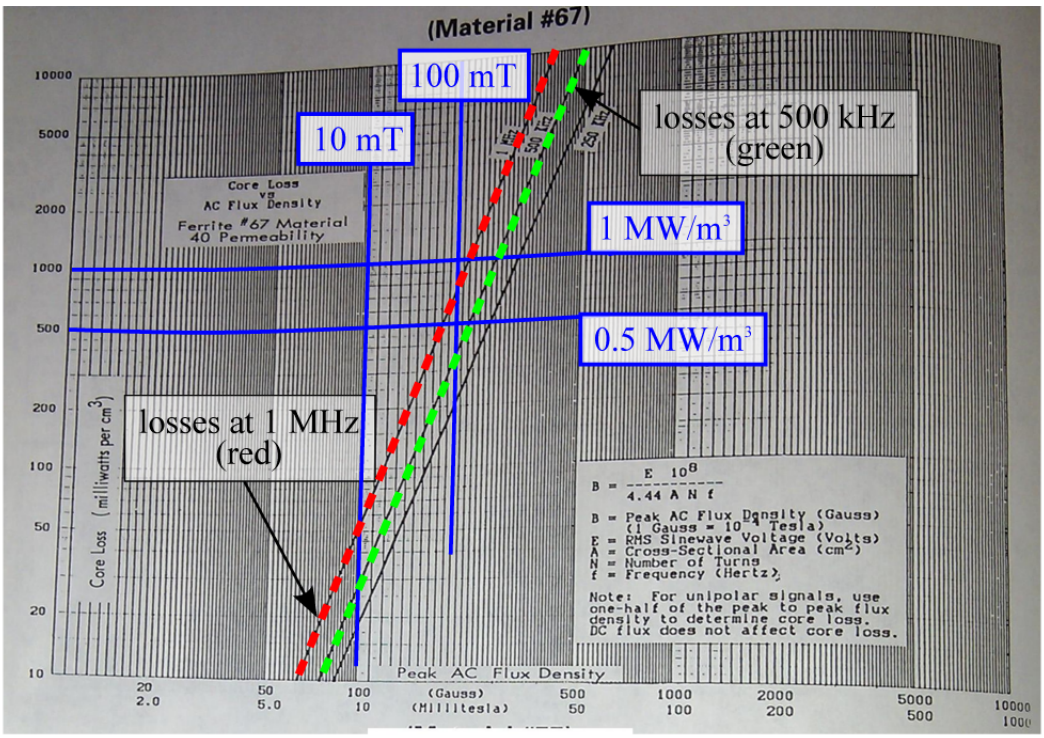

Figure 93: Vendor-Provided Loss Data for Qualification Sample

Table 8: Qualification Sample Loss Parameters (P_0 = 1 W/m^3, B_0 = 1 T)

Frequency	$\log \alpha$	β
250 kHz	11.18254	3.457265
500 kHz	11.71969	3.632442
1 MHz	12.30456	3.785981

5.5.2.3 Measured Losses

The qualification sample's losses were characterized with the balanced-reactivity loss measurement system. The results are shown in Figure 94, along with expected losses, calculated with Equation 107, at the nearest frequency and estimated operating flux.

The extracted core losses agree reasonably well with the manufacturer-reported loss behavior. The greatest departure between the manufacturer-reported (*i.e.* expected) losses and the measured losses occurs at low power dissipation levels. This can be attributed to the difficulty of accurately characterizing winding losses, which, at low excitation levels, clearly outweigh the core losses.

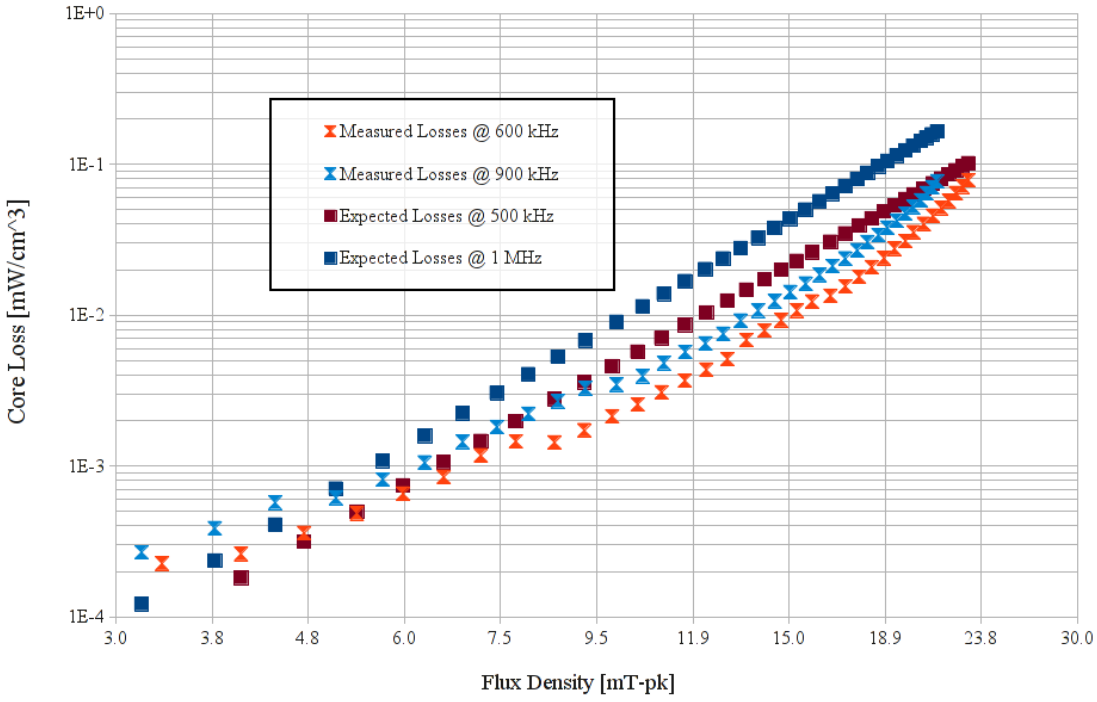

Figure 94: Qualification Sample Losses, Measured and Expected

Table 9: Samples Selected for Core Loss Characterization

Sample	Turns	Per-Layer Plating Time	Layer Thickness	Layer Count
AC14	14	no core		
SFH6A	14	1 minute	0.17 μm	199
SFH5B	14	2 minutes	0.33 μm	99
SFH4B	14	6 minutes	1.00 μm	66 (2 stacked cores)
SHT3A	14	8 minutes	1.33 μm	25
SHT2B	14	10 minutes	1.66 μm	20

5.5.3 Selected Samples

The laminated-core samples that were characterized for loss behavior are listed in Table 9. All samples had identical lateral geometry, as well as the same *nominal* magnetic material volume. In particular, the product of layer count × plating time was held invariant across all samples, producing a total nominal magnetic layer thickness of 33 μm . Also, all samples were packaged in the same bobbin structure, as shown in Figure 78, through which 14 turns of 7 × 46 AWG litz wire were threaded, forming a 14-turn inductor.

Choosing a uniform lateral geometry, total volume, and winding geometry allowed for ideally straightforward comparison of behavior between cores. Nonidealities in the cores' fabrication that would be dependent on layer thickness or per-layer plating time are not controlled for by this device-level geometric invariance. But, this does not detract from the validity of simple comparison between cores, since the essential difference between the cores is, in the first place, in their fabrication.

5.5.4 Winding Loss Extraction

As was shown with the qualification sample testing, accurate compensation for winding losses is absolutely crucial to accurate loss characterization. As a result, a substantial amount of attention was devoted to modeling and understanding the winding loss in these inductors, and the associated mechanisms.

Winding losses were modeled based on the air core inductor from the set of characterized samples. In particular, the winding on this air core inductor had precisely the same geometry as the windings on all other inductors. Without any core present, it was assumed that all losses in this device could be fully attributed to winding losses. Further, it was assumed that windings in samples with cores would exhibit similar losses. Therefore, the approach chosen for estimating winding losses was to first obtain an accurate model of the air core inductor losses, assuming dependencies only on frequency and winding current, and then apply this model in the analysis of all other inductors.

5.5.4.1 Measured Losses

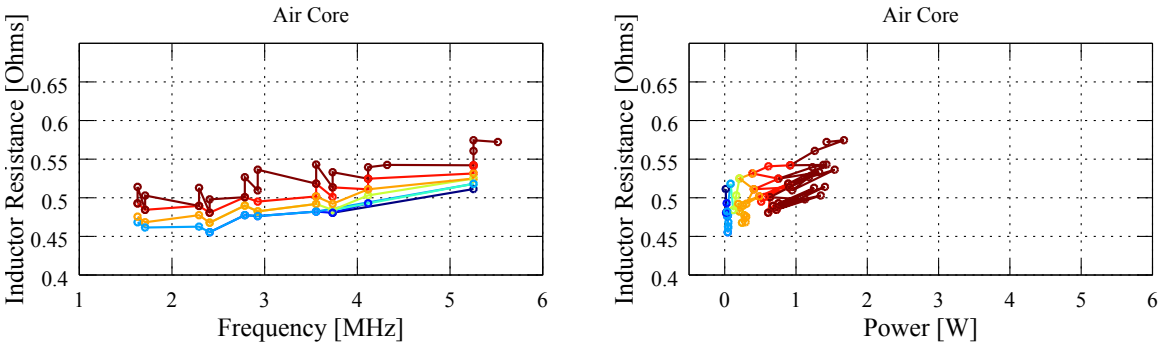

Figure 95: Air Core Inductor Losses

The total power loss for the air core inductor was, at each point, converted to an equivalent resistance value. These resistance values, presumably equivalent to resistive losses in the winding conductor plus any secondary parasitics, such as eddy currents induced in nearby metal structures, is shown in Figure 95. The winding resistance is seen to increase slightly with frequency, as might be expected due to a very weak skin effect or environmental parasitics. However, the winding resistance is also seen to increase with total power dissipation, perhaps suggesting that self-heating of the windings may be non-negligible.

5.5.4.2 Extracted Relationships

The relationship between winding resistance and power dissipation is, as can be seen, roughly linear. However, as the flux-dependent coloring helps disclose, the air core winding resistance shows bilinear behavior with respect to both power dissipation and frequency. The linear coefficient of winding resistance with respect to frequency was found through inspection to be approximately 18.22 $\mu\Omega$ per MHz. The intercept and linear coefficient with respect to power were found through linear regression to be, respectively, 416.7 $\mu\Omega$ and 36.83 $\mu\Omega$ per Watt.

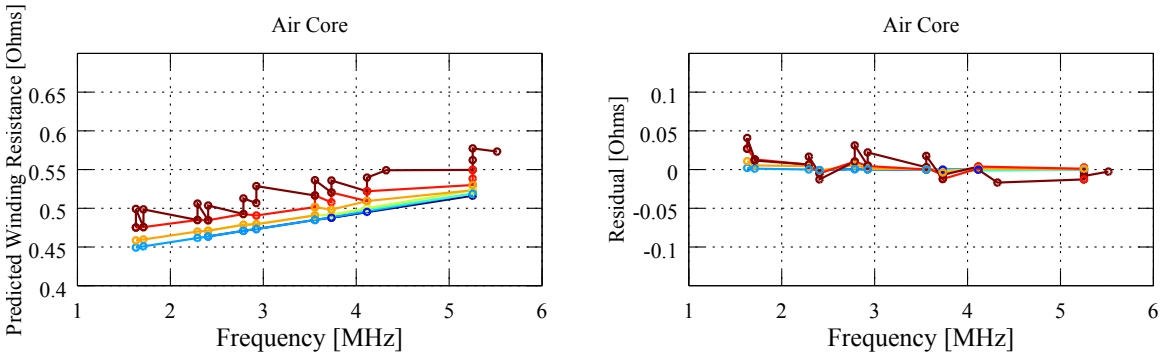

Figure 96: Air Core Inductor Losses

Figure 96 shows the agreement between the observed values for winding resistance and the extracted formula

$$R_w = 416.7\text{m}\Omega + 36.83\frac{\text{m}\Omega}{\text{W}}P + 18.22\frac{\text{m}\Omega}{\text{MHz}}f \tag{108}$$

where R_w is the winding resistance, P is the power dissipated in the inductor, and f is the frequency.

5.5.4.3 Discussion

Skin Effect

In general, the substantial increase in winding resistance is unexpected. The conductor used, 46 AWG Litz wire, has a diameter of 20 μm . This diameter is roughly equal to the skin depth of copper at 10 MHz, while most observations were taken well under 10 MHz. Thus, it can be concluded that skin effect-induced current crowding cannot be taken as a plausible explanation for this increase in resistance with frequency.

Hot Winding Effects

One of the most straightforward explanations for the observed increase in resistance with power is the increase in copper's intrinsic resistivity with temperature. The power dissipated by the device creates a temperature rise throughout the entire structure, including the copper windings. Copper's resistivity increases linearly with temperature:

$$\rho(T) = \rho_0 [1 + \alpha (T - T_0)] \tag{109}$$

where ρ is resistivity, T is temperature, α is the thermal coefficient of resistance (3800 ppm per C for copper), and ρ_0 is the resistivity at T_0.

The apparent increase in winding temperature can be derived by inverting Equation 109:

$$\Delta T = T - T_0 = \frac{1}{\alpha}\left(1 - \frac{\rho(T)}{\rho_0}\right) \tag{110}$$

$$= \frac{1}{\alpha}\left(1 - \frac{R(T)}{R_0}\right) \tag{111}$$

At this point, Equation 108 can be introduced, along with the assumption of $R_0 = R_w$ at zero power dissipation, to produce

$$R_w = 416.7\text{m}\Omega + 36.83\frac{\text{m}\Omega}{\text{W}}P + 18.22\frac{\text{m}\Omega}{\text{MHz}}f \tag{112}$$

$$\frac{R_w}{R_0} = 1 + 0.0884\text{W}^{-1}P + 0.0434\text{MHz}^{-1}f \tag{113}$$

$$\tag{114}$$

and

$$\Delta T = \frac{1}{\alpha}\left(1 - \frac{R_w}{R_0}\right) \tag{115}$$

$$= \frac{1}{\alpha}\left(0.0884\mathrm{W}^{-1}P + 0.0434\mathrm{MHz}^{-1}f\right) \tag{116}$$

$$= \frac{1}{3800E - 6\mathrm{C}^{-1}}\left(0.0884\mathrm{W}^{-1}P + 0.0434\mathrm{MHz}^{-1}f\right) \tag{117}$$

$$= 23.2\frac{\mathrm{C}}{\mathrm{W}}P + 11.5\frac{\mathrm{C}}{\mathrm{MHz}}f \tag{118}$$

This implies a temperature rise in excess of 150 C for the maximum produced power dissipation in the air core device. Given the vigorous flow of air onto the sample during testing (shown in Figure 90) this would entail a conductor temperature no lower than 175 C. However, it is unlikely that the winding insulation would be able to survive at this high temperature.

This implies that the change in winding resistance is not fully attributable to temperature rise in the copper conductors. Accordingly, it is likely there are additional loss mechanisms present in the air core device. These might include eddy currents in surrounding structures as well as dissipative losses in the bobbin's polymers.

Other Effects

The thermally-referred increase in resistance with frequency – 11.5 C per MHz – in Equation 115 is more difficult to interpret as a strictly thermal phenomenon. The equivalent increase in resistance – 18.22 $\mu\Omega$ per MHz – may be more easily explained by parasitic loss mechanisms in structures surrounding the device under test. In particular, the HFHF circuit board contains large expanses of copper as well as lossy dielectric ("FR-4") that would produce noticeable losses at high frequencies.

Applicability

Using this analysis to predict winding losses in inductors with cores requires the acceptance of two equating assumptions. First, it must be assumed that the effective temperature rise and its effects will be the same in both the air-core and with-core configurations. In the air-core configuration, most power dissipation is likely occurring in the windings themselves, while in the with-core configuration, most power dissipation is happening within the core. The first-order effects of this difference act in opposing senses and, qualitatively, neutralize each other.

Second, it must be assumed that all changes in winding resistance are due to only temperature rise of the copper conductors and isolated parasitic losses. Since "coil-cutting" field effects are known to be a major source of losses in cored inductors, this assumption is not clearly valid. However, it is not clearly invalid either, as heat-related failures (*i.e.* melted insulation) have been readily observed in windings.

5.5.5 Laminated-Core Characterization Results

5.5.5.1 Basic Results

The observed quantities and basic extracted quantities are presented in this section. To best facilitate the separation of observed core losses into multiple contributions from different loss mechanisms, all cores should be tested at some set of standard frequencies and some set of standard flux levels. In this case, at each combination of frequency and flux level, it would be possible to directly compare core loss as a function purely of lamination thickness, while effectively keeping frequency and flux level constant. Such a well-regimented data set would facilitate both qualitative and quantitative analysis of the data.

However, this was not possible to achieve in practice for a number of reasons. First, and most importantly, the inductance of the samples varied significantly across lamination thicknesses, even though the nominal inductance values were invariant. For a particular inductor sample, there was a one-to-one mapping between capacitance banks and frequencies at which resonance could be reached, and thus at which core loss measurement could take place. Given that there was a finite, relatively small set of capacitance banks available – less than 10 possible capacitance values – this meant that each inductor would, generally, get measured at a unique set of frequencies. Therefore, it was generally not possible to measure two different inductors at the same frequency.

As well, it was not possible to maintain a constant core flux level from sample to sample. This was due entirely to the practical difficulties found both in estimating core flux levels in real time and in implementing a closed-loop algorithm for tuning the inductor current to achieve a targeted flux level.

As a solution to this issue, core loss was measured at many different, but regularly-spaced, drive levels. Each drive level corresponded to a unique flux density, although the flux density was

calculated later, as a dependent variable. This approach provided for the collection of a very large, dense set of data, which would allow for accurate execution of statistical curve fitting algorithms such as least squares regression, described in Appendix B. For analyses in which constant core flux was needed, the full data set (for each inductor) was separated into bins according to flux level. The density of the data ensured that these bins would, with few exceptions, contain enough data points to validate any analyses or curve fitting carried out only on its contents (*ex.* Figure 101).

Figure 97 shows these dense data sets for each inductor, except the air core device AC14, summarizing the ranges of frequency and core flux over which each inductor was characterized. The density of the collected data is clearly represented, but Figure 97 also shows the division of each inductor's data into bins according to flux level.

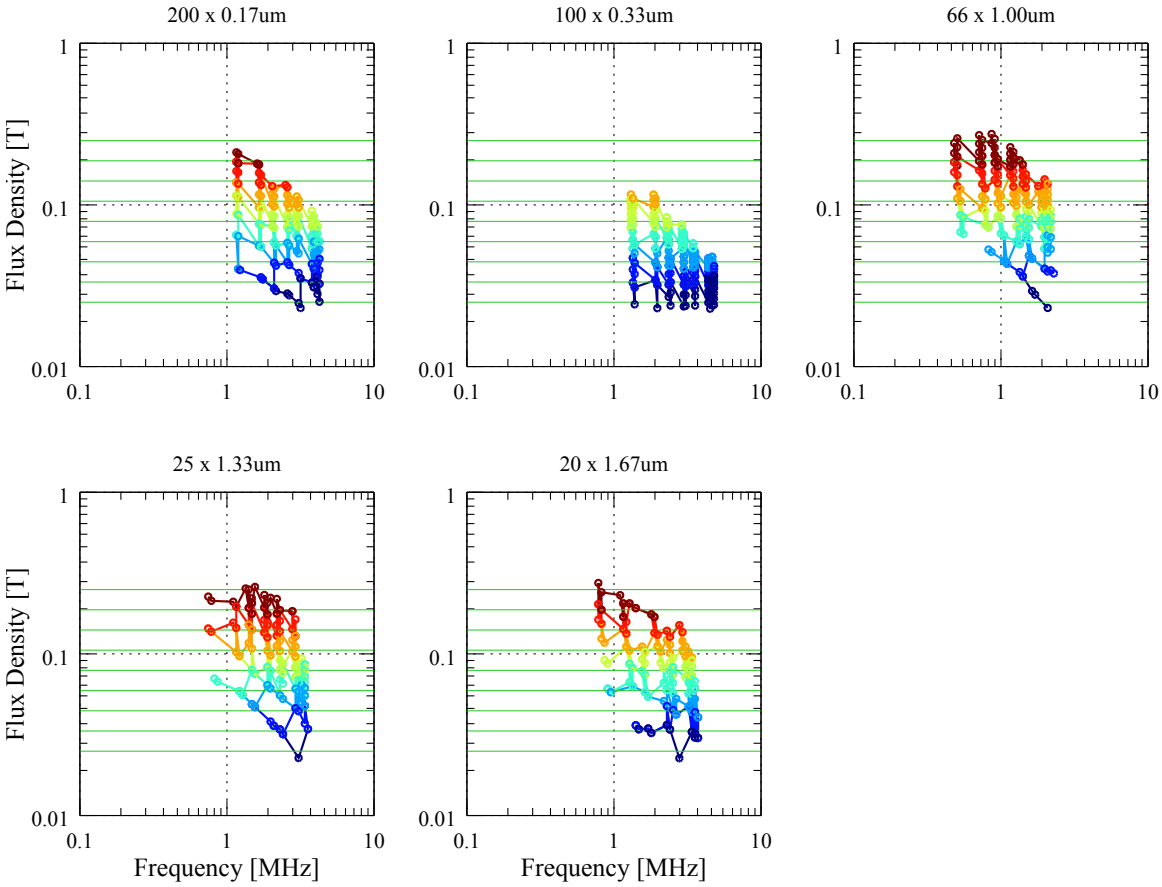

Figure 97: Flux density and frequency coverage for core loss measurement. This plot introduces the color convention used in the proceeding plots, and may be used as a key to interpret the coloration in Figure 99 through Figure 101.

In Figure 97, the convention is introduced of colorizing data points according to the associated

magnetic flux density within the core. The color blue corresponds to the lowest flux levels, red corresponds to the highest core levels, and intervening flux levels are assigned colors that intervene along the standard hue axis. The green horizontal lines demarcate the boundaries between adjacent flux density coloration levels.

This colorization based on core flux is redundant in Figure 97, since the basis for colorization is the same as the y axis. However, exactly the same colorization scheme is maintained in all subsequent plots up to and including Figure 101. Thus, Figure 97 may be referred back to as a key for interpreting the colorization in those plots. These aspects of this set of plots are illustrated in Figure 98.

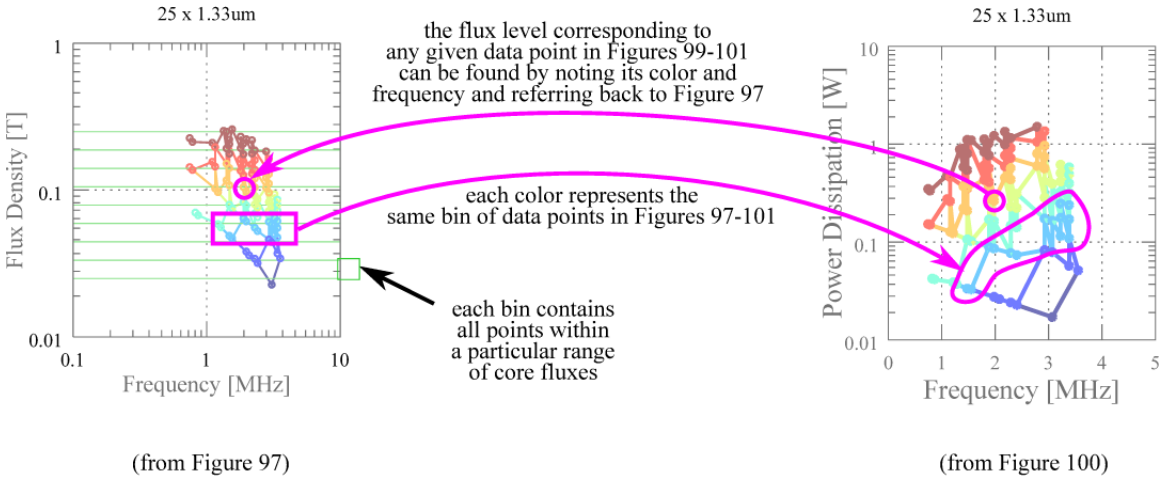

Figure 98: Colorization and Binning of Core Loss Data Based on Core Flux Level. In Figure 97, the division of all datapoints into bins according to core flux level, and colorization of the datapoints by the same rule, is introduced. This binning and colorization is maintained through Figure 101. This allows the inference of relative core flux levels according to color, as well as a way to explicitly check the core flux level corresponding to a particular data point.

Figure 99 shows the inductance of all characterized samples. The dependence of inductance on frequency is shown explicitly, while the coloration of the datapoints illustrates the dependence of inductance on flux density.

Not surprisingly, the air-core inductor AC14 does not show any change in inductance with either frequency or flux density. In fact, the presence of any such dependence would suggest that there were substantial parasitic effects in the HFHF measurement system. The remaining inductors – those with magnetic core material present – show a consistent increase in inductance with increasing flux density, as the datapoint color transition from blue to red.

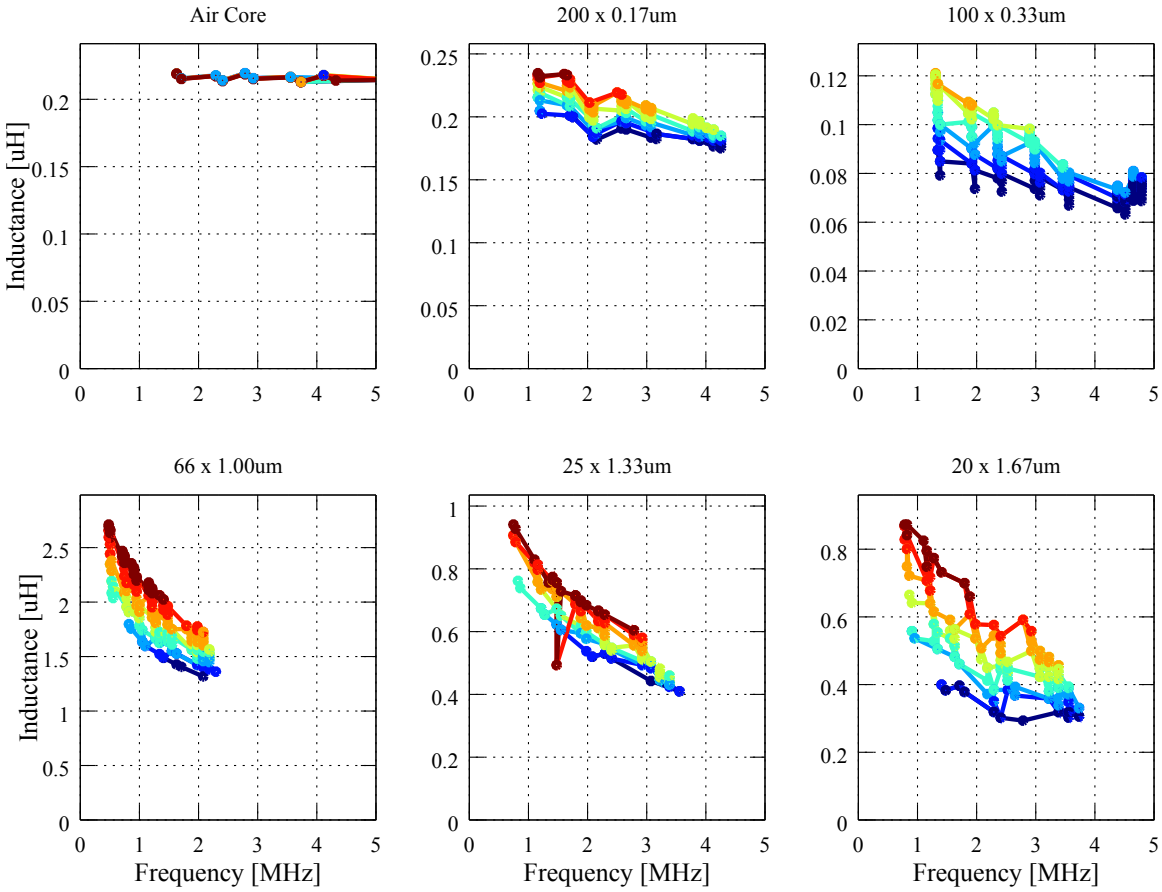

Figure 99: Inductance per sample, colorized by core flux density. The coloration of the data points is based on the core flux binning illustrated in Figure 97.

Figure 100 shows the total power loss within the inductor for each sample, again illustrating dependencies on frequency and on flux density. As expected, power losses increase with frequency and with core flux level.

However, the data in Figure 100 is total power loss in the inductor, P_{tot}. The quantity of interest for this analysis is core loss, denoted simply as P in the remainder of this chapter. This can be calculated from total inductor power loss simply by subtracting the winding losses P_w.

$$P_{tot} = P + P_w \tag{119}$$

$$P = P_{tot} - P_w \tag{120}$$

Equation 108 was used to estimate the winding resistance R_w for each datapoint, based on the frequency f and total device dissipation P_{tot} corresponding to that datapoint. This yields

$$P = P_{tot} - I^2 R_w\,(f, P_{tot}) \tag{121}$$

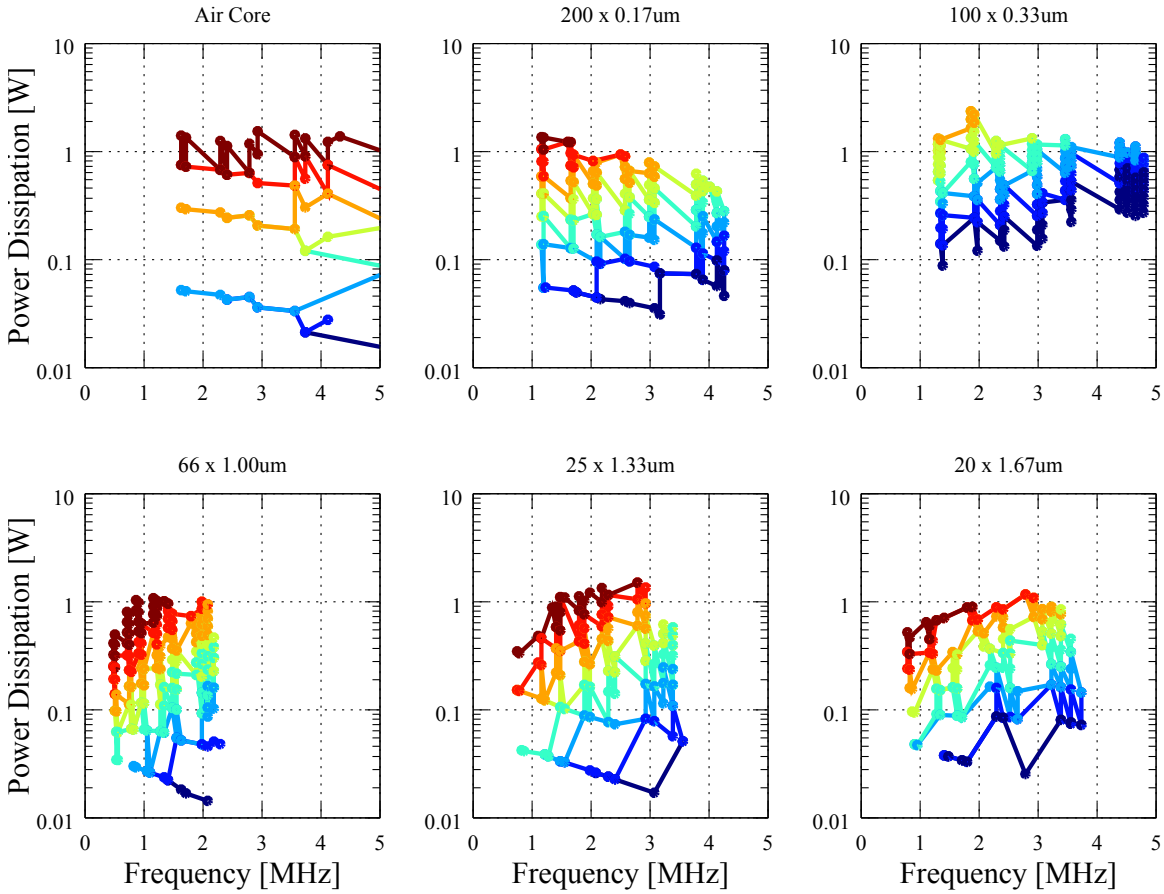

Figure 100: Measured Power Loss for Laminated-Core Inductors The coloration of the data points is based on the core flux binning illustrated in Figure 97.

where I is the inductor's measured RMS current and $R_w(f, P_{tot})$ is evaluated according to Equation 108.

Precisely evaluating the frequency- and flux-dependent behavior the core losses P across varied lamination thickensses calls for appropriate analytical framing. This is presented in the proceeding sections.

5.5.5.2 *First-Order Analysis*

The initial analysis of the core loss measurements consists of graphically characterizing the loss's dependence on frequency. Standard core loss models call for hysteresis power loss P_h that is proportional to frequency, core volume, and the square of the flux density, and eddy current power loss P_e that is proportional to core volume and to the squares of frequency and flux density. This

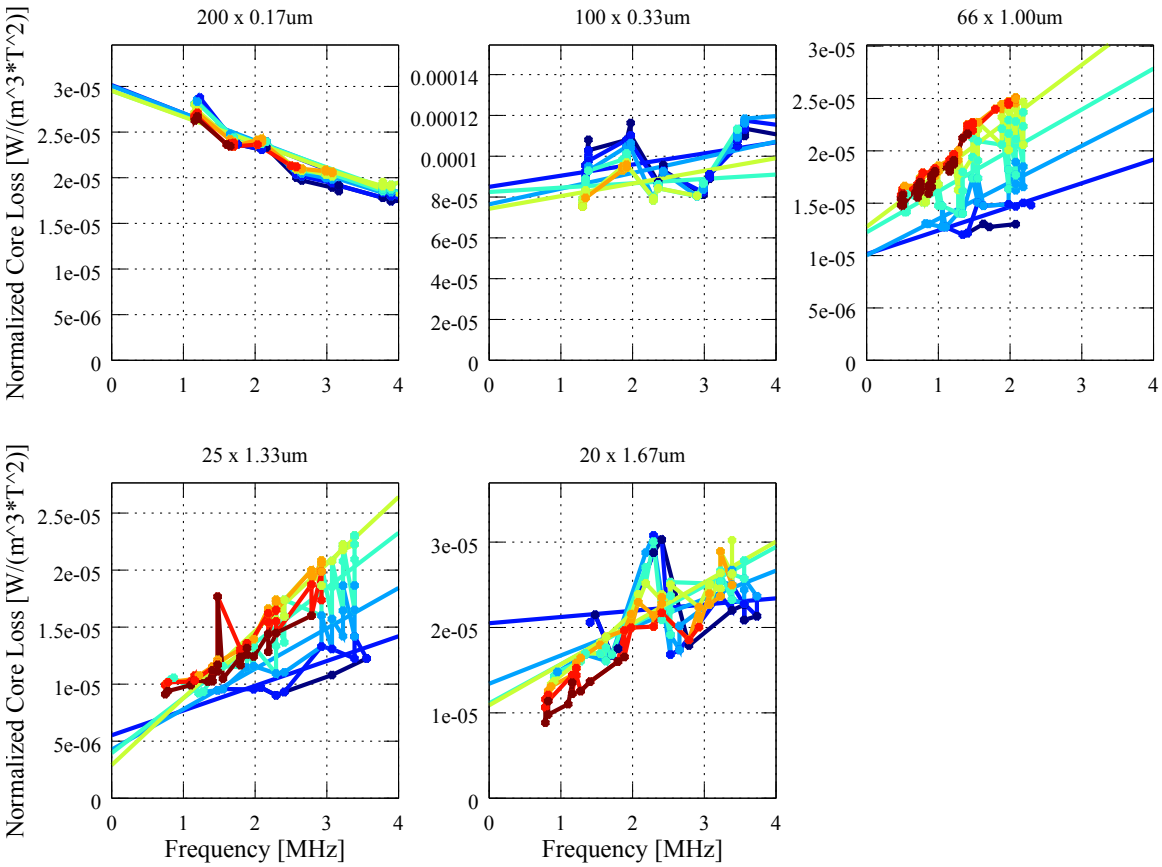

Figure 101: Normalized Core Loss per Cycle, with Simple Regression Lines. The coloration of the data points, as well as the simple regression lines, is based on the core flux binning illustrated in Figure 97. Each regression line was generated from the datapoints within only one core flux bin, as indicated by like coloration.

relationship is captured symbolically in

$$P = P_h + P_e \tag{122}$$

$$= K_h \left(VB^2 f\right) + K_e \left(VB^2 f^2\right) \tag{123}$$

$$= VB^2 f \left(K_h + f K_e\right) \tag{124}$$

$$\frac{P}{VB^2 f} = K_h + f K_e \tag{125}$$

where P is the total core power, V is the core volume, f is frequency, and B is flux density. Equation 125 can be rewritten in volumetric terms as

$$P_V = P_{Vh} + fP_{Ve} \tag{126}$$

$$= B^2 f\left(K_h + fK_e\right) \tag{127}$$

$$\frac{P_V}{B^2 f} = K_{Vh} + fK_{Ve} \tag{128}$$

where P_V, P_{Vh}, and P_{Ve} are the quotients of P, P_h, and P_e divided by core volume V.

Equation 128 describes a simple linear relation between $P_V/\left(B^2 f\right)$ and f, where K_h is the y-axis intercept and K_e is the slope. Thus, plotting $P_V/\left(B^2 f\right)$ as a function of frequency should produce a highly linear curve, assuming that K_h and K_e behave as constants. A core with zero eddy current losses, as an example, would produce a K_e of zero, and therefore a horizontal line for its plot of $P_V/\left(B^2 f\right)$ as a function of frequency.

Figure 101 shows these plots for each inductor that was characterized. As was expected, the slope of the plot increases with increasing lamination thickness. The slightly negative slope seen in the 166 nm sample (SFH6A) plot is surprising, and may be attributed to over-compensation for winding losses.

5.5.5.3 Advanced Analysis

The analysis in subsubsection 5.5.5.2 is straightforward and intuitive. However, it is clear from Figure 101 that it does not fit the data collected in an ideal fashion. As well, it does not account for so-called anomalous losses, also known as excess losses, which canonically scale as both frequency and flux density taken to the exponent of 1.5. Analytically, this is expressed as

$$P_{Va} = \frac{P_a}{V} \tag{129}$$

$$= \frac{1}{V} K_a f^{1.5} B^{1.5} \tag{130}$$

where P_a is the anomalous power loss throughout the entire core, P_{Va} is the volumetric power loss, V is the core volume, and f and B are frequency and core flux density, respectively.

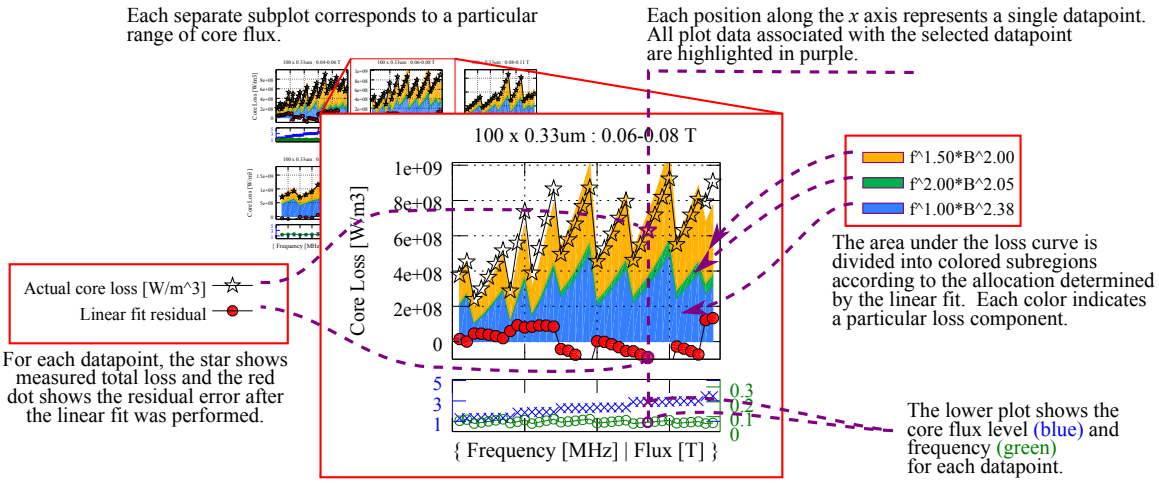

Figure 102: Explanation of Fit Plot Format. Each discrete location along the x axis corresponds to a single measured datapoint. Data plotted at the same x location on the plots provides complete information about the results of the linear fit, as it pertains to the given datapoint.

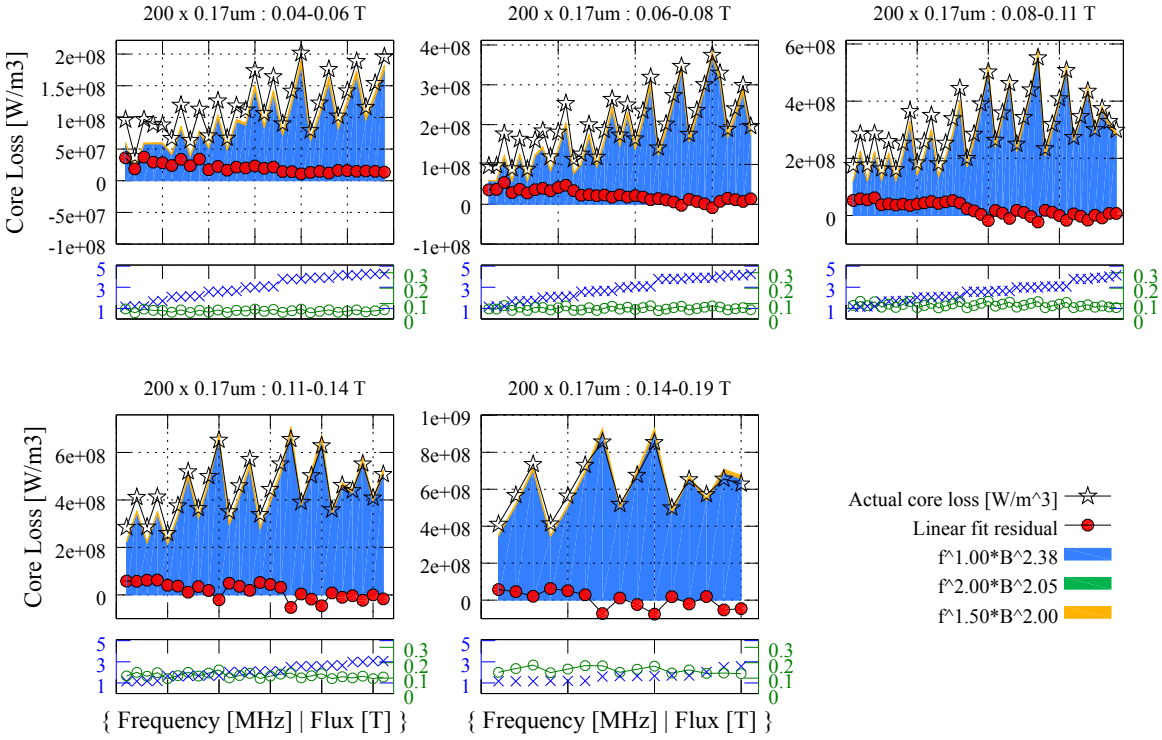

Figure 103: Least-squares decomposition of sample SFH6A (166 nm laminations). $f^{1.00} * B^{2.38}$ is the term associated with hysteresis losses, $f^{2.00} * B^{2.05}$ is the term associated with eddy current losses, and $f^{1.50} * B^{2.00}$ is the term associated with anomalous losses

Generalizing the core loss expression to include anomalous losses gives

$$P_V = P_{Vh} + P_{Ve} + P_{Va} \tag{131}$$

$$= K_h B^2 f + K_e B^2 f^2 + K_a B^{1.5} f^{1.5} \tag{132}$$

$$= \begin{bmatrix} K_h \\ K_e \\ K_a \end{bmatrix} \cdot \begin{bmatrix} B^2 f \\ B^2 f^2 \\ B^{1.5} f^{1.5} \end{bmatrix} \tag{133}$$

This expression describes a rich linear decomposition for core losses. The algorithm of least-squares fitting was used to determine the linear coefficients K_h, K_e, and K_a for each of the six samples characterized. This produced good fits between the measured core losses and the linear decomposition.

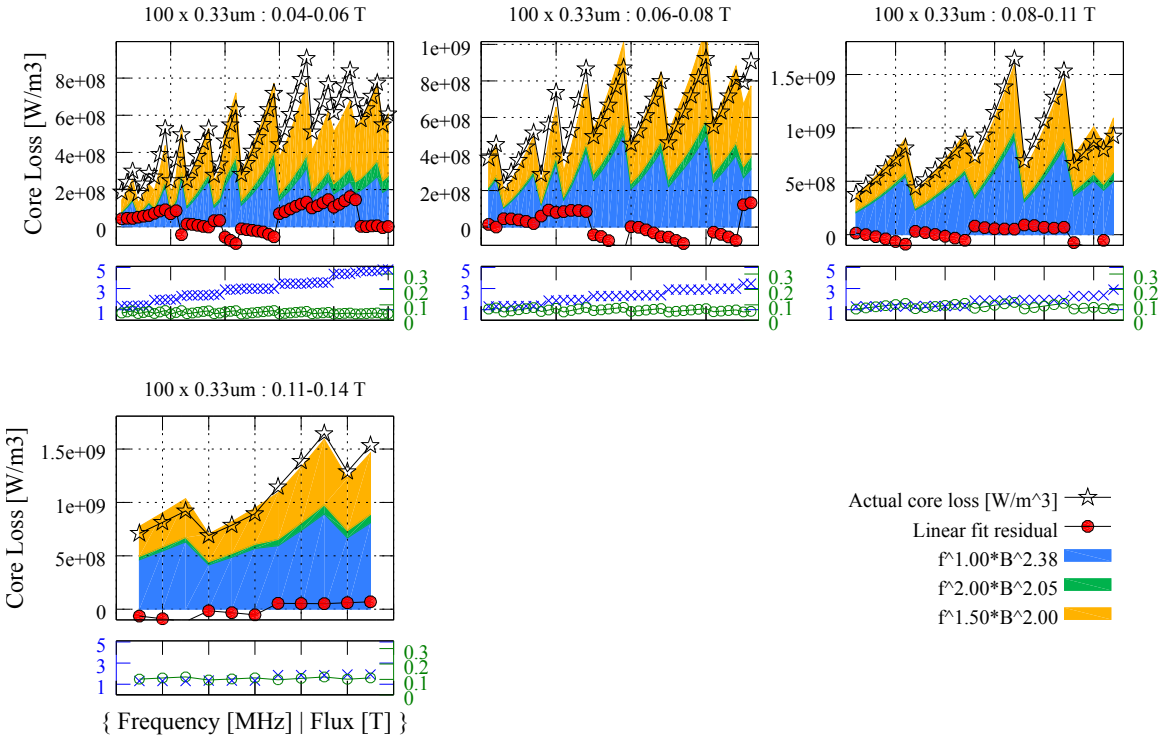

Figure 104: Least-squares decomposition of sample SFH5B (333 nm laminations). $f^{1.00} * B^{2.38}$ is the term associated with hysteresis losses, $f^{2.00} * B^{2.05}$ is the term associated with eddy current losses, and $f^{1.50} * B^{2.00}$ is the term associated with anomalous losses

However, it was found that minor variations in the exponents to which f and B are raised in Equation 133 could produce even better fits. This is expressed mathematically by replacing each

exponent with a parameter β_i.

$$P_V(\beta_1 ... \beta_6) = \begin{bmatrix} K_h \\ K_e \\ K_a \end{bmatrix} \cdot \begin{bmatrix} B^{\beta_1} f^{\beta_2} \\ B^{\beta_3} f^{\beta_4} \\ B^{\beta_5} f^{\beta_6} \end{bmatrix} \tag{134}$$

Each unique set of β_i would dictate a corresponding set of linear-fit loss coefficients K_h, K_e, and K_a, as well as an associated residual r. In this way, a function-like mapping from β_i to linear residual r could be defined.

A gradient-descent non-linear least squares algorithm was used to minimize the residual r with respect to the β_i. In this way, the optimal values for the loss coefficients could be discovered. The optimized loss decomposition is expressed as

$$P_V = \begin{bmatrix} K_h \\ K_e \\ K_a \end{bmatrix} \cdot \begin{bmatrix} B^{2.38} f \\ B^{2.05} f^2 \\ B^{1.5} f^{1.5} \end{bmatrix} \tag{135}$$

This formulation was applied to the loss data measurements fo each inductor and, in all cases, produced a fit with an R^2 value of 0.94 or higher.

The fit of each sample's measured core losses into this decomposition is shown in Figure 103 through Figure 107. The format used for these plots is somewhat unique and requires some explanation. Figure 102 outlines the general structure and interpretation of these plots. These plots, one for each inductor, show all datapoints gathered for the device. It was desired to simultaneously depict the entirety of the data taken, the quality of the linear fit, and, most importantly, the results of the fit.

Each of these figures consists of a number of subplots, with each subplot consisting of two sets of axes, vertically stacked. In these figures, each flux range, as defined in Figure 97, is plotted in a separate subplot with the corresponding flux range shown in the title above the upper set of axes. Within each subplot, the upper set of axes shows the measured and extracted core losses. The directly measured total core losses are plotted as black stars, meaning that the distance between the black star and the x axis (*i.e.* $y = 0$) is proportional to core loss at that point. This distance, at each point, is then divided according to the allocation of core losses, as determined by Equation 135, and the space between the black star and the x axis is colorized according to the terms in Equation 135.

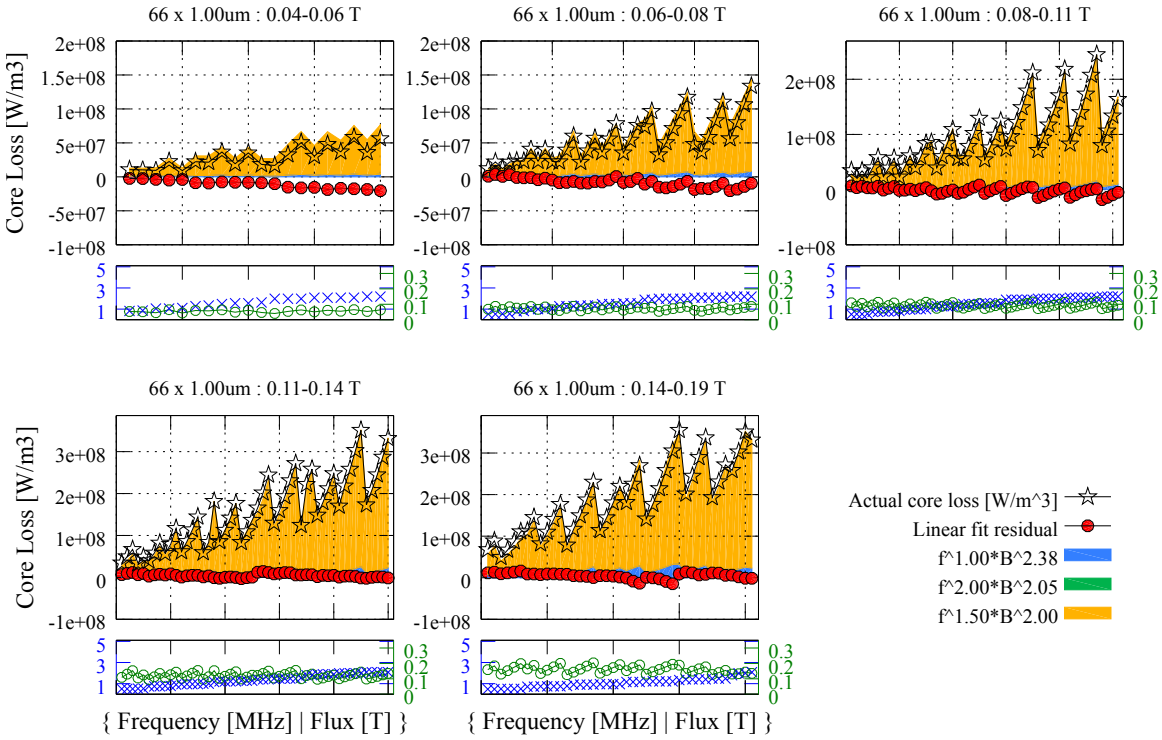

Figure 105: Least-squares decomposition of sample SFH4B (1 μm laminations). $f^{1.00} * B^{2.38}$ is the term associated with hysteresis losses, $f^{2.00} * B^{2.05}$ is the term associated with eddy current losses, and $f^{1.50} * B^{2.00}$ is the term associated with anomalous losses

Across a given set of points, this division and colorization creates a plot where area is proportional to power dissipation, and colorization of the area under the curve visually indicates the discovered allocation of core losses. Finally, the residual, or difference between measured total losses and losses predicted by Equation 135, is depicted at each measurement point as a red dot.

The x axis for each subplot serves only to serialize the set of datapoints in that subplot; it corresponds to a generally arbitrary order of iteration through the data. However, since each point on the x axis corresponds to a single datapoint, each point on the x axis must have an associated frequency and core flux level. The lower set of axes in each subplot shows precisely these quantities – frequency in blue, ruled by the left hand axis and core flux in green, ruled by the right hand axis. Figure 102 highlights all of these features of these plots.

The results of this least-squares fit, across all inductors, is shown in Figure 108. The graphic shows a clear trend of increasing eddy and anomalous current losses as the lamination thickness increases. This is generally as expected, but the magnitude of losses is unexpected. This is discussed in subsection 5.5.6.

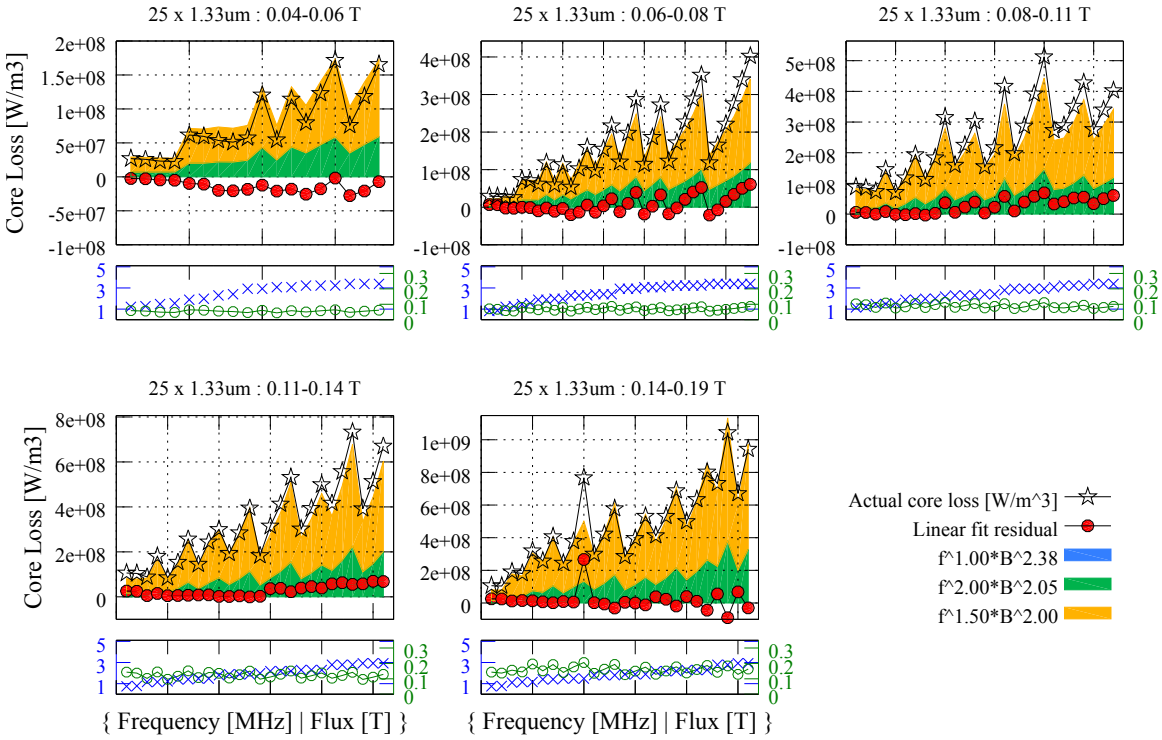

Figure 106: Least-squares decomposition of sample SHT3A (1.33 μm laminations). $f^{1.00} * B^{2.38}$ is the term associated with hysteresis losses, $f^{2.00} * B^{2.05}$ is the term associated with eddy current losses, and $f^{1.50} * B^{2.00}$ is the term associated with anomalous losses

5.5.5.4 *Comparison with Expected Losses*

While the analysis in subsubsection 5.5.5.3 provides quantitative insight into the reduction of eddy current losses, it does not establish agreement with theoretically-predicted values. Towards this end, the analysis presented in [4], as developed in subsection B.3.1, is used to connect experimental observations with analytically-calculated values. The linear core loss decomposition based on the expected loss values P_{Ve} and P_{Vh}, corresponding respectively to the volumetric eddy current and hysteresis core losses:

$$P_{Ve} = \frac{\pi}{4} P_{magV}(f, B_a)\, G_{Ve}\left(\frac{a}{\delta}\right) \tag{136}$$

$$P_{Vh} = \frac{S}{2} P_{magV}(f, B_a)\, G_{Vh}\left(\frac{a}{\delta}\right) \tag{137}$$

would be

$$P_V = \begin{bmatrix} K_e \\ K_h \end{bmatrix} \cdot \begin{bmatrix} \frac{\pi}{4} P_{magV}(f, B_a)\, G_{Ve}\left(\frac{a}{\delta}\right) \\ \frac{S}{2} P_{magV}(f, B_a)\, G_{Vh}\left(\frac{a}{\delta}\right) \end{bmatrix} \tag{138}$$

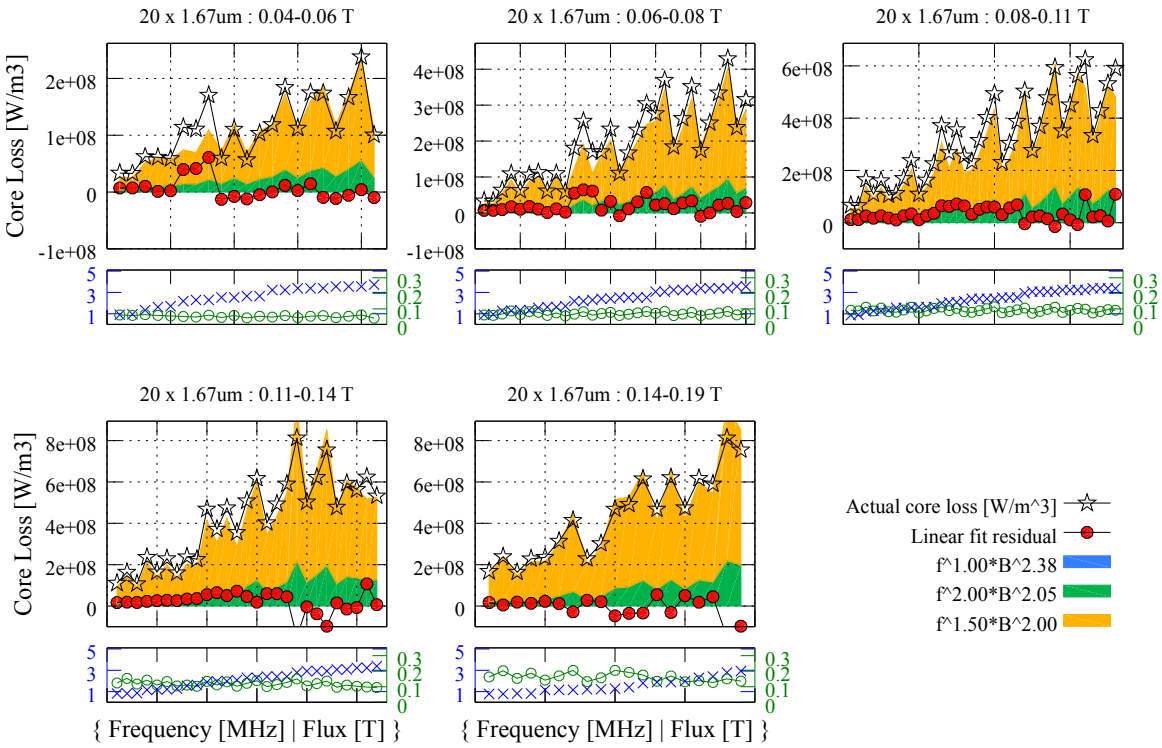

Figure 107: Least-squares decomposition of sample SHT2B (1.66 μm laminations). $f^{1.00} * B^{2.38}$ is the term associated with hysteresis losses, $f^{2.00} * B^{2.05}$ is the term associated with eddy current losses, and $f^{1.50} * B^{2.00}$ is the term associated with anomalous losses

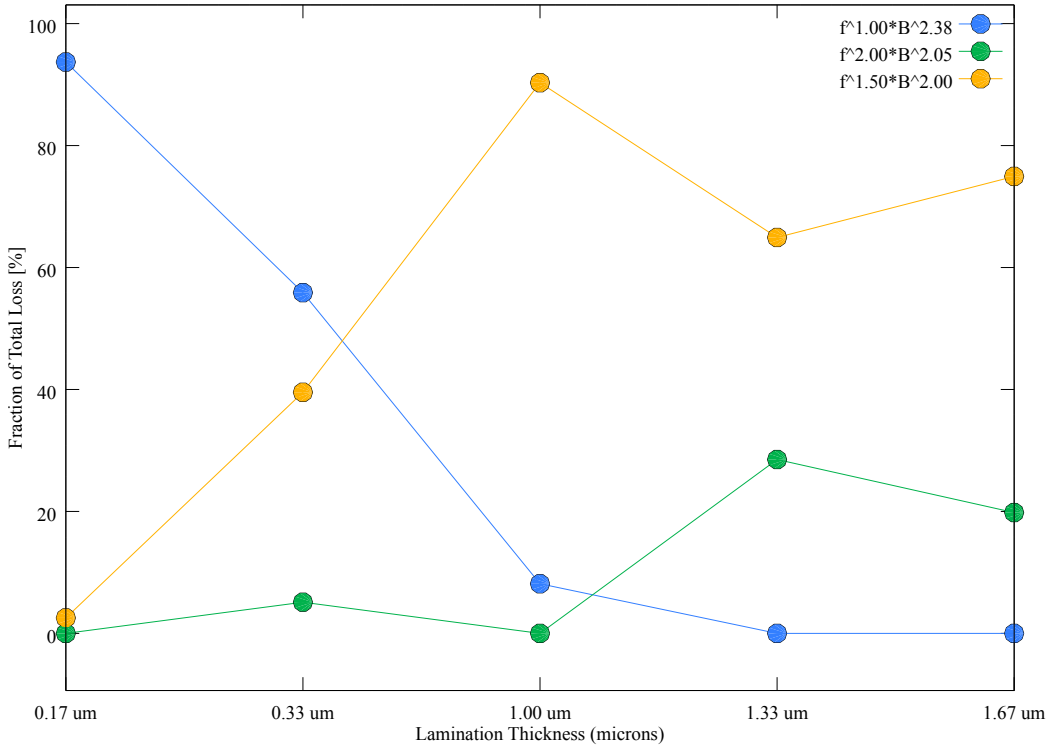

Figure 108: Summary of least-squares fit for core losses

Further, within this decomposition, the linear fit coefficients K_e and K_h also indicate the quantitative match between the observed and expected losses.

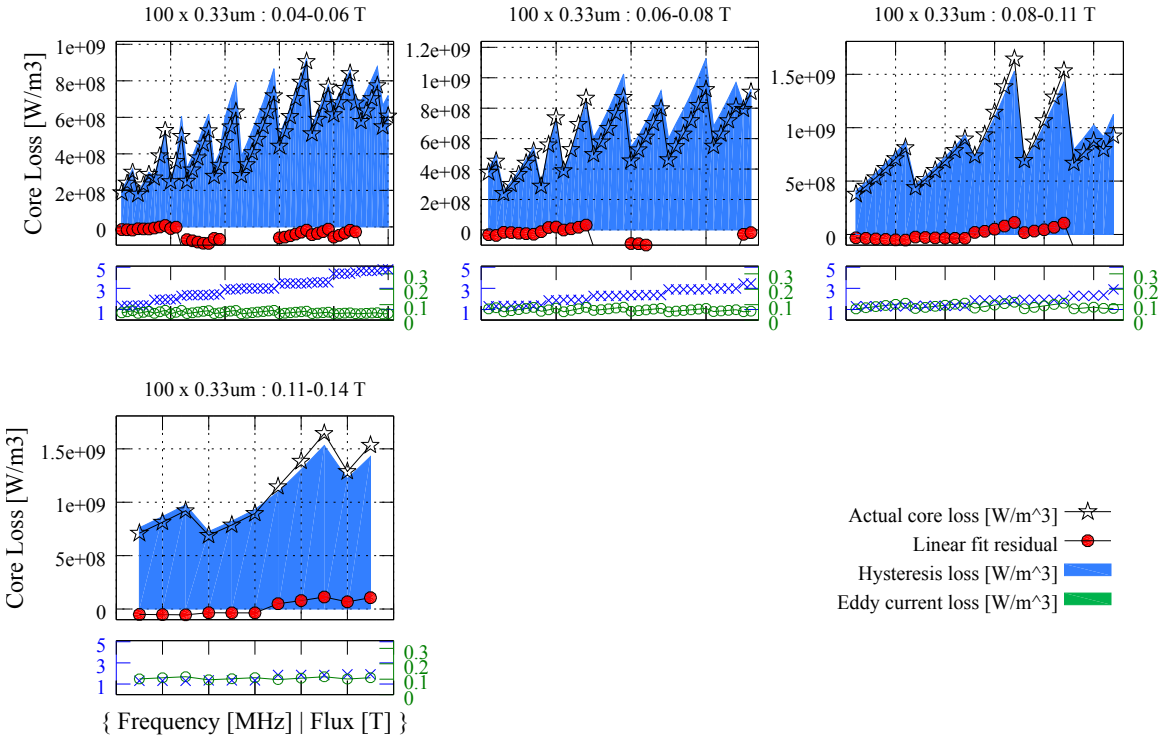

Figure 109: Least-squares decomposition of sample SHT5B (0.333 μm laminations)

Accordingly, a least-squares fit was performed for each sample onto the functions P_{Ve} and P_{Vh} above. The expected core losses were calculated using Equation 217. P_{mag} was estimated from observed quantities using Equation 248 and S was estimated from vibrating sample magnetometer (VSM) data in accordance with Figure 156. Detailed examples of this fit are shown in Figure 109 and Figure 110.

The fit results are summarized in Figure 111 and Figure 112. Figure 111 is analogous to Figure 108, in which the power loss attributed to each loss mechanism is summed across all data points for each sample, and expressed as a fraction of the total power loss (also summed across all data points for the sample). This illustration shows the increase in observed eddy current losses and the decrease in hysteresis losses as lamination thickness increases. This is, qualitatively, exactly as expected.

Although the trend may seem less dramatic than the one seen in Figure 108, it is likely more

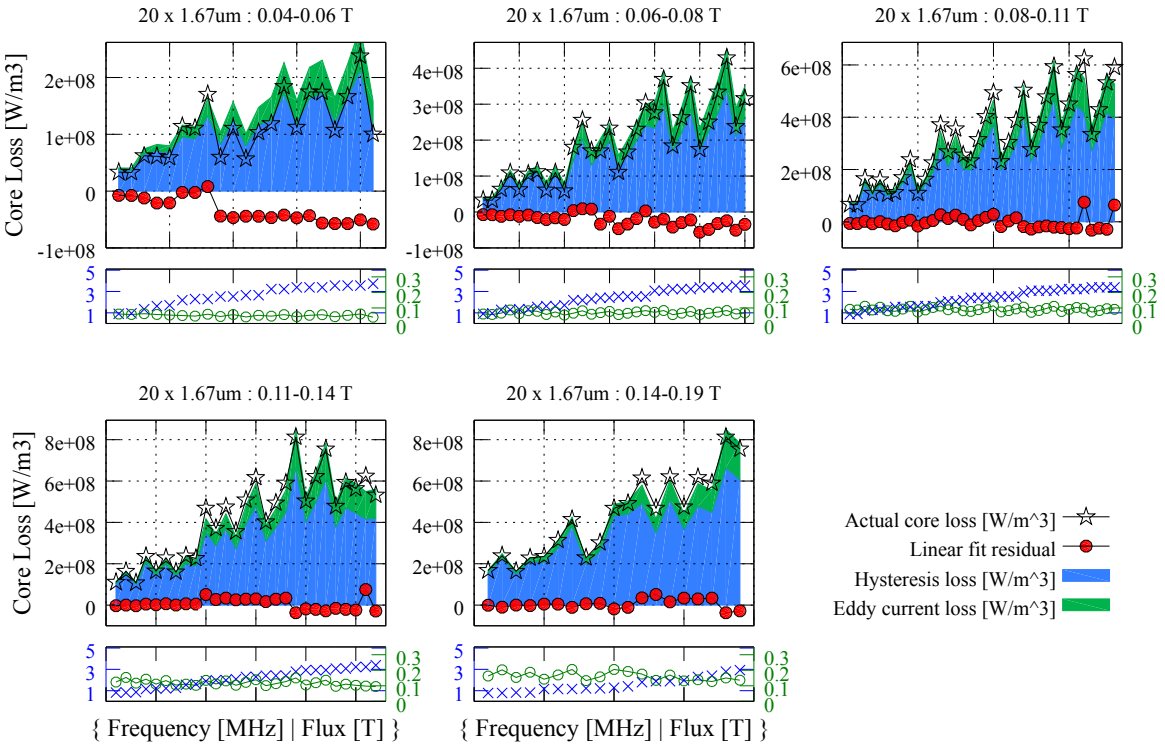

Figure 110: Least-squares decomposition of sample SHT2B (1.66 μm laminations)

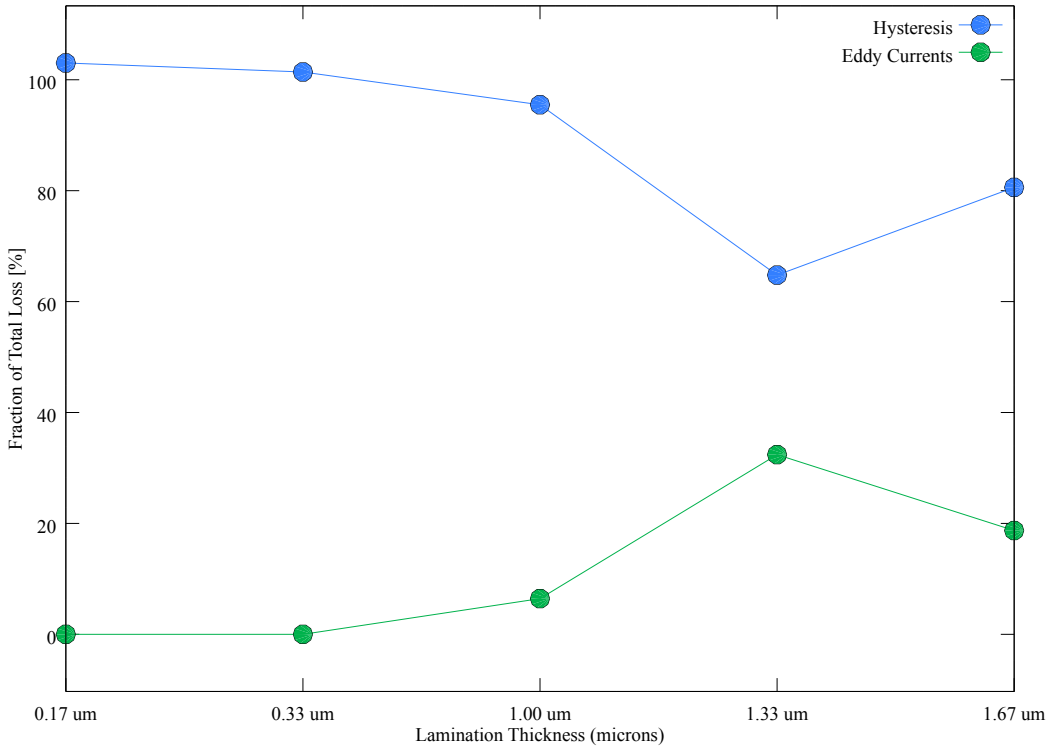

Figure 111: Allocation of measured core losses among core loss components

realistic. This is due to the improved fidelity of the analytical modeling underlying the loss decomposition, which considers hysteresis and eddy current losses within the same framework.

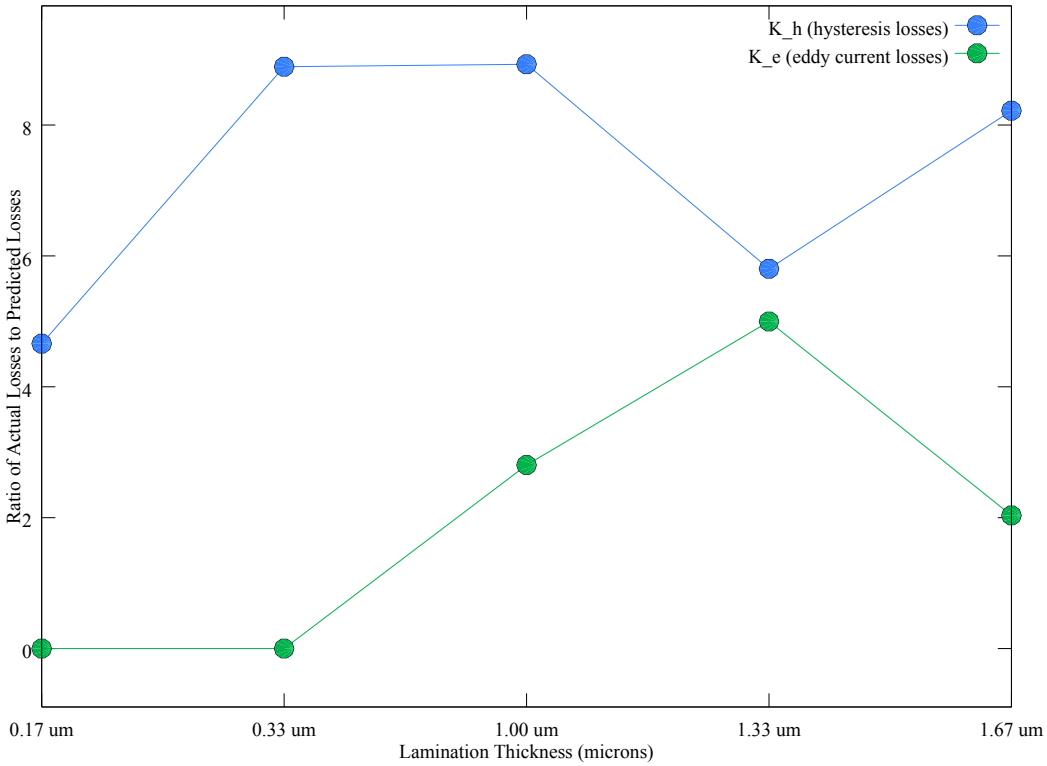

Figure 112: Ratio of extracted measured core loss to predicted core loss

Figure 112 shows, graphically, the actual coefficients K_e and K_h determined by the linear fit for each loss term, for each sample. This result indicates that the hysteresis current losses were consistently underestimated by the chosen analytical model. As well, for thicker laminations where eddy losses are more pronounced, the model also significantly underestimated the losses.

These departures, as well as the surprising aspects of Figure 108, are discussed in subsection 5.5.6.

5.5.6 Potential Sources of Discrepancy between Expected and Measured losses

Figure 112 reveals a substantial mismatch between expected and measured core losses, of both hysteresis and eddy current origins.

5.5.6.1 *Anomalous Losses*

Anomalous losses are defined as observed core losses not accounted for by eddy current or magnetization hysteresis phenomena. These are addressed as a high-level source of mismatch between measured losses and expected losses because the model used in the analysis above, from [4], does not account for anomalous losses. Therefore, any anomalous losses present would necessarily, in any loss separation methodology, alias to the presumed loss mechanisms of eddy currents and hysteresis. At the same time, the loss separation summarized in Figure 108 does account for anomalous losses.

Analytical models necessary to understand anomalous losses is very complex and requires modeling magnetization dynamics at the micro- and nano- scales. Domain-level modeling [8] [52] is a relatively simple approach to understanding anomalous behavior. Figure 113 shows the cross-section of a lamination and the induced eddy currents, as well as the associated core loss behavior, under two different modeling assumptions. The "classical eddy currents" case is the same as the analysis used in [4] and the expected core loss behavior is the same.

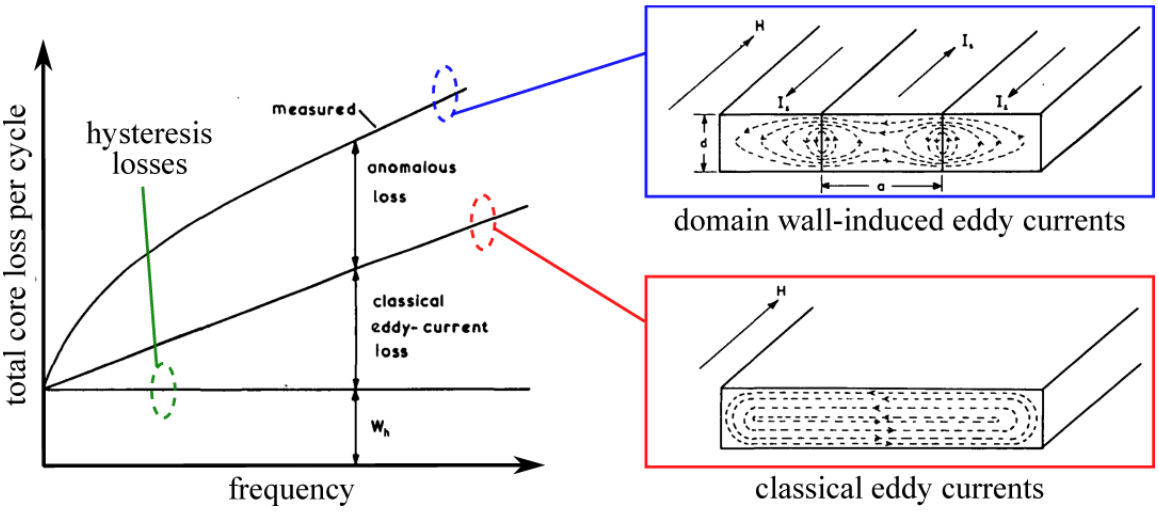

Figure 113: Core Loss Constituents, Including Anomalous Losses. [8]

The "domain wall-induced eddy currents" case depicts individual domains, presumed to have a longitudinally extruded shape as well as longitudinally directed magnetization. These domains have parallel and antiparallel magnetization directions, but for at least some portion of the magnetization trajectory, the domains do not change their magnetization direction; rather, they grow and shrink

to achieve the externally-mandated total flux. As the fully-magnetized domains shrink and grow, the material in the vicinity of the domains' walls experience a sudden, full-amplitude reversal in magnetization. This high rate of change in magnetic flux induces eddy currents which dissipate power.

The behavior of the anomalous losses over frequency and magnetization intensity will depend heavily on the geometry of the domains and their behavior under dynamic conditions. The analyses in [8] and [50] are based on well-studied domain structures in rolled silicon steel, a material with well-known behavior. The simplicity of these structures allows the development of some experimentally-verifiable models, even in the presence of modestly complicating assumptions such as nonuniform domain wall spacing ([50]).

Developing analogous models for the electrodeposited laminations studied in this work would not only require information about the microstructure and composition of the deposited material, but also information about the magnetic behavior of the encountered combinations of microstructure and composition. Additionally, since any such physical characterization would be destructive in nature, techniques for ensuring adequate control over the composition and microstructure of the deposited ferromagnetic films would have to be developed. Such fundamental studies are logical next steps for this research, but without the resulting insight into the microscale magnetic behavior of the magnetic cores, accurate theoretical modeling of anomalous losses will not be possible.

5.5.6.2 Classical Eddy Currents

Even without detailed models of anomalous losses, it is possible to gain some insight to how the presence of losses related to domain magnetization dynamics could be interpreted as classical eddy current losses. Figure 114 summarizes the result of one such theoretical treatment. The ratio of actual losses to expected classical eddy current losses (W/W_C) is shown to be a roughly linearly-increasing function of the ratio between average domain size $2L$ and lamination width d. [50].

Assuming all else to remain constant, this effect, if present in this work, would result in an overestimation of eddy current losses that decreases with increasing lamination thickness. This can be seen, to some extent, in Figure 112.

Neglecting microscale magnetic domain behavior, and confining the analytical modeling to

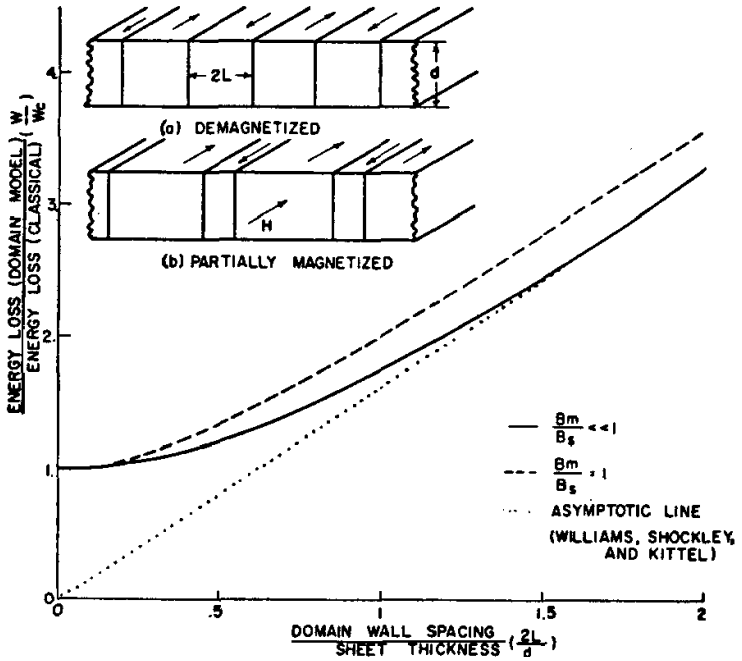

Figure 114: Ratio of of Predicted Eddy Current Losses (W) to Classically-Predicted Eddy Current Losses (W_C) in Thin Lamination (height d) with Longitudinal Domains (width L). [50]

"classical" terms, other conditions that might be present in the microfabricated cores. One such condition might be any curvature in the layers. Depending on the radius of curvature, this might introduce substantial misalignment between the magnetic flux and the layer, in which case the layer's effective cross-section, as seen by the magnetic flux, would be significantly larger than its thickness. This would impede the lamination's ability to prevent eddy currents, leading to higher-than-expected eddy current magnitudes and related losses. As well, any bending of the layers would induce flux crowding at the inner margins of the bends, again leading to higher-than-expected losses.

Another highly plausible condition that could lead to high eddy current losses would be electrical connectivity between laminations. Like bending in the layers, this condition would be due to nonidealities in the sacrificial release process. Clearly, as the division of the core into laminations is specifically performed to limit eddy current magnitudes, any electrical contact between layers would produce higher-than-expected eddy currents.

5.5.6.3 *Classical hysteresis*

As seen in Figure 112, the hysteresis losses were consistently underestimated by a factor of 4x to 8x. It is a well-known fact that the *B-H* hysteresis loop area increases significantly at higher frequencies over its DC value. As well, hysteresis losses scale proportionally with the *B-H* loop area. A hysteresis constant S, relating the *B-H* loop area to the rectangular area inscribed by the *B-H* loop, of 0.2 was used for the hysteresis loss calculations.

However, this loop was characterized using a VSM, and therefore obtained at DC. Accordingly, the increase in hysteresis loop area at high frequencies can at least contribute heavily to, if not fully account for, the underestimation of hysteresis losses.

As well, the presence of iron oxide and other oxides along the surfaces of the plated laminations, or perhaps even as inclusions during electrodeposition, may contribute to the excessive hysteresis losses seen. These oxides typically exhibit high hysteresis losses per unit magnetization, and can thus bear substantial power losses even with small overall volumes. [58]

This would create a physical configuration where the material at the surfaces of the laminations experienced higher specific hysteresis losses than did the material in the interior of the laminations. The model presented in [4] establishes that, at high frequencies, magnetic flux is concentrated at the lamination surfaces and lowered in the lamination's interior. Meanwhile, the hysteresis parameter S was measured under quasistatic conditions, in which the entire volume of the lamination is allowed to reach uniform field intensity and full saturation flux. Therefore, for a given amount of total flux through the magnetic core, the parameter S will underestimate the total hysteresis losses, as it does not account for the dynamic concentration of magnetization at the relatively lossy surfaces.

5.5.6.4 *Underestimation of magnetic core flux*

A wide-reaching source of error in the expected loss calculations might be an underestimation of the magnetic flux density present inside the core. Naturally, power losses of all types increase monotonically with magnetic flux levels, so higher-than-expected flux levels would result in higher-than-expected losses. As well, the exponents of roughly 2 to 2.4 applied to the core flux value in the loss extraction models mean that any amount of relative error present in the flux estimation would be more pronounced in the core loss calculations.

This underestimation could be present at many scales, ranging from the entire cross-section of the core, down to the size of individual magnetic domains.

Macroscale error in flux estimation

The magnetic core's flux density was estimated according to

$$\Phi = \frac{(L - L_w)\, I}{N_t} \tag{139}$$

$$B = \frac{(L - L_w)\, I}{N_t A} \tag{140}$$

where Φ is the total magnetic flux, L is the total measured inductance, L_w is the winding's self-inductance I is the coil current, N_t is the number of turns, A is the core's cross-section area, and B is the magnetic flux density. Conceptually, L_w represents inductance due to all flux paths that *do not* pass through the core at all.

One source of error in the magnetic flux estimation might be overestimation of L_w. A higher-than-expected value of L_w would result in a decrease in B. The air core inductance, as measured with sample AC14, was 220 nH. A value for L_w was assumed as roughly half of this value. The flux paths contributing to L_w were assumed to be present in two distinct regions of the windings:

- around the small hole in the center of the core, where the litz wires are literally abutted to each other, ensuring that even minimal-length flux paths will intersect neighboring turns; and
- along the inductor's two leads, or the two lengths of wire (approximately 2cm each) necessary to connect the inductor to test equipment, acting effectively as isolated straight wires in space.

By varying the value of L_w in Monte Carlo fashion, it was found that the qualitative behavior exhibited in Figure 108 and Figure 111 was not sharply sensitive to the value of L_w. However, any error present in the assumed value of L_w would certainly contribute to the error present in the estimated values for B.

Microscale error in flux estimation

All models used for core loss mechanisms assume that the magnetic material is linear, or that $B \propto H$ in all cases. However, magnetic materials such as the permalloy used in these cores are very clearly nonlinear. They exhibit both saturation and hysteresis behaviors, as a guaranteed minimum set of departures from the linear model.

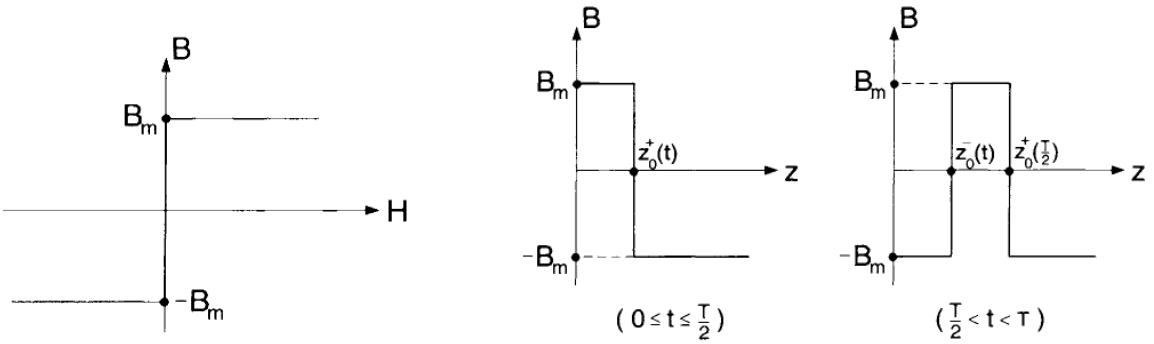

B-H curve
(B_m is saturation flux density)

diffusion of sine wave with
period T along depth coordinate z

Figure 115: Magnetic Diffusion in Nonlinear Material. The material is assumed to have an ideal *B-H* characteristic. [40]

In [40], an examination of the magnetic diffusion principles upon which [4] is based is carried out with a more realistic *B-H* curve, shown in Figure 115.

The assumed *B-H* curve does not account for hysteresis, but does model saturation, while assuming a very high permeability. The results of this analysis suggest that AC magnetic fields diffuse from the surface of a lamination inward not with the exponentially decaying profile seen in classical eddy current analysis (*ex.* [4]), but rather with sharp spatial edges. This is shown in the right-hand figures of Figure 115.

To the extent that this analysis describes the actual case, the upshot of this discrepancy would be a radically different modeling scenario. Instead of all magnetic material being partially magnetized, all material would be fully magnetized in one direction or another. As well, as frequency and magnetomotive force are varied, according variations in the core's magnetization would be achieved by changes in the volume of magnetized material. Lower frequencies would allow the magnetization waves to penetrate deeper into the laminations, while the number of laminations magnetized and the lateral extent of magnetization would modulate according to the MMF.

While difficult to model, this would entail that all magnetized material is operating at saturation flux levels, and that all AC magnetization is either of zero magnitude or maximum swing (*i.e.* from $-B_{sat}$ to $+B_{sat}$). For the same total core flux, such a maximally-concentrated flux distribution would result in higher losses due to the exponents of 2 (and higher) to which the flux density is raised in core loss models. Mathematically, this can be seen as the departure developed between the linear

summation of local flux densities into total core flux and the sum-of-squares summation of local flux densities into core loss terms.

5.5.6.5 Unanticipated Real Losses

In addition to underestimation of well-known loss mechanisms, it is possible that some less common, and perhaps novel, loss mechanisms are active in the characterized devices.

Nonequilibrium thermodynamics

For heat transfer by conduction, the characteristic diffusion length δ is given by

$$\delta = 2\sqrt{\frac{k}{\rho c_P}t} \tag{141}$$

where k is the thermal conductivity, t is time, *rho* is density, and c_P is heat capacity. [68] For time-varying heat dissipation caused by magnetization at frequency f, the characteristic length becomes

$$\delta = \sqrt{\frac{k}{\rho c_P}\frac{2}{f}} \tag{142}$$

Using standard values for permalloy of $k = 17$ W/(m·K), $c_P = 520$ J/(kg·K), and $\rho = 8450$ kg/m^3, this becomes

$$\delta = 2.78\ \mu\text{m}\sqrt{\frac{1\text{ MHz}}{f}} \tag{143}$$

meaning that at 1 MHz, the characteristic length of thermal diffusion is only a few microns. While this is still significantly greater than the lamination thicknesses at the frequencies used in characterization, it is significantly less than any other dimension of the laminations.

This means that excitation of the core at such high frequencies does not allow thermodynamic equilibrium to be achieved. This, in turn, creates the potential for unexpected interaction between magnetization and other modes of energy storage. Well-explored examples of such interactions include the magnetocaloric effect and the thermoelastic effect. Any of these nonequilibrium effects that would occur in appreciable amounts would also induce a related amount of damping and, therefore, energy loss. Unfortunately, detailed analyses of such phenomena are very complex, and also require exacting empirical measurements to enable any attempt at correlation.

Note that these effect would be dramatically more pronounced for magnetic multilayers than in typical single-film cases, where the film is deposited onto a substrate, perhaps with a very thin

adhesion or interface layer. A single thin film, then, will have have excellent thermal relief to the substrate, assuming that any adhesion or interface layer will be much thinner than the aggregate characteristic thermal diffusion length. For this reason, it is plausible that such effects could be very active in the results presented in this thesis, yet largely unexplored by the broader magnetics community. This effect would also be exacerbated by any full-swing magnetization models, such as [50]. In these, very high effective thermal frequencies are seen as domain walls move rapidly and induce eddy current losses.

Magnetostriction-related losses

Magnetostriction is a transduction mechanism that relates magnetization with mechanical deformation. This phenomenon could allow a number of loss pathways that are not treated by the analytical models presented.

Certainly, any time-varying stress distributions within the laminations, regardless of their origin, could potentially lead to actuation of the layers. If the layers are deformed during the release process, or if the stress distributions are nonuniform in such a manner to create warping or buckling, mechanical displacement could result. This displacement, or passive actuation, would necessarily require that some portion of the mechanical work performed by the magnetostrictive actuation would be converted to heat and manifested as core losses. If the stress distribution is viewed as compressive stress that is released principally through vertical strain (*i.e.* laminations growing thicker and thinner over time) squeeze film damping could become a nontrivial effect, especially if temperatures are reached at which the interlayer dielectric becomes liquefied. Finally, any amount of time-varying strain will induce some amount of acoustic energy dissipation due to "internal friction" within the laminations. [7]

5.6 Conclusions

5.6.1 Intractability of Electromagnetic Theory

While modeling of ideal magnetic materials and regular structures yields germane analytical solutions, taking account of microscale magnetic dynamics as well as structural nonidealities quickly leads to an unwieldy analysis. Particularly when highly nonlinear dynamics as well as memory

effects such as hysteresis are considered, even numerical simulation becomes impractical. The result is a set of inadequate analytical relations that must be used to make sense of very complicated experimental data.

More convenient analytical forms would go a long way towards advancing the interpretation of observed systems, as would the availability of magnetic simulation tools that can efficiently handle hysteresis and saturation behavior. It is also possible that numerical simulation tools which accurately model microscale magnetic phenomena such as domain wall movement and spin interactions would provide great clarity into the behavior of these highly-laminated cores.

5.6.2 Utility of Automated Characterization

In completing the work detailed in this chapter, the effort required to automate the loss characterization was greatly rewarded. Empowering a modern microprocessor to run the experimentation as much as possible resulted in results that were more accurate, more complete, and more quickly obtained. As well, countless hours of operator attention were liberated for other concerns.

It was found to be of prime importance that the HFHF loss measurement system was architected, from the ground up, to support automation. As well, wherever possible, the selection of automation-friendly, high-speed computer interfaces for instruments, is important to derive maximum dividends from automation.

5.6.3 Fragility of Microfabricated Cores

It was also found that these highly-laminated cores were extremely susceptible to heat damage and other forms of mechanical damage, such as bending or tearing of the individual layers. This dictates an important directive for future research: more robust physical packaging for cases where the core will be separated from its silicon substrate. Bobbins formed from laser-cut, low-k ceramics such as silica should provide the required mechanical stability and precision, as well as thermal relief, without introducing excessive loss mechanisms or nuisance electromagnetic effects (such as displacement currents or dielectric dissipation). In the meantime, this also dictates that during experimentation with the cores, the survival of the samples be held sacrosanct.

The development of the bobbin structure reflects much of this learning. Additionally, it would have been very benificial to build a thermal sensor, or a fuse, into each inductor's package. Naturally,

the delivery of any power to the inductor itself would then be interlocked to the thermal sensor or fuse, so that when the inductor reaches some thermal limit, the test equipment or power converter in question immediately stops delivering power to the inductor.

CHAPTER VI

APPLICATION IN POWER CONVERTER

Although the core losses of the highly-laminated inductors followed expected trends, fully establishing the relevance of the sacrificial multilayer plating process to advanced power converters requires demonstration of an advanced power converter based on a highly laminated-core inductor. This chapter presents precisely this element of the research.

6.1 Advanced Power Converter Design

6.1.1 Design Goals

The power converter circuits used for this demonstration were both DC/DC converters, and have very similar architectures. There were a few major goals that guided the design and implementation of these converters. These goals resulted in converters that do not conform tightly to typical power converter goals of maximized efficiency, compactness, and optimized closed-loop control.

First, these converters were designed to support comparative evaluation of inductors at high switching frequencies. A switching frequency range of 0.1-12 MHz was targeted, as this would allow access to much of the operation space targeted by high-frequency SMPSs. This goal also led to the use of discretely-packaged, high-performance MOSFETs, to achieve high power dissipation capacity and very low on-state resistance. This, in turn, required the use of discrete MOSFET gate drivers to achieve fast, consistent switching over the entire frequency range. The use of MOSFETs with high gate capacitance and low on-state resistance ensured that variations in the electrical properties of the inductors under test would amount to minimal variations in the switching and operation of the MOSFETs.

Second, these converters were intended to provide as much flexibility as possible. Most importantly, this meant allowing arbitrary selection of any frequency and duty cycle. This goal also favored the use of a user-friendly control interface for the converters, rather than the resistor-programming interface used by many SMPS controller ICs. The use of a Cypress PSoC microcontroller as the converter controller was found to provide a complete solution to the need for

flexibility. In combination with an UART-to-USB interface IC (FTDI FT232R), the microcontroller could readily be accessed from a PC. As well, the PSoC microcontroller not only provides built-in, fully-programmable PWM generation up to 24 MHz, but also has an array of configurable on-board analog and digital functional blocks and interconnect.

The final design goal for these converters was to serve as an easily-instrumented platform from which not only inductor performance could be assessed, but also from which future designs could be derived. These goals culminated in the relatively sparse layout as seen in Figure 116 and Figure 117, including the large cutout to support the plastic bobbins used to produce the highly laminated core inductors. This sparse layout incorporates multiple test points and circuit configuration options, while also providing room not only to accommodate test equipment probes but also to add components and functionality in future revisions.

6.1.2 Circuit Design

Given the chosen approaches to address the design goals as stated above, the circuit design was relatively straightforward. The microcontroller circuit design was borrowed from the microcontroller-based PCB designs used for the multilayer plating robot. The microcontroller circuit provided power supply decoupling, signal breakout, and the 5-pin programming header required to download programs to the microcontroller. The UART-to-USB daughter board was based on a reference design, and functioned as a self-contained unit.

The same MOSFETs were used for both converters, and they were ST Microelectronics STP22NF03L N-channel MOSFETs, with under 50 $\mu\Omega$ of on-state resistance, 22 A drain current rating, and 6.5 nC of gate charge. The Microchip MCP1404 gate driver was used for the boost converter. It provides up to 4.5A of gate current per output, and its two outputs were ganged together to drive a single MOSFET. The National Semiconductor LM2725 high and low side driver IC was used for the buck converter. This gate driver, while it only provides 1.2 A of gate current, allows the use of matching N-channel FETs for both switches in the buck converter.

6.1.3 Implementation

The schematic and implementation of the buck converter are shown in Figure 116. A synchronous switching approach was chosen to minimize switching losses and to maximize the symmetry of

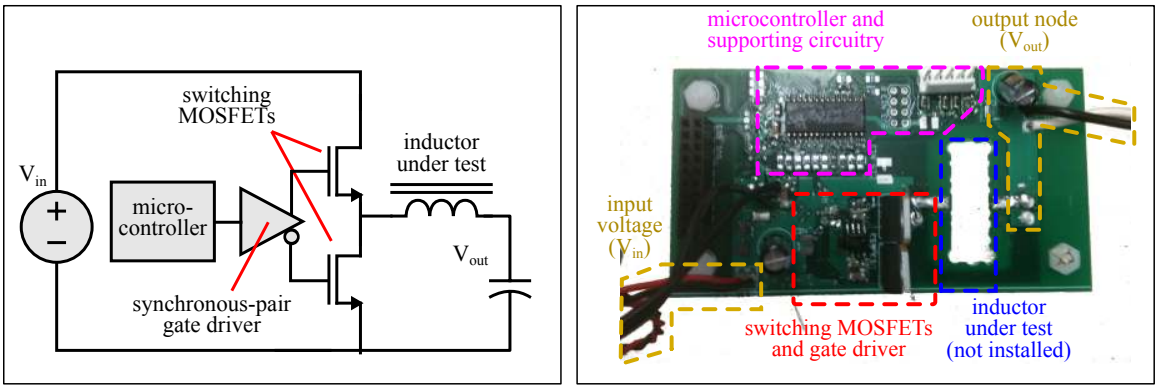

schematic implementation

Figure 116: Buck Converter, schematic and implementation

the switching waveforms. The schematic shown was implemented in a 2-layer, 3.75” x 2” printed circuit board. At the point in the circuit where the inductor under test was connected, a 0.3” x 1” cutout was included in the PCB design to accommodate the packaged inductors, including any hardframe structures.

The PCB design was intentionally sparse, to allow maximum adaptability and to minimize parasitic interactions between subcircuits. The converter divides into five subcircuits, shown in Figure 116. The high-power switching elements, particularly the MOSFETs and their gate drivers, constitute the most important subcircuit. The input and output subcircuits consisted simply of appropriate terminations and noise bypass and filter capacitors.

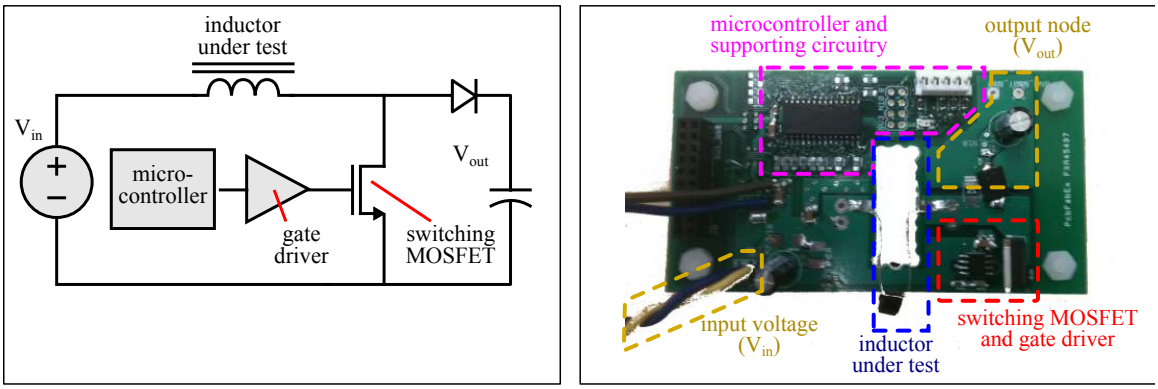

schematic implementation

Figure 117: Boost Converter, schematic and implementation

The schematic and implementation of the boost converter are shown in Figure 117. This converter, like the buck converter, was implemented as a 3.75” x 2” printed circuit board, including the

same 0.3” x 1” cutout to accommodate packaged inductors. As well like the buck converter, this converter divides naturally into five subcircuits, as depicted in Figure 117.

6.1.4 Inductors Used

Three different inductors were used in the operation of the two power supplies above. Details are given below. All 3 inductors were characterized using the HFHF core loss characterization system described in Chapter 5.

6.1.4.1 Ferrite Inductor

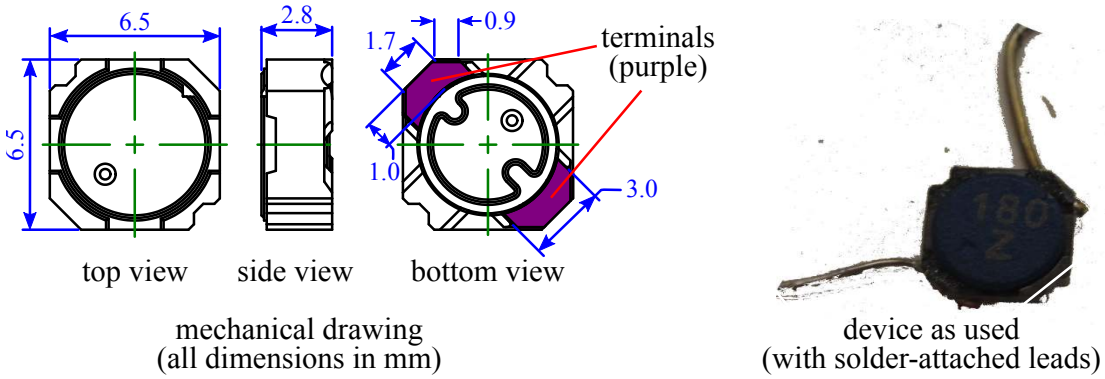

Figure 118: Ferrite inductor used for power supply testing

The ferrite inductor used with these power converters was an 18 μH Sumida power inductor, part number CDR6D28MNNP-180NC. Its physical construction, based on a pot core construction, is detailed in Figure 118. Its DC resistance was 150 $\mu\Omega$ and its rated saturation current was 1.3 A.

6.1.4.2 Low-Volume Laminated-Core Inductor

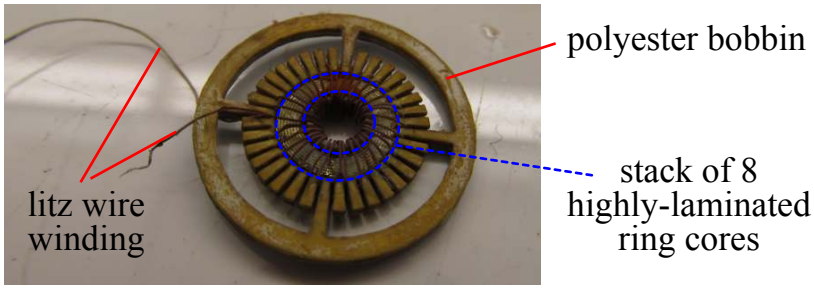

Figure 119: Low-Volume laminated-core inductor used for power supply testing

The laminated-core inductor used for the bulk of the power supply testing is shown in Figure 119. This inductor, fabricated by Florian Herrault, Joon-Cheol Kim, and Min-Soo Kim using the multilayer process described in Chapter 3, comprised 48 turns of litz wire (7 strand, 46 AWG) as its winding. Its core consisted of 8 stacked cores, each consisting of 40 1 μm thick permalloy layers.

6.1.4.3 High-Volume Laminated-Core Inductor

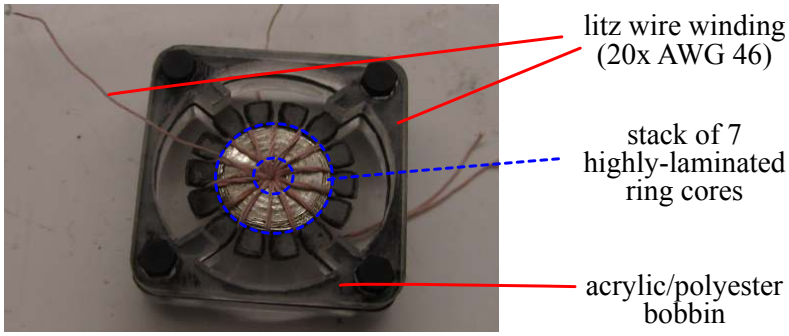

Figure 120: High-Volume laminated-core inductor used for power supply testing

A second laminated-core inductor was used to extend the power supply's operation to higher output powers. The inductor used for this effort is shown in Figure 120. This inductor's core was a stack of seven remnant cores from the work described in Chapter 5, with layer thicknesses of 1.33 and 1.66 μm . This inductor used the same packaging approach, including the same bobbin design, used for the inductors characterized in Chapter 5. However, to boost inductance and thus SMPS efficiency, this inductor's winding consisted of 28 turns of 20 $\times$ 46AWG wire, rather than the 14 turns of 7 $\times$ 46AWG litz wire used as the standard winding in Chapter 5.

6.2 Experimental Setup

6.2.1 Equipment Used

The equipment used in the power converter testing is shown schematically in Figure 121. An adjustable DC power supply was used to provide input power, while a large adjustable resistor was used as the load. Current sense resistors were included at the input and output of the power supply to allow accurate determination of the input and output currents. For dynamic characterization, oscilloscope probes were placed at the output and switching nodes, and a high-frequency current

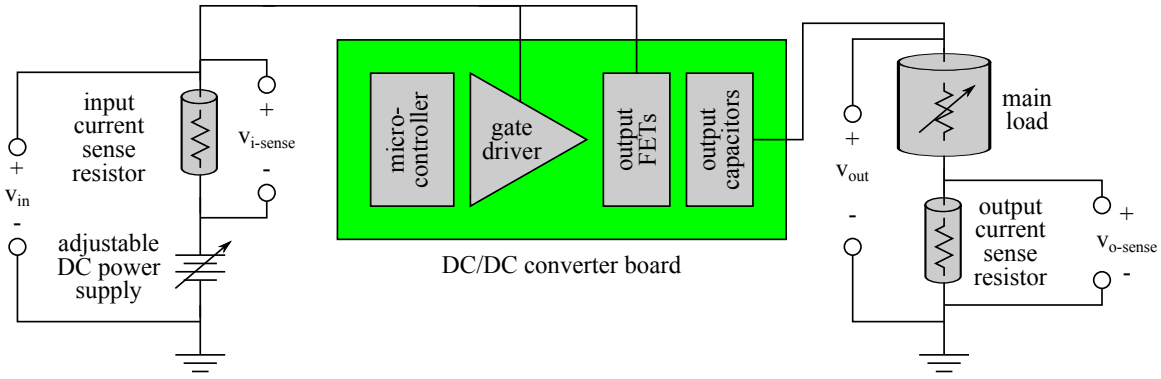

Figure 121: Operational Setup for Power Converter Testing (schematic)

probe was latched around an inductor lead to capture the inductor's current waveform.

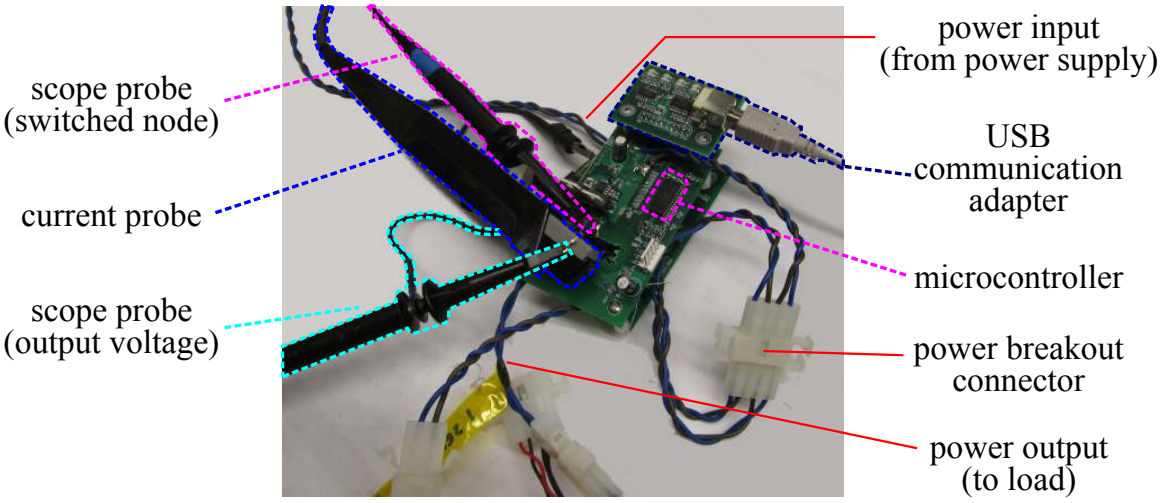

Figure 122: Operational Setup for Power Converter Testing (implementation)

The setup used is shown in Figure 122. The power breakout connector, not shown in Figure 121, was added to allow independent control over the 5 Volt supply and the power rail supply.

6.2.2 Efficiency Calculation

The operating efficiency of the power supply was calculated as η:

$$\eta = \frac{P_{out}}{P_{in}} \tag{144}$$

where P_{out} and P_{in} are the output and input power, respectively. These quantities were calculated in a straightforward manner from the noted quantities in Figure 121.

$$P_{in} = V_{in} I_{in} \tag{145}$$

$$= \left(V_{in-tot} - \frac{V_{in-sense}}{R_{in-sense}} \right) \frac{V_{in-sense}}{R_{in-sense}} \tag{146}$$

$$P_{out} = V_{out} I_{out} \tag{147}$$

$$= V_{out} \frac{V_{out-sense}}{R_{out-sense}} \tag{148}$$

It should be noted that the power delivered by the 5 Volt was not included in the input power figure. This is consistent with the use of these power supplies as highly flexible high-frequency testbenches. The components that enabled these qualities – the microcontroller and high-speed FET gate drivers, respectively – are responsible for a majority of the power consumed on the 5 Volt rail. In a design better suited for practical applications, both of these components would be replaced by lower-cost, less capable alternatives that would also consume far less power, and therefore not present a dominant efficiency limitation.

6.3 Dynamic Observations

6.3.1 Operation with Ferrite Inductor

Waveform captures of the buck converter's operation with the ferrite inductor are shown in Figure 123, Figure 124, and Figure 125. The waveforms are generally as would be expected for a buck converter. The switched-mode voltage has a square waveform, the inductor current has a triangular waveform, and the output voltage remains very constant.

However, there is considerable ringing on the switched-voltage waveform, visible in all captures. This ringing did not depend on which inductor was used, and was assumed to be the unavoidable interaction between parasitics on the buck converter PCB and the extremely fast edges on the switched-voltage waveform.

The PCB needed to have a sparse layout to facilitate testing and accommodate the setup shown in Figure 121, leading to substantial interconnect parasitics. At the same time, very fast switching times for the MOSFET outputs were required to support the high switching frequencies targeted.

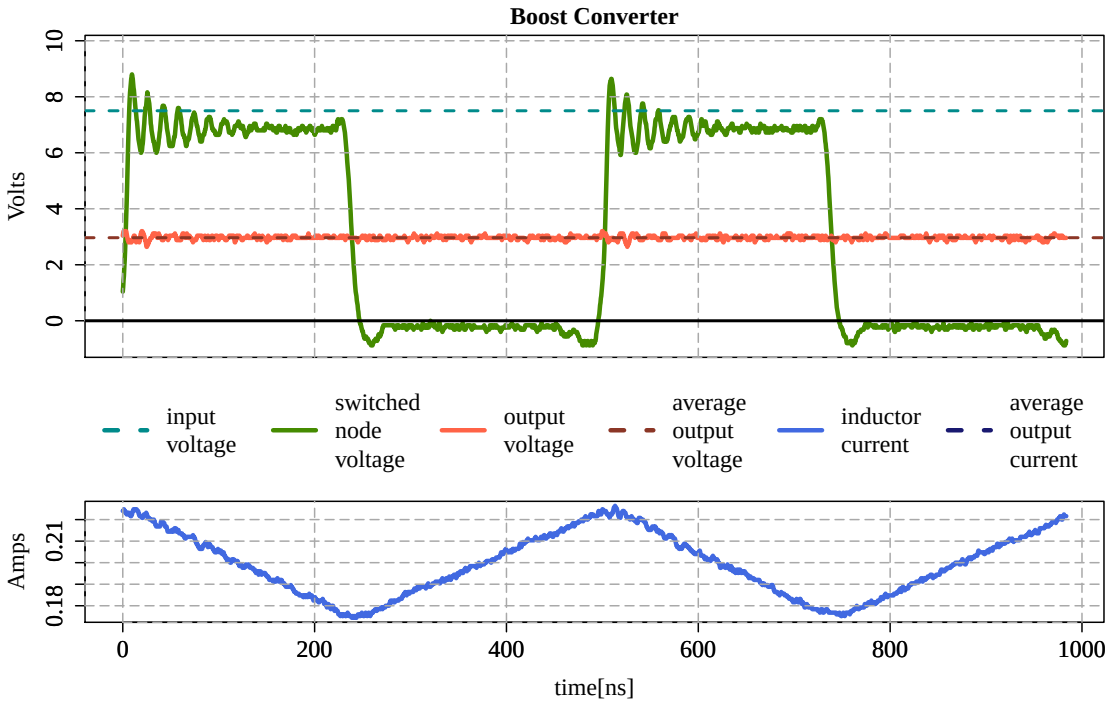

Figure 123: Buck Converter Waveforms with Ferrite Inductor, 2MHz

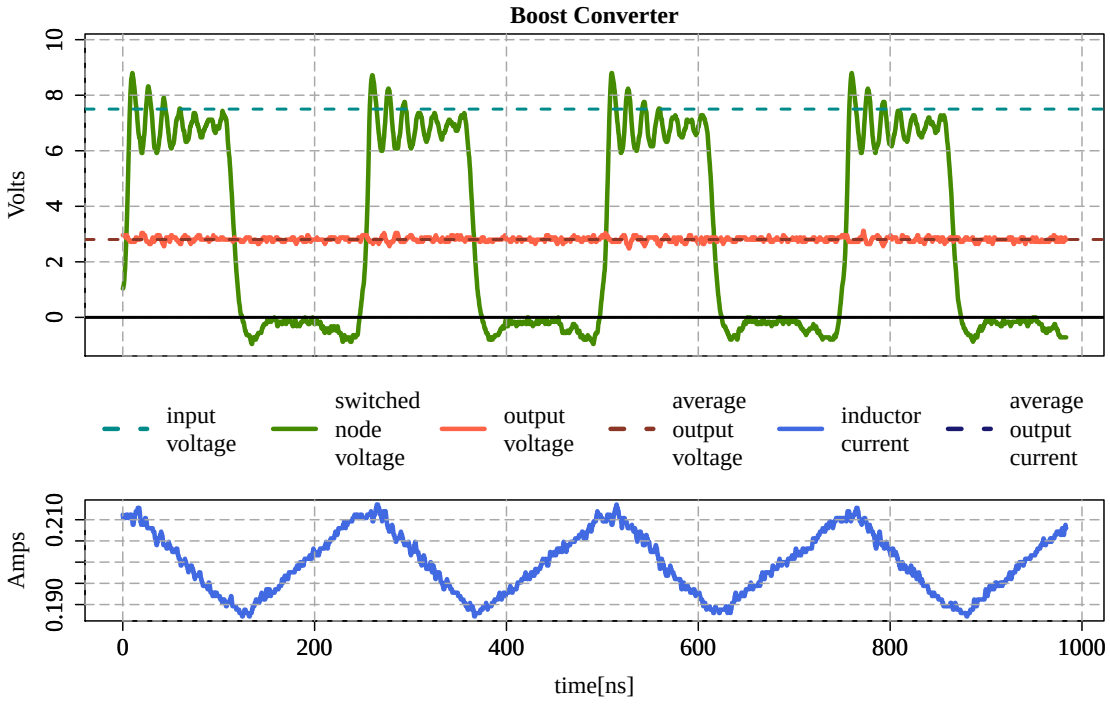

Figure 124: Buck Converter Waveforms with Ferrite Inductor, 4MHz

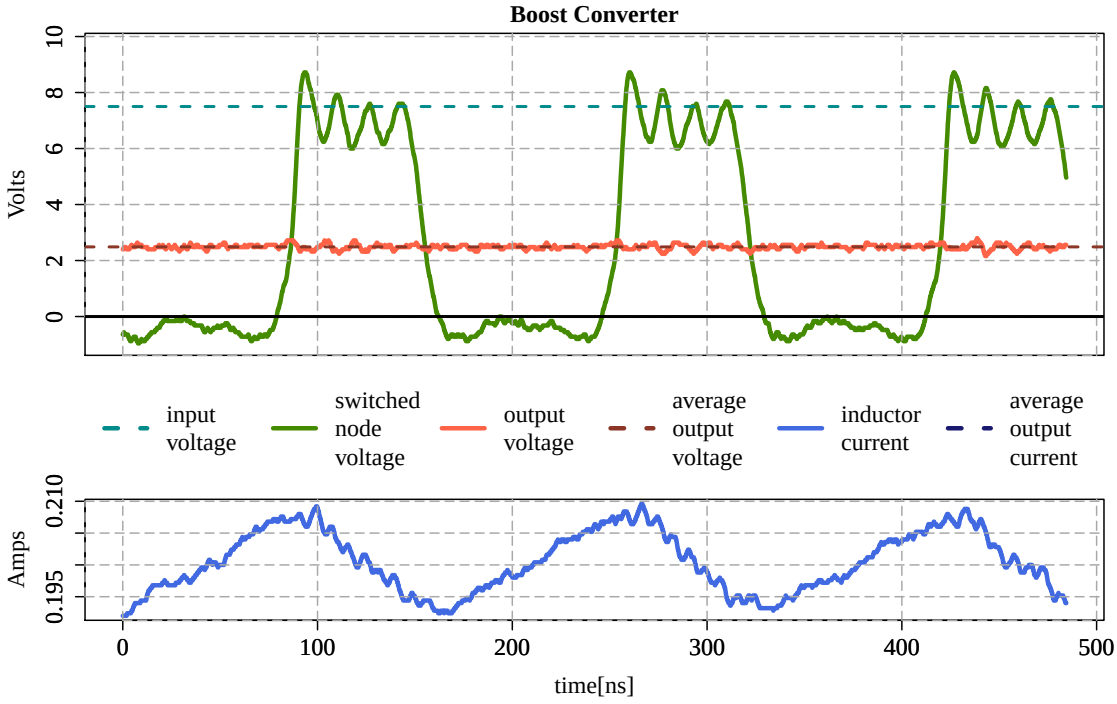

Figure 125: Buck Converter Waveforms with Ferrite Inductor, 6MHz

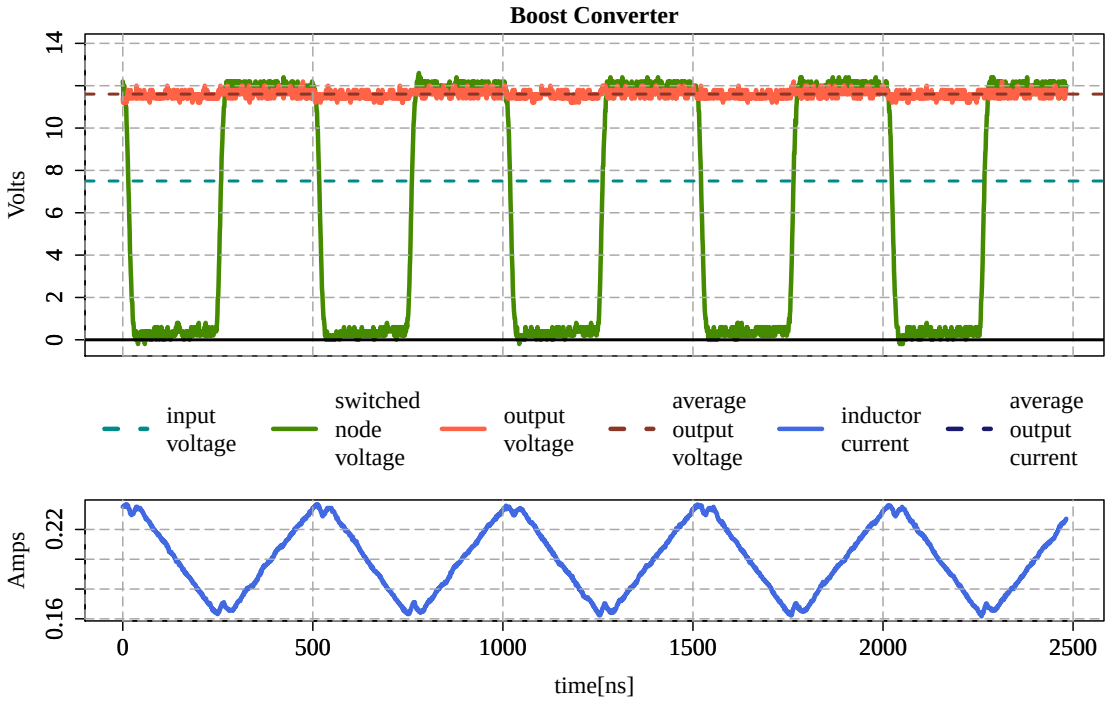

Figure 126: Boost Converter Waveforms with Ferrite Inductor, 2MHz

Waveform captures of the boost converter's operation with the ferrite inductor are shown in Figure 126, Figure 127, and Figure 128. These waveforms also are qualitatively as expected. The switched-mode voltage waveform is square, the inductor current waveform is triangular, and the output voltage is constant.

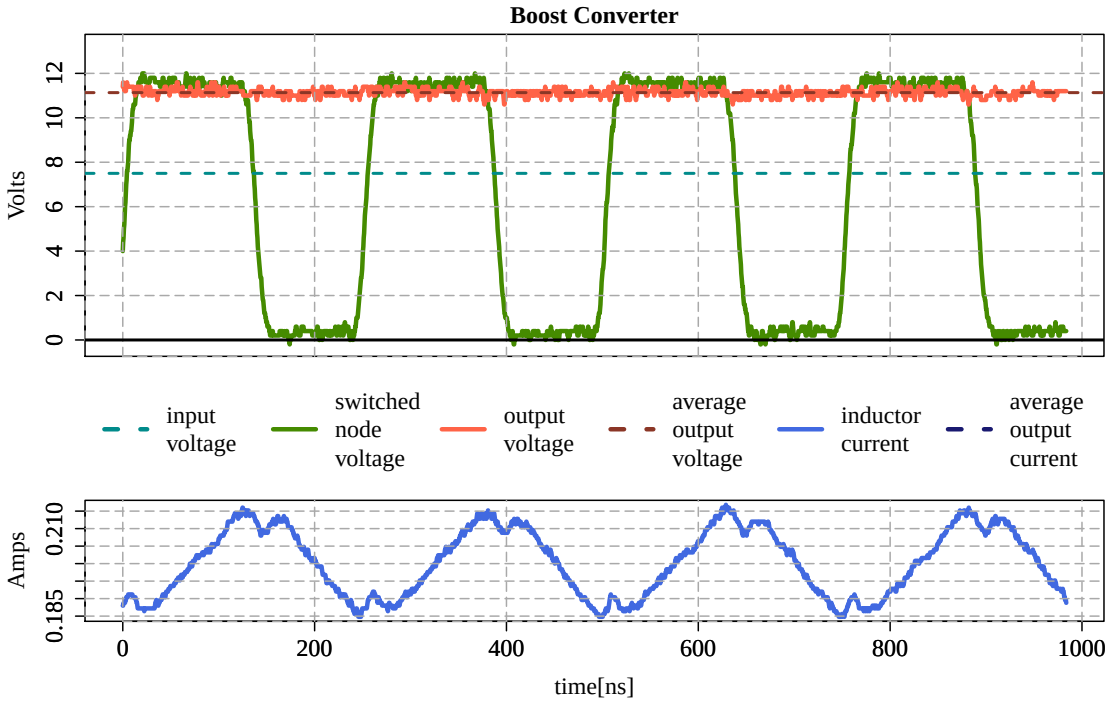

Figure 127: Boost Converter Waveforms with Ferrite Inductor, 4MHz

The principal nonideality in this converter's operation is the presence of "notches" at the positive and negative peaks of the inductor waveforms. These are minor effects in the waveforms taken at 2 MHz switching frequency (Figure 126), but dominant features at 6 MHz switching frequency (Figure 128). These notches were also only seen with the ferrite-core inductor. They are speculated to potentially be the result excessive winding capacitance.

6.3.2 Operation with Laminated-Core Inductor

Waveform captures of the boost converter's operation with the laminated-core inductor are shown in Figure 129, Figure 130, and Figure 131. At 2 MHz (Figure 129), the inductor current waveform shows a slight concavity, likely a result of its higher winding resistance. As well, the AC amplitude of the inductor current waveform is considerably larger for this inductor than for the ferrite

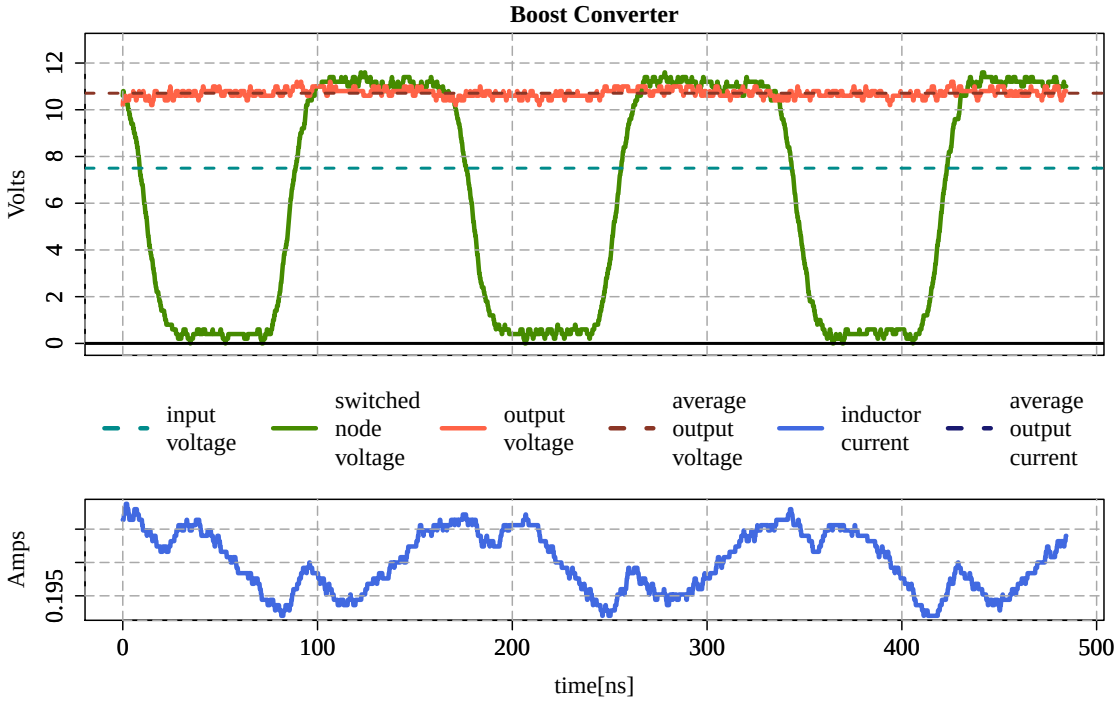

Figure 128: Boost Converter Waveforms with Ferrite Inductor, 6MHz

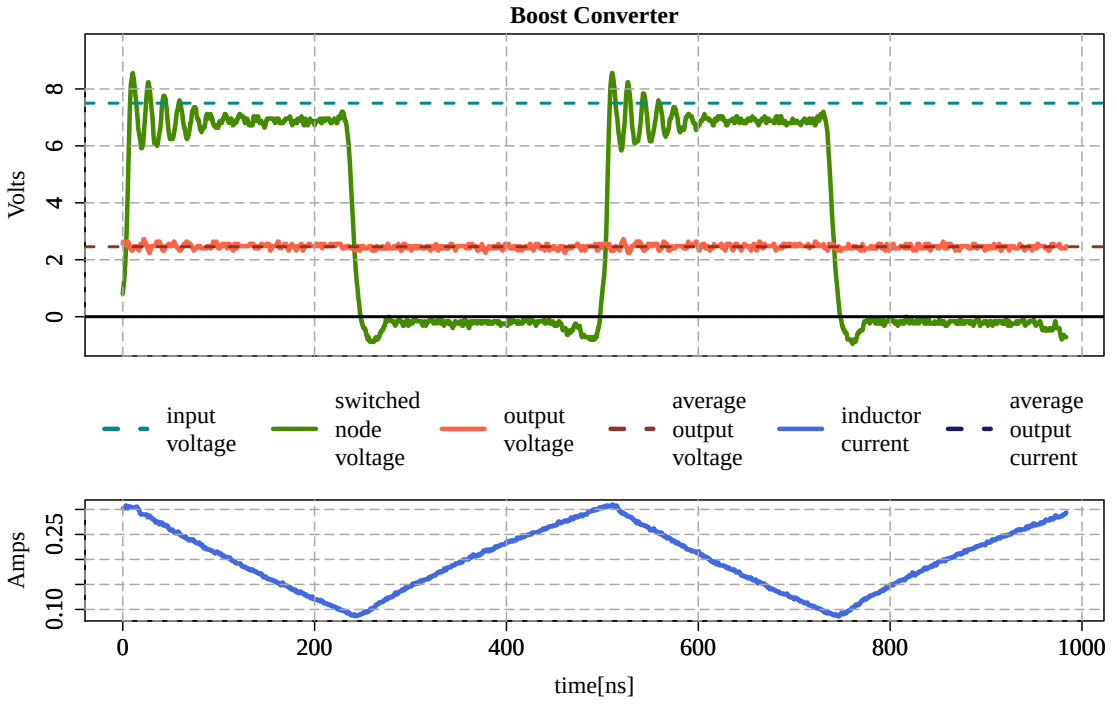

Figure 129: Buck Converter Waveforms with Laminated-Core Inductor, 2MHz

inductor. This is due not only to this inductor's lower nominal inductance, but also to its decrease in inductance with increasing operating frequency.

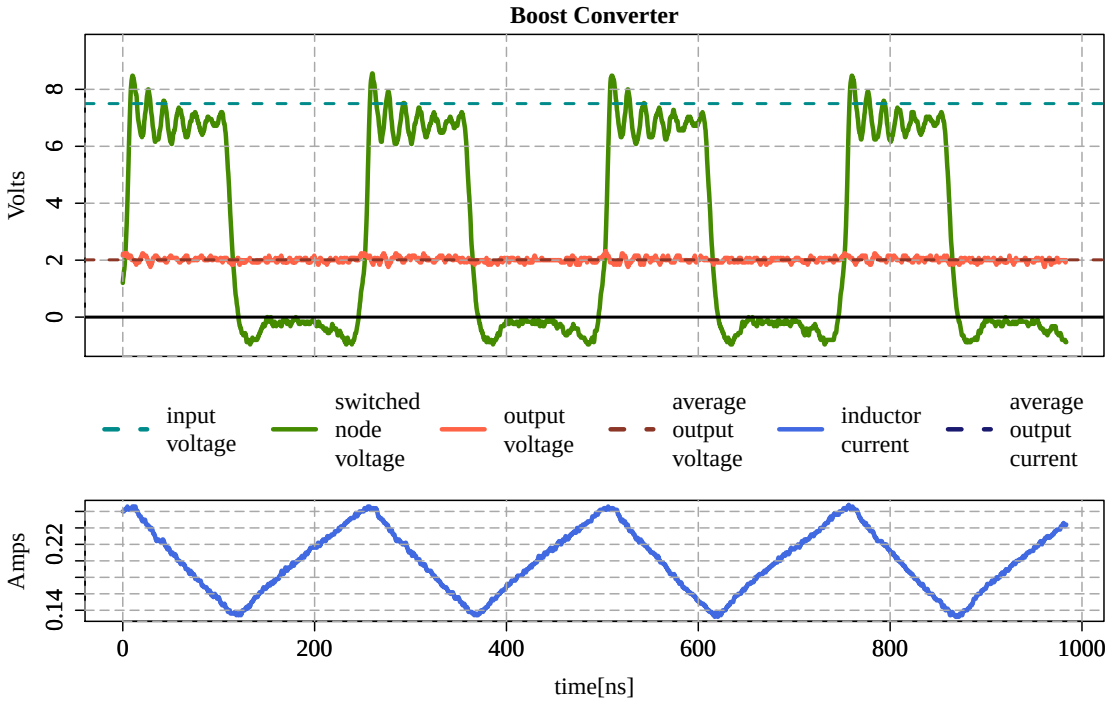

Figure 130: Buck Converter Waveforms with Laminated-Core Inductor, 4MHz

Waveform captures of the boost converter's operation with the laminated-core inductor are shown in Figure 132, Figure 133, and Figure 134. As with the buck converter, the amplitude of the inductor current waveform is higher with this inductor than with the ferrite inductor.

There is also a pronounced rounding of the positive and negative peaks of the inductor current waveform. This artifact is unexplained, but is likely beneficial as it serves to curb the high-frequency content of the winding current.

6.3.3 Operation with High-Volume Laminated-Core Inductor

Waveform captures of the boost converter's operation with the high-volume laminated-core inductor are shown in Figure 135 and Figure 136. The purpose of the high-volume inductor was to expand the operating envelope of the converters. Accordingly, the boost converter was operated at very high duty cycles, to produce a maximum output voltage. An intermediate switching frequency of 4.8 MHz was chosen.

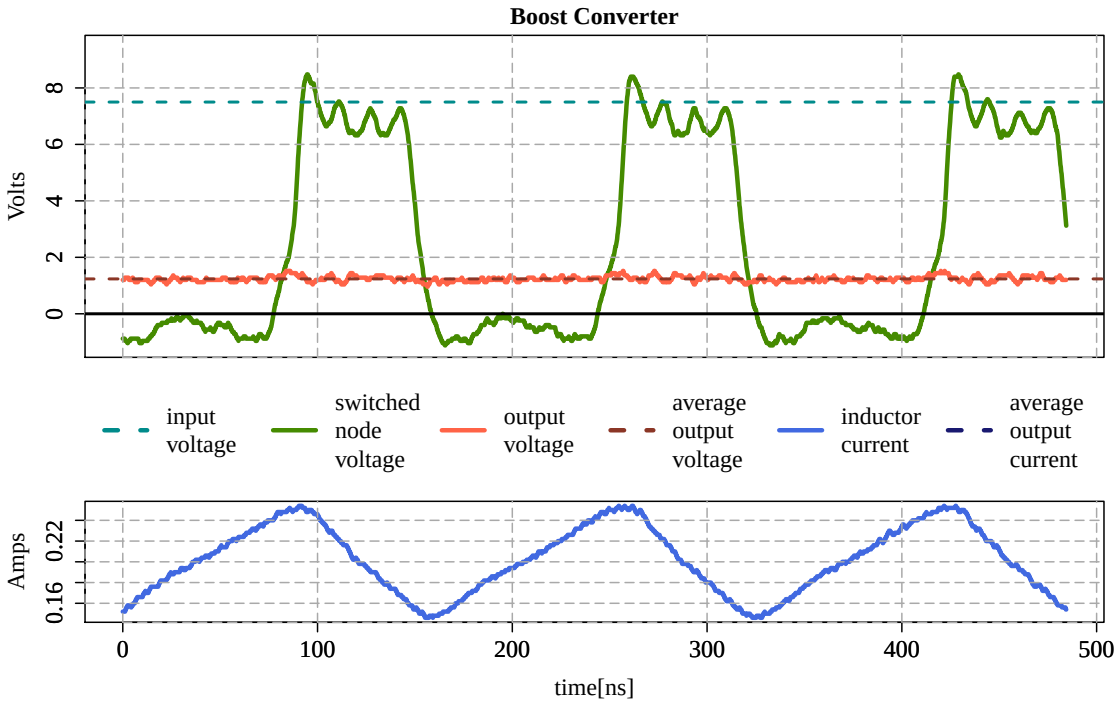

Figure 131: Buck Converter Waveforms with Laminated-Core Inductor, 6MHz

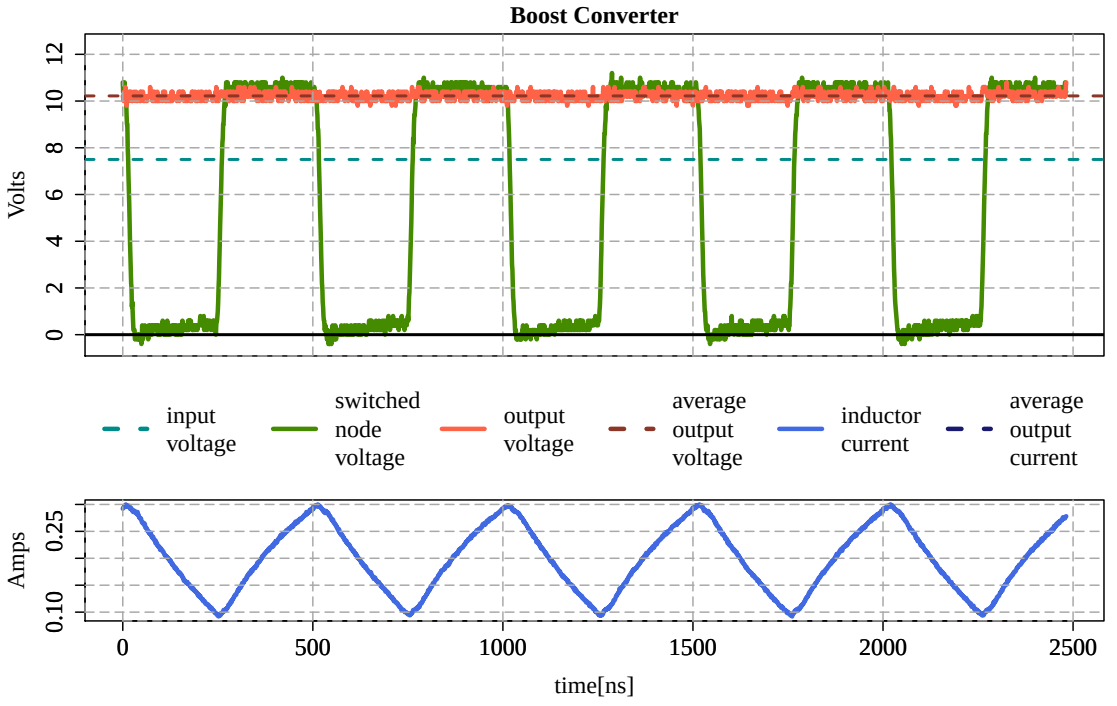

Figure 132: Boost Converter Waveforms with Laminated-Core Inductor, 2MHz

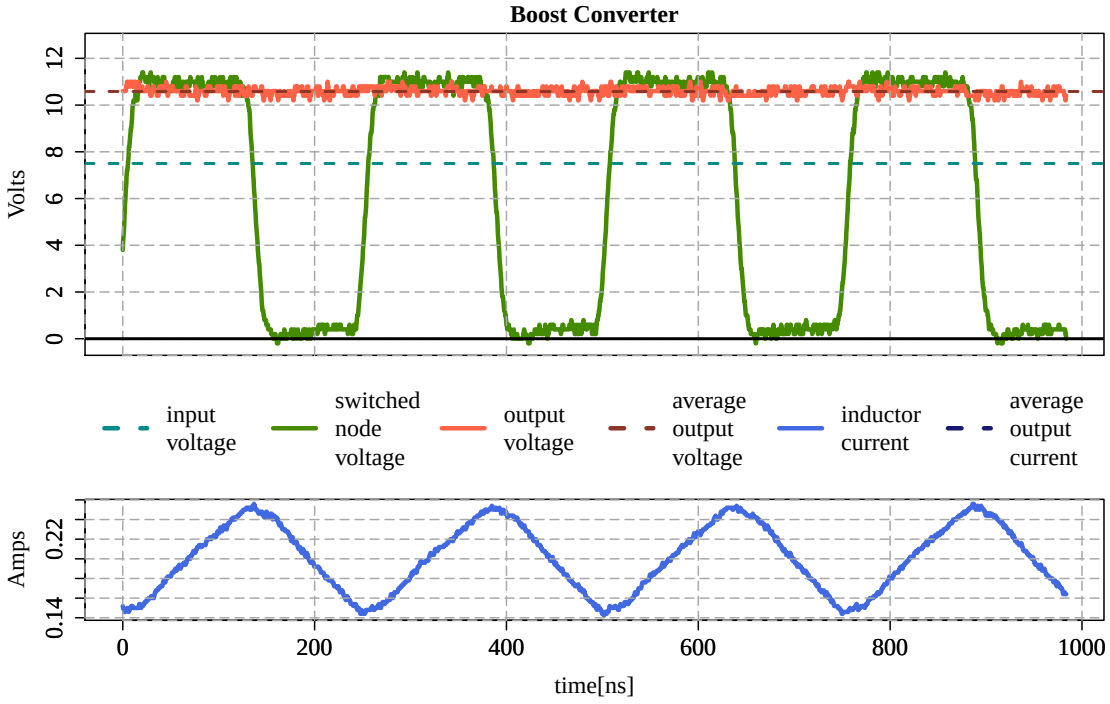

Figure 133: Boost Converter Waveforms with Laminated-Core Inductor, 4MHz

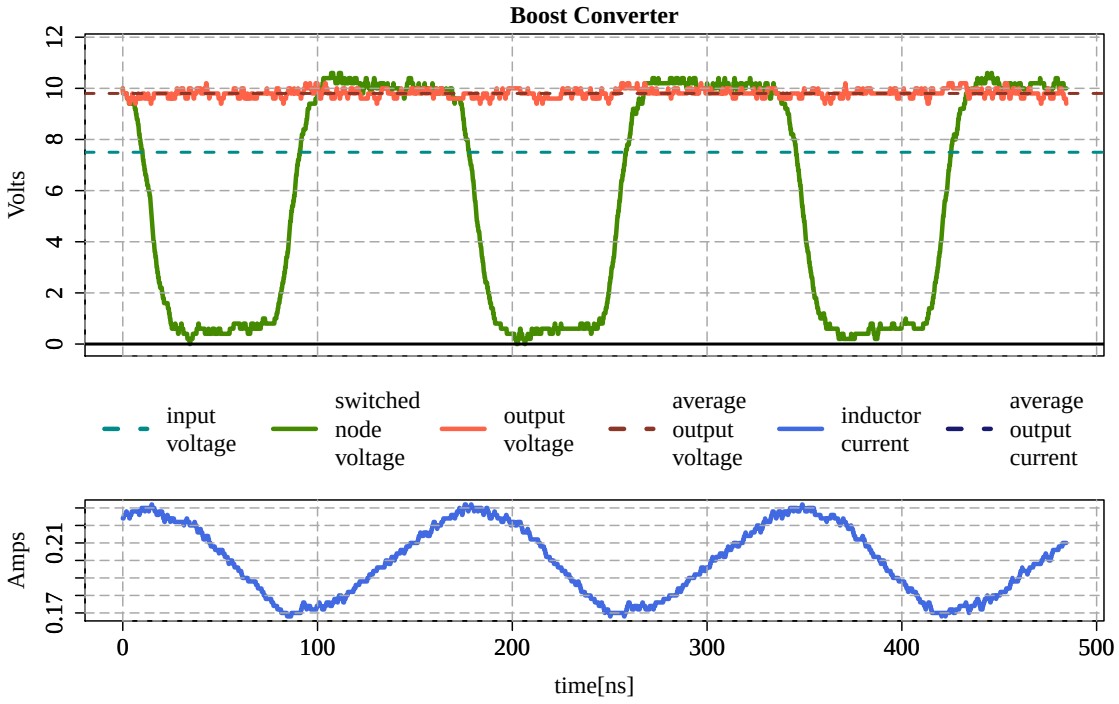

Figure 134: Boost Converter Waveforms with Laminated-Core Inductor, 6MHz

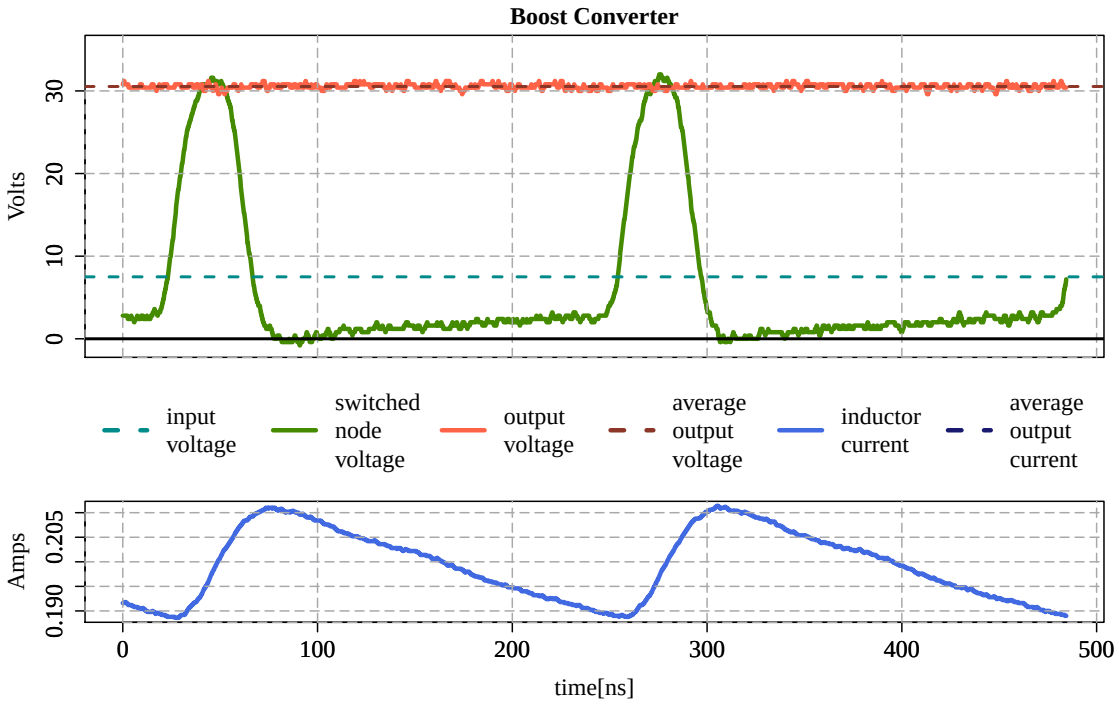

Figure 135: Boost Converter Waveforms with High-Volume Laminated-Core Inductor, 4.8 MHz

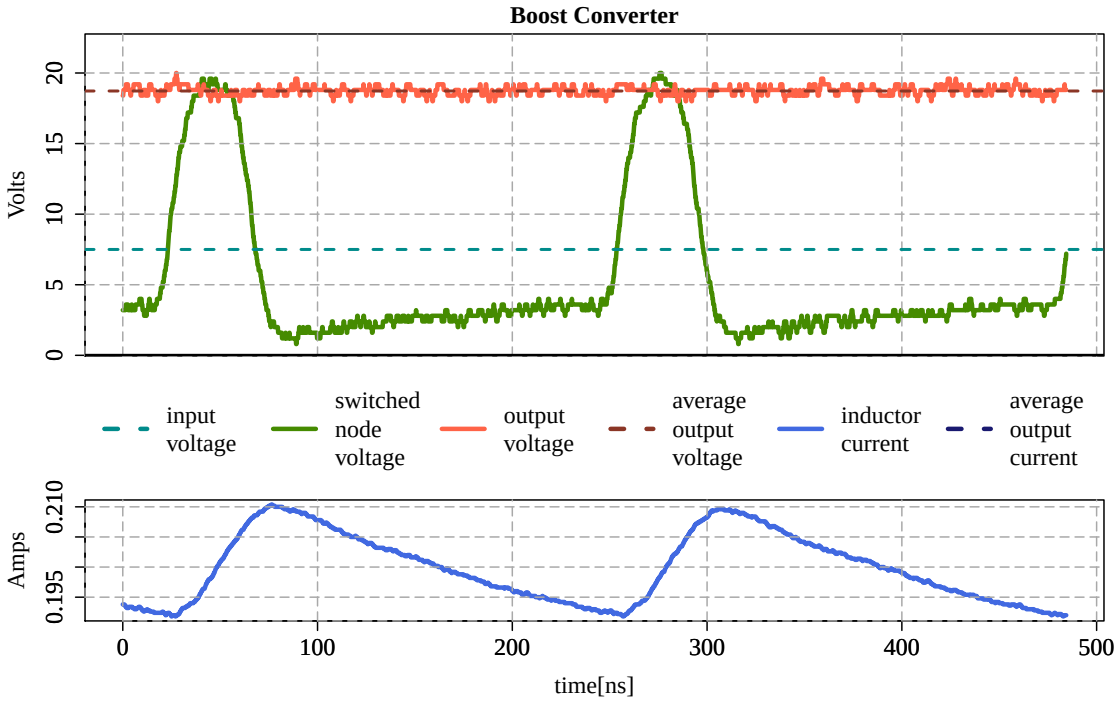

Figure 136: Boost Converter Waveforms with High-Volume Laminated-Core Inductor, 4.8 MHz

As shown, 30 V output voltage was readily obtained with qualitatively correct waveforms. Some nonidealities are clearly visible, including rounding of the inductor current and switching voltage waveforms as well as a slow ramp upwards during the inductor's charging period.

The buck converter was also operated with the high-volume inductor, but waveforms were not captured. An output power of 5 Watts at a switching frequency of 6 MHz was achieved, at which point the MOSFETs exceeded their thermal capacity.

6.4 Performance Analysis

6.4.1 Compensation for Inductor Losses

All inductors used with these power supplies were comprehensively characterized with the HFHF characterization system described in Chapter 5. The loss measurements of these inductors can be correlated with their performance in the converters presented in this chapter.

Unfortunately, precise characterization of winding parasitics and leakage inductance was not available for these devices, so their intrinsic core losses could not be readily explored. However, the device-level loss measurements could still be fitted to core loss models by incorporating a loss term that is proportional to winding current, and otherwise assuming that core flux is proportional to winding current.

More precisely, the loss model that was fit to each inductor's HFHF data was:

$$P_V = P_w + P_h + P_e + P_a \tag{149}$$

$$= K_w i^2 + K_h i^2 f + K_e i^2 f^2 + K_a i^{1.5} f^{1.5} \tag{150}$$

$$= \begin{bmatrix} K_w \\ K_h \\ K_e \\ K_a \end{bmatrix} \cdot \begin{bmatrix} i^2 \\ i^2 f \\ i^2 f^2 \\ i^{1.5} f^{1.5} \end{bmatrix} \tag{151}$$

where the subscripts w, h, e, and a pertain to losses associated, respectively, with winding loss, hysteresis, eddy currents, and anomalous losses. Also, f is frequency, i is the RMS value of winding current, which in the HFHF testing was sinusoidal with no DC component.

One primary source of uncertainty in this extrapolation is the difference between the frequencies attainable in the HFHF (roughly 1 MHz and under) and the converters' switching frequencies (2-6

MHz). The combination of the inductors' very large inductances and the HFHF's range of matching capacitances, which was optimized for inductances of 0.5-2.0 μH , is what limited these devices' HFHF data to such low frequencies. During the extrapolation from the HFHF frequencies to the SMPS frequencies, this 2x-10x mismatch in frequency is squared, magnifying any mismatch or physical behaviors not captured in the lower-frequency HFHF data.

Another source of uncertainty in the use of HFHF data to compensate observed converter behavior is the difference in the spectral content of the winding current. In the HFHF, the winding current is a very pure sine wave with no DC content; in the converters, the winding current is a sawtooth wave with a DC content that is ideally much larger than the sawtooth wave's magnitude.

Although it would seem accurate to consider some number of harmonics of the sawtooth wave and individually predict the loss associated with each harmonic, this would introduce excessive uncertainty into this correlation. The principal issue with this approach would be its implicit assumption that the principle of linear superposition applies to core loss prediction. Core losses are closely related to the patently nonlinear behavior of magnetic materials, and so the validity of invoking superposition would require its own investigation. As well, the frequency mismatch between the fundamental frequency and the characterization data would be substantially exacerbated at the harmonic frequencies.

In light of these issues, the correspondence between HFHF test currents and converter operating currents was established solely on the basis of the RMS value of the AC component of each current. The presence of a DC component in the converters' inductor current was not addressed. The RMS value of a sawtooth wave is $A_{pp}/\sqrt{12}$, where A_{pp} is the wave's peak-to-peak amplitude. The current probe, shown in Figure 121, was used to estimate the amplitude of the sawtooth wave.

In all cases, the compensated efficiency and output voltage are shown along with the actual observations. Compensating the efficiency results for the inductor losses was simply a matter of adding the estimated inductor losses to the output power before dividing by input power

$$\eta = \frac{P_{out}}{P_{in}} \tag{152}$$

$$\eta_C = \frac{P_{out} + P_{ind}}{P_{in}} \tag{153}$$

$$\tag{154}$$

where η is actual efficiency, η_C is compensated efficiency, P_{out} is output power, P_{in} is input power, and P_{ind} is estimated power loss in the inductor.

The voltage drop due to the inductor's resistance was calculated by referring the inductor's power loss to an equivalent resistance, based on the converter's output current:

$$v_{outC} = v_{out} + \frac{P_{ind}}{i_{out}^2} \tag{155}$$

where v_{out} and v_{outC} are the actual and compensated output voltages, P_{ind} is the estimated power loss in the inductor, and i_{out} is the converter's output current.

6.4.2 Output Loading

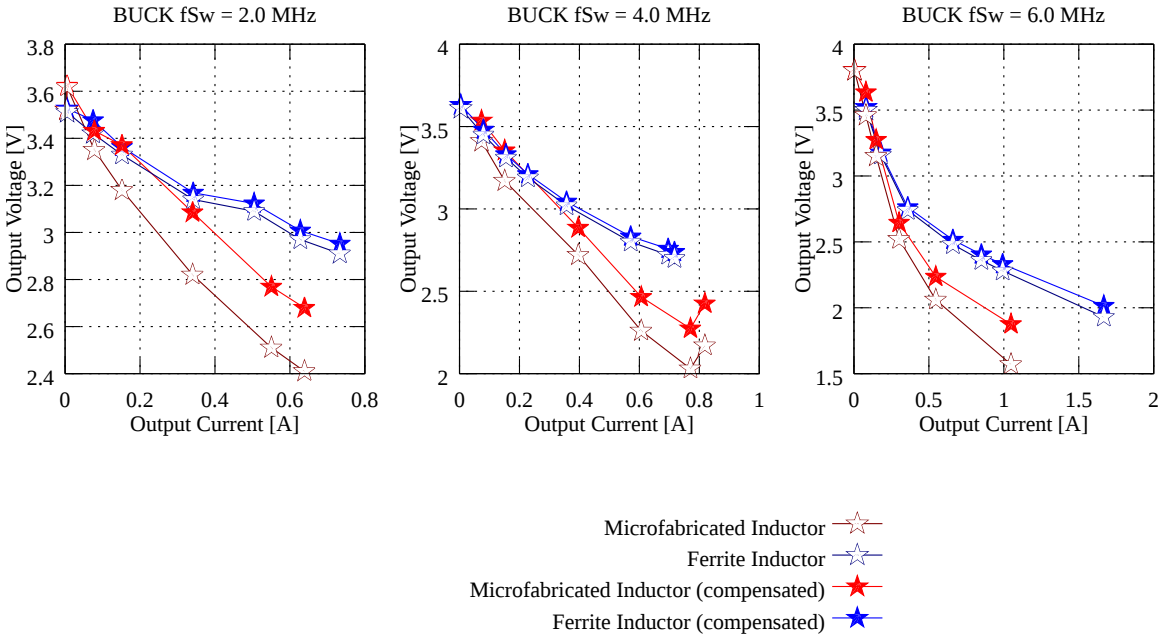

Figure 137: Buck converter load characteristic, with and without compensation for voltage drop due to inductor losses

The power converters were operated at constant duty cycle: 50 % for the buck converter and 75 % for the boost converter. This necessarily means that there was no closed-loop control for the output voltage. This allowed the output loading characteristic (output impedance behavior) of the converters to be more clearly examined.

Figure 137 and Figure 138 show the output loading curves for the buck and boost converters, respectively. While both the ferrite and laminated-core inductor showed similar output impedance for the boost converter case, there is a noticeable difference in the buck converter. There are two

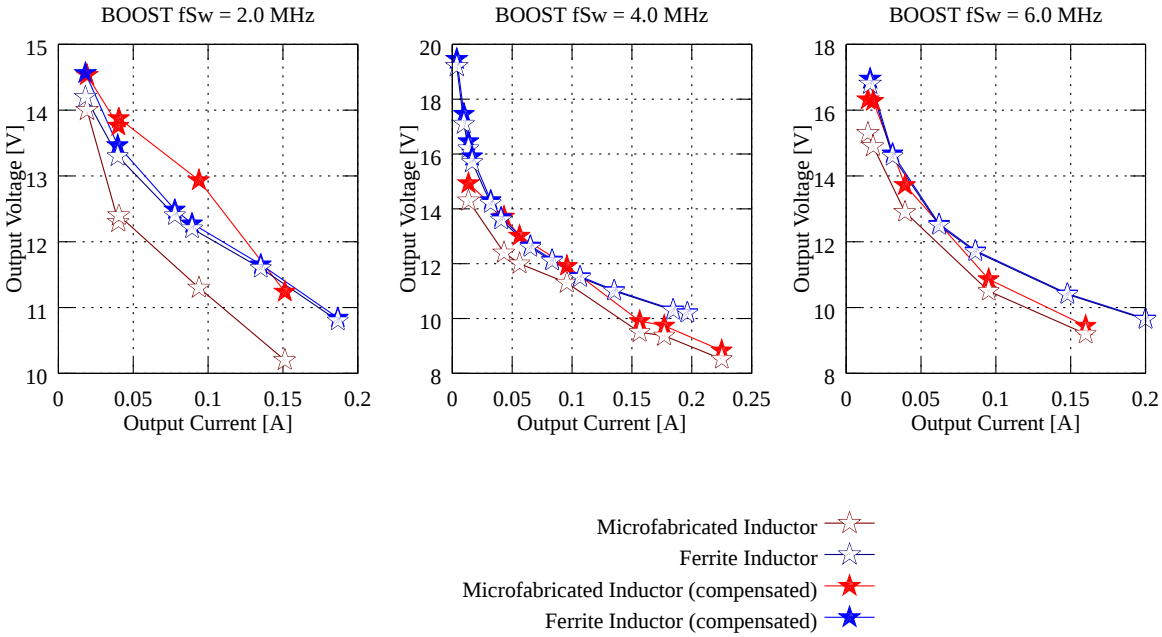

Figure 138: Boost converter load characteristic, with and without compensation for voltage drop due to inductor losses

causes for this difference. First, the laminated-core inductor has a significantly higher DC resistance than the ferrite inductor, 320 $\mu\Omega$ vs. 150 $\mu\Omega$. Second, the laminated-core inductor also has higher core losses than the ferrite inductor. This can be attributed to the significantly higher flux density in the laminated-core inductor, as well as to the same type of fabrication defects that would explain the higher-than-expected losses discussed in Chapter 5.

The fact that the difference in output impedance is more noticeable in the case of the buck converter can be explained by the fact that the boost converter apparently has significant intrinsic (*i.e.* common to both inductors) loss mechanisms, which contribute a dominant quantity of output losses.

6.4.3 Efficiency Results

The efficiency results for the buck and boost converters are shown, respectively, in Figure 139 and Figure 140. These results correlate strongly with the output loading results.

Note that a higher total output power was achieved with the ferrite inductor. This was due not only to its relatively high mass, and therefore high thermal mass, but also to its relatively low value. Since the laminated-core inductor represented a large amount of fabrication effort, and time

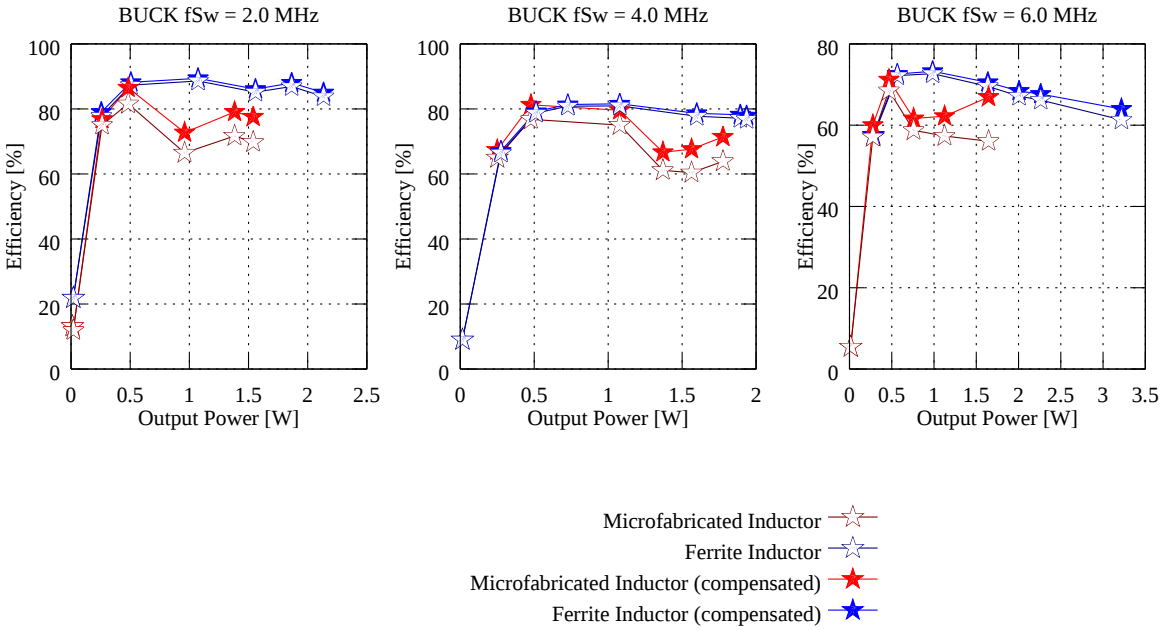

Figure 139: Efficiency of buck converter, before and after compensation with predicted inductor losses

to recreate, its temperature was more closely monitored than that of the ferrite inductor.

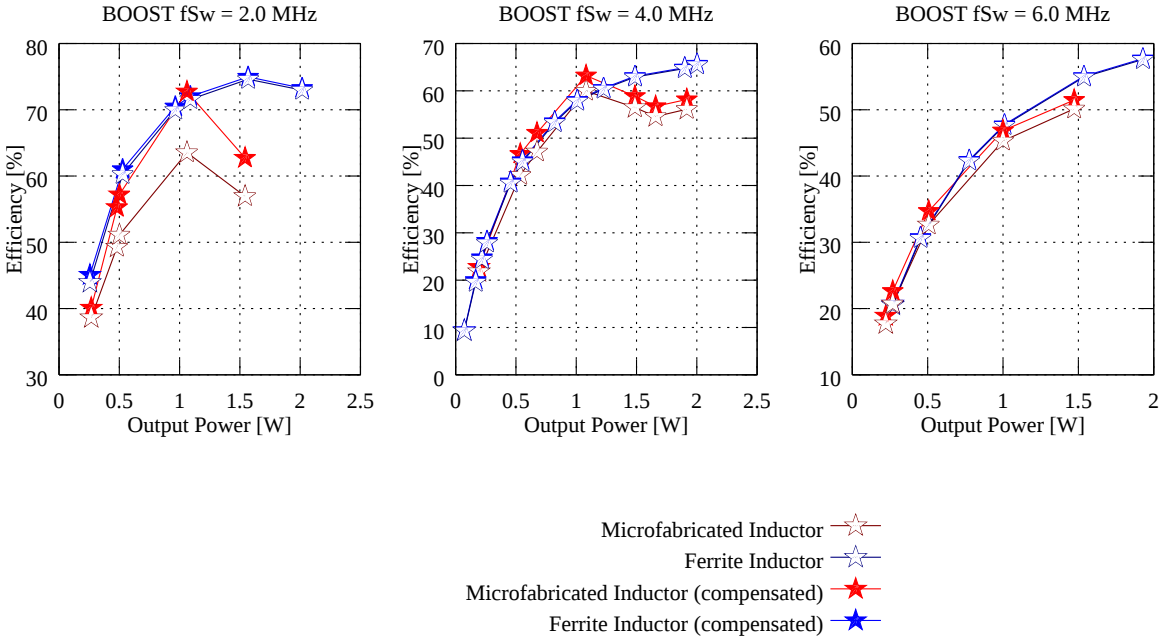

Figure 140: Efficiency of boost converter, before and after compensation with predicted inductor losses

Compensating the efficiency results for the predicted inductor losses accounts for a substantial portion of the performance difference between the inductors. However, the microfabricated inductor

shows an unexplained dip in efficiency around 1 W output, corresponding to roughly 300 mA output current. Surprisingly, increasing the output power increases the efficiency past this point.

This dip is postulated to be due to heating effects within the inductor. As the switching frequency increases, the dip occurs at lower output power. If this was a saturation effect, as switching frequency increases, the dip would occur at higher output power.

6.5 Conclusions

6.5.1 Performance Differences

The results in this chapter clearly establish the utility of the laminated cores in high frequency switching power converters. The ferrite inductor chosen did significantly outperform the microfabricated-core inductors, but this discrepancy should be viewed in the context of two major considerations.

6.5.1.1 Maturity of Technologies Used

First, ferrite inductor technology is extremely mature. Thus, the ferrite core itself can be expected to be made of highly-engineered materials with low losses. At the same time, the material constituting the microfabricated cores was grown in low-volume lab environments with low-cost, custom-made equipment.

As well, the packaging of the ferrite inductor is clearly superior to that of the laminated-core inductors. This results in lower winding resistance – both AC and DC resistances – as well as superior thermal performance. The reduced winding resistance directly improves converter efficiency, since the inductor is, in both cases, in series with the output current. The thermal performance is also very important because the performance of both ferrite and ferromagnetic materials deteriorates at higher temperatures. The ferrite inductor's superior ability to shed heat certainly helped the converters operate at higher efficiency.

6.5.1.2 Core Volume

Another, equally important, difference between the inductors used in this work is the core volume. The estimated core volume of the ferrite inductor, assuming a minimum ferrite feature size of 1mm, is approximately 40mm^3. Meanwhile, the core volume of the laminated-core sample is less than 7mm^3. The saturation current of the two inductors, although not precisely measured, are roughly

equivalent, around 1.5A.

The inductors presented in this thesis are compatible with most microfabrication and microelectronic packaging processes, and that is a major element of their value. As electronics systems continue to push towards ever-smaller packages with tighter integration, this processing compatibility is a key enabling trait. And it is along this direction that this cited reduction in core volume achieved with the move from ferrite magnetics to laminated-core magnetics compounds its impact.

6.5.2 Additional Advances in Power Converter Design

Although the demonstration circuits used in this chapter focused on high-frequency operation as the essential front along which power converter technology is progressing, there are many other directions in which advances have been made and are being made. These include quasi-resonant conversion techniques, advanced snubber design, multiphase architectures, and power factor correction. While none of these may offer a similarly direct path to improvement as seen with advanced core lamination, some of these design features may still present opportunities to apply sacrificial multilayer plating or other MEMS-based microfabrication techniques.

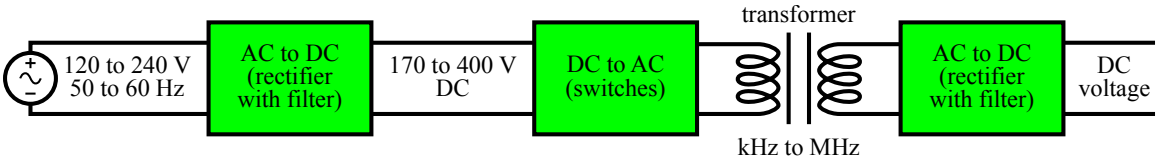

Figure 141: Typical Offline AC/DC Converter Architecture

One example might be found in power factor correction (PFC). PFC applies to any AC-connected load, and it refers to design features intended to make an load appear as much like an ideal resistor as possible to its AC input. Missing or ineffective PFC can cause inefficient or unpredictable behavior within power grids. For SMPSs, where the architecture is typically as shown in Figure 141, PFC would be applied in the first rectifier stage, and its goal would be to minimize the harmonic content of the SMPS's input current waveform.

This is typically achieved by implementing the first rectifier as a variable-conversion-ratio DC/DC converter, and modulating its conversion ratio such that it draws an input current that is proportional to the input voltage. When implemented in this manner, the first rectifier stage is referred to as the PFC converter. PFC converters are frequently boost or SEPIC topologies, and they provide an

opportunity to convert input power from a wide variety of grid voltages (*ex.* 120 V to 240 V, 50 Hz to 60 Hz) to an invariant format, typically around 400 V DC.

6.5.3 Higher Efficiency

While the demonstration of successful power conversion with these microfabricated cores is important, significantly more impact could be gained from demonstrating high-efficiency conversion, close to 90%, on par with typical commercial converters.

This work establishes that laminated-core inductors, even with extremely low core volumes, can support high-frequency switched-mode power conversion. The most straightforward path to such high-efficiency conversion would be to increase the inductance by including higher-volume cores. This would reduce the amplitude of the AC component of the inductor current, reducing its AC losses. Achieving more core volume could, naturally, be the result of stacking more cores together, or of plating thicker cores.

Another requirement to achieve higher efficiency in these converters would be to reduce the inductors' DC winding resistance. Particularly for low output voltage applications, this becomes very important. For instance, in the buck converter results presented in this work, where the output voltage was approximately 3V, the laminated-core inductors' DC winding resistance of 320 $\mu\Omega$ precluded any operating efficiencies of 90% or higher.

Lower DC winding resistance could be achieved either by including more strands of litz wire, or by using copper foil windings. In either case, the available volume for the windings would present a dominant limit.

6.5.4 Importance of Preserving Samples

As mentioned above, great care was taken to protect the integrity of the microfabricated cores and to prevent their overheating. During this research, a great many samples were lost to overheating during characterization and during operation in the power converters. It seems that, for future work, incorporating measures to monitor the core's temperature and initiate an automatic shutdown of the power converter would be of great value. This would not only prevent the loss of valuable samples, but also, by effectively reducing the required margins, allow the full operating envelope of the power converter to be explored.

6.5.5 Value of Tight Integration

One final aspect of these converters which could be improved to great advantage would be the integration of the inductor as well of as the entire system. This would give better performance, particularly by reducing the inductor's lead resistance and inductance. The latter – leads inductance – is of immense importance for the implementation of isolating topologies, since the self-inductance of the leads translates into leakage inductance, which directly limits the amount of energy that can be transferred into the core.

CHAPTER VII

CONCLUSIONS

This chapter draws together, across the scope of the entire thesis, key results and broad-based conclusions. As well, the principal intellectual contributions are highlighted.

7.1 Review of Major Results

Chapter 2 presented the fabrication of a compact inductor, microfabricated on a silicon substrate. The inductor was characterized both at the device level and in a high-speed switched-mode power supply (SMPS).

Sacrificial multilayer electroplating served as the central microfabrication technique for this work, and was the subject of numerous extensions and enhancements. These are detailed in Chapter 3, with particular attention devoted to improvements to the essential fabrication steps. Several process enhancements based on experiential learning are presented. Most substantial among these is the development of a second-generation automated multilayer electroplating system. Additional details of this system are presented in Appendix C.

Based on the sacrificial multilayer electroplating process, a new type of microfabricated capacitor was designed, fabricated, and characterized. This capacitor also served as to demonstrate the capacity of the process to produce high surface area structures. This work is presented in Chapter 4.

Also utilizing the sacrificial multilayer electroplating process, highly-laminated magnetic cores were fabricated. With lamination thicknesses ranging from 166 nm to 1.66 μm , these cores served as the means to demonstrate the use of the process to suppress eddy current losses in the produced structures. The design and fabrication of the cores, along with their integration and packaging into finished inductors is presented, in Chapter 5.

The successful suppression of eddy current losses is also detailed in Chapter 5, although the overall losses were found to be substantially higher than predicted by analytical modeling. The task of accurately measuring inductor core losses at high frequency and high magnetic flux entailed the design and implementation of a new core loss characterization system, which was based on

the principle of resonant loss measurement. As well, advanced numerical techniques were used to divide the observed core loss into eddy current, hysteresis, and anomalous components. Many of the fundamental principles underlying both the core loss measurement system and the analysis of the empirical data are detailed in Appendix B.

The final major result was the successful operation of the laminated-core inductors in high-frequency DC/DC converters, as described in Chapter 6. Both buck and boost topologies were implemented and comprehensively observed at switching frequencies from 2 MHz to 6 MHz and output loading of up to 2 W. Using a high-volume laminated-core inductor, outputs of 3 W and 5 W were achieved, respectively, with the buck and boost converters.

7.2 Intellectual Contributions

The contributions of this work comprise multiple microfabricated devices, the development of new equipment for the systematic construction of these devices, and the development of a methodology for the characterization and analysis of these devices' performance.

First, microfabrication techniques were gainfully applied to build a new type of power inductor in a fashion that is fully suitable for post-CMOS, wafer-level integration. Not only were aspects of the microfabrication processes used to introduce performance-enhancing features in the inductors, but it was also shown that the inductors offered excellent performance, both at the component level, and when operated in a prototype high-speed switching power supply.

Building on these results, a second type of power inductor was investigated. Although not a new type of inductor, these devices were built using electroplated magnetic cores with unprecedented degrees of lamination, reaching layer counts of 200, corresponding to a layer thicknesses of 166 nm. As well, the fabrication of these cores incorporated numerous advances in the unit fabrication processes, a new approach to mechanically supporting and electrically isolating the core laminations, and device-level packaging, all to create devices with good enough process yield and mechanical robustness to allow all subsequent characterization.

The premise behind implementing such extreme lamination geometries was the reduction of eddy current losses. Systematic and empirically conclusive experimentation to validate this premise was enabled by the development of two new, complex systems. First, an automated system for

mulitlayer electroplating was designed and implmented as a key infrastructure element. This system allowed reliable fabrication of multilayer devices, under very repeatable conditions, producing highly regular structures with well-controlled geometry. This increased process yield and accelerated the associated experiental learning, but, more importantly, the controllability and repeatability provided a plausible basis for comparing the observed behavior of these devices with analytical predictions.

A second fully custom system, for measuring the inductors' core losses at the frequencies and excitation levels found in advanced power converters, was also designed and implemented. This system stands as a highly evolved implementation of resonant core loss measurement, and a new way to measure losses in microfabricated cores. However, this system also made essential contributions to the successful confirmation of the suppression of eddy current losses in thinner laminations.

By incorporating numerous aspects of precise calibration and automation, it allowed the collection of a very large number of data points over a wide range of dynamic conditions. The precision and quantity of measurements obtained through the use of this system allowed the use of advanced numerical techniques, such as nonlinear least squares regression, to provide accurate and conclusive analyses of the core loss data.

The final results of these numerical analyses constitute the principal contribution of this work: a successful venture into core geometries with very thin lamination that delivered the expected benefit of suppressing eddy current losses. This result can also be broadened, without any loss of soundness or importance, to an ultimate contribution: that microfabrication processes may be borrowed from the MEMS field to create new types of power electronics devices with new levels of performance.

This is further strengthened by this work's auxiliary contribution of the design and fabrication of a new type of capacitor. The development and characterization of this capacitor illustrates the use of sacrificial multilayer plating, along with the complementary combination of unrelated microfabrication processes, to create a new device. Additionally, the investigation into the capacitor builds a basis for future research that, perhaps with improved packaging, could eventually yield a very high-performance, high-density power capacitor.

A final contribution of this work is the successful demonstration of the laminated-core inductors in very high-speed switching power supplies. This result establishes the relevance of the chosen

inductor core fabrication technology, including its capacity to limit eddy current losses, to practical power supply implementations. As well, it creates an avenue for the direct application of ferromagnetic electrodeposition technology in power magnetics devices, eventually in fully-integrated power conversion systems.

7.3 Potential Future Applications

Sacrificial multilayer plating has been applied to produce capacitors (Chapter 4) and high-performance power inductors (Chapter 5). However, there are numerous other applications in which this technique, particularly its ability to produce magnetic structures with high surface area, can deliver unique benefits. These may serve as fruitful areas of future investigation based on the techniques and results presented in this thesis.

7.3.1 Additional roles in SMPSs

The high-speed magnetic cores produced by sacrificial multilayer electroplating allow for SMPSs to operate at high speeds with reduced magnetic core volumes. As well, they might find expedient application in at least two other types of subcircuits commonly found in practical SMPSs.

7.3.1.1 High-side MOSFET Gate Drivers

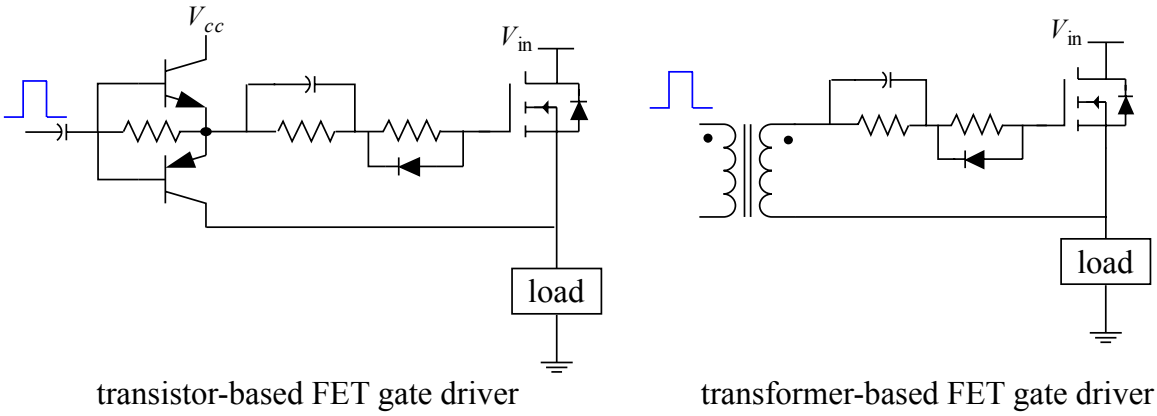

Figure 142: High-Side MOSFET Gate Drive Circuits [72]

MOSFETs are the most common type of switch element in silicon-based SMPSs. Though MOSFETs are available with both *n*- and *p*-doped channels, the former, also known as N-channel MOSFETs or just NFETs, are strongly preferred for use as power switches. [72]. This is due almost

entirely to the higher intrinsic mobility of electrons relative to holes, given the NFET's employment of electrons as primary charge carriers. The result is reduced conduction loss within the channel as well as faster switching, both of which directly improve the efficiency of the SMPS. A secondary advantage of using only NFETs in a SMPS is found in subcircuits where behavioral symmetry between two (or more) NFETs is important for performance, such as a typical power pole or a current mirror.

The chief difficulty with using only NFETs in SMPSs is the fact that an NFET's gate must be driven to a voltage higher than its source; in the case of a high-side NFET, shown in Figure 142, this means that the FET's gate must be driven higher than the power rail. There are a number of techniques for solving this problem, including the maintenance of a separate gate-drive rail (V_{CC} in Figure 142), often generated by a separate charge pump circuit. A capacitor provides the DC shift from the control circuitry's voltages to the high-side NFET's gate and source voltages, near V_{in}.

Another technique utilizes a pulse transformer to implement the DC voltage shift. Utilizing a transformer-based NFET drive circuit obviates the need for the separate gate drive V_{CC} power rail as well as the need for the high-voltage, high-speed gate drive transistors. The removal of these two elements substantially simplifies the SMPS's entire implementation.

One challenge with transformer-based NFET driver circuits is that the voltage waveforms required at the MOSFET's gate have very high slew rates, requiring that the gate drive transformer be capable of transferring energy at very high frequencies. Any excessive high-frequency losses in the gate drive transformer core will reduce, overall, the gate drive signal, but will especially attenuate the higher frequencies, leading to a lowered slew rate at the MOSFET gate. The fine laminations of magnetic cores produced with sacrificial multilayer electroplating would allow suppression of high-frequency eddy current losses, enabling the use of these cores in high-speed gate drive transformers. Moreover, the high saturation current of typical electrodeposited ferromagnetic alloys ensures that the gate drive transformers will be able to transfer the required amount of energy to the MOSFET gate while still maintaining a small size.

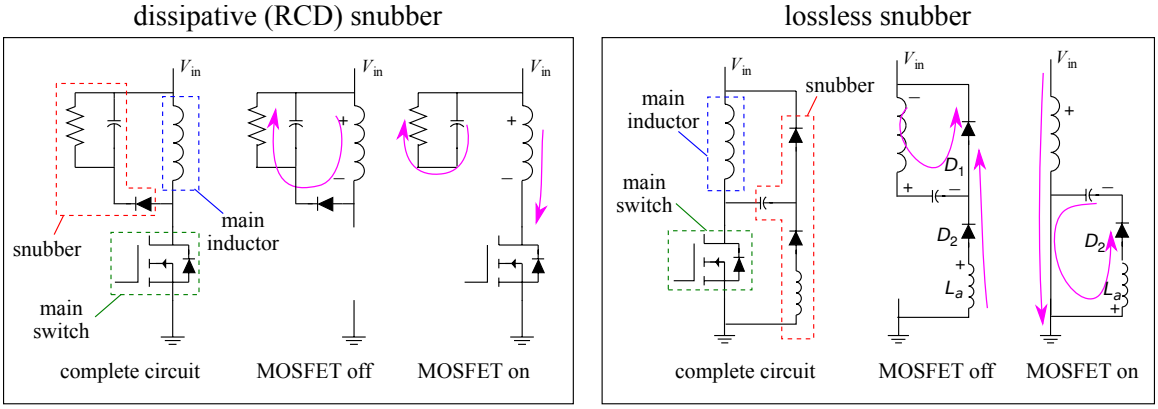

Figure 143: Lossless snubber circuit. Purple arrows indicate current flow. When the MOSFET turns off, the main inductor's flyback charges the snubber capacitor, while the snubber inductor discharges. When the MOSFET is on, the snubber capacitor discharges into the snubber inductor. [72]

7.3.1.2 Lossless Snubbers

Another common supporting circuit found in SMPSs is the snubber circuit. Snubber circuits, shown in Figure 143 are used to absorb the inductive flyback experienced when a power switch attempts to de-energize an inductive element. The voltage spike associated with inductive flyback is a major concern in practical designs, as it introduces increased voltage stress on the switch components as well as increased localized heat dissipation during the flyback. [72].

Snubbers necessarily dissipate some amount of power, essentially absorbing the high-frequency components of the inductive flyback with one or more resistors. However, overall converter efficiency can be improved by incorporating so-called "lossless" snubbers, which use energy storage elements, especially smaller inductors, to absorb the fast transient, and later return the absorbed energy to the input power rail. The two snubber types are illustrated in Figure 143. The high-speed, microfabricated magnetic cores produced with sacrificial multilayer electroplating might make excellent constituents of lossless snubbers.

7.3.2 Magnetic refrigeration

The magnetocaloric effect (MCE) describes a well-known behavior of magnetic materials, and is widely exploited to achieve refrigeration at temperatures near 0 Kelvin. MCE refrigeration is commonly used to cool superconducting materials to the very low temperatures at which superconduction occurs. Recently, however, materials and systems for magnetocaloric refigeration near

room-temperature, as a replacement for conventional refrigeration and cooling systems, are being researched. Room-temperature magnetocaloric refrigeration offers the possibility of higher efficiency and reduced environmental toxicity relative to organic refrigerant-based refrigerators. [33]

The magnetocaloric effect can be viewed as a consequence of the second law of thermodynamics, and is manifested any time a magnetic material undergoes a change in its magnetization state, particularly the change from fully saturated to fully randomized. The change in magnetic ordering amounts to a change in entropy, and by the second law of thermodynamics dictates a corresponding change in temperature under adiabatic conditions.

For a ferromagnetic material, the magnetocaloric effect is most pronounced near the material's Curie temperature T_C, at which the transition between ferromagnetism and paramagnetism occurs. This transition between these two magnetic regimes is termed a "magnetic phase transition," and bears many similarities to physical phase transitions, such as condensation or evaporation. Moreover, such magnetic phase transitions can be induced solely through modulation of an external magnetic field. [62]

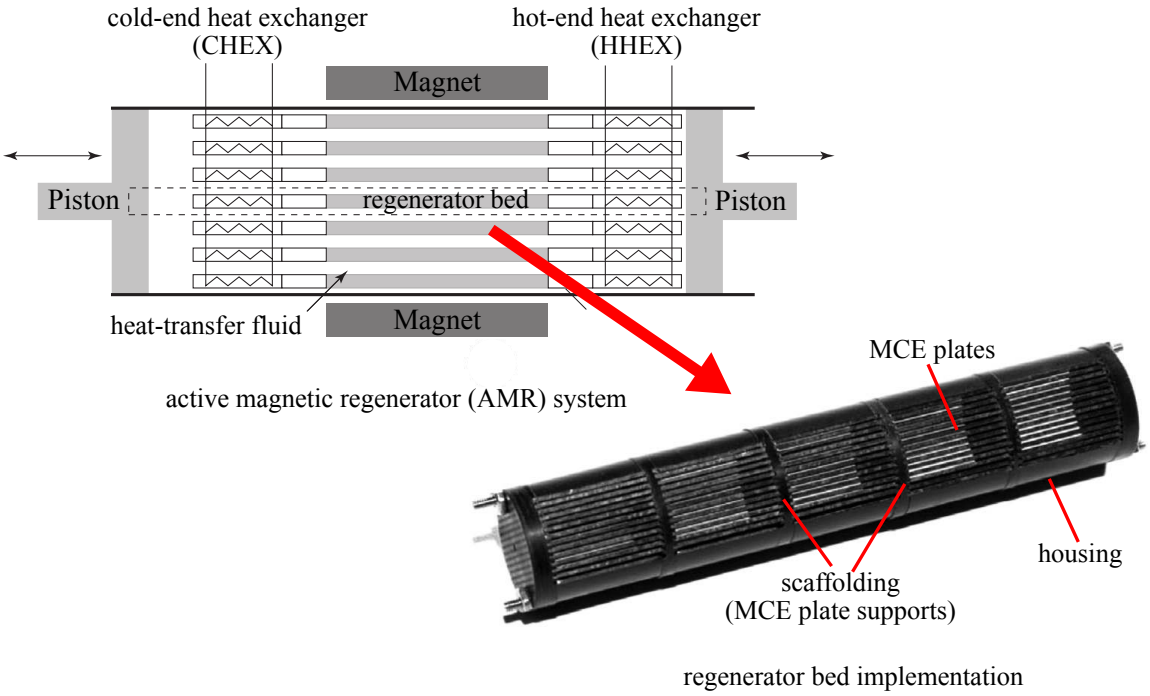

Figure 144: Magnetocaloric Refrigerator (AMR) and Regenerator Bed. after [65] and [48]

The magnetocaloric effect is harnessed to effect refrigeration by bringing a thermally conductive fluid into contact with some fully-magnetized material, and then subsequently removing the magnetizing field. As the magnetic material spontaneously loses its magnetic order, and its entropy increases, it will absorb heat from the fluid. The fluid, now at a lower temperature, is then removed from the magnetic material. Subsequent re-magnetization of the material will release this heat from the magnetic material. However, since the temperature change associated with the magnetic phase change is very low, less than 4 Kelvin in most cases, a regenerative configuration must be used. [2]

One example of a regenerator is shown in Figure 144, where two pistons force a heat-conduction fluid to undergo reciprocal motion through a closed tube, over a bed of magnetocaloric material. One end of the tube is designated as the hot end and the other the cold end; each end has a dedicated heat exchanger. An external magnet modulates the regenerator bed's magnetization in phase with the fluid's direction of flow so that the MCE-induced temperature changes tend to raise the temperature of the hot end lower the temperature of the cold end. This results in an end-to-end temperature difference greater than the temperature change achieved during any one pass of the fluid over the regenerator bed. [65] [48]

A typical regenerator bed is shown in Figure 144, in which thin plates of magnetocaloric material are suspended parallel to each other, so that the heat transfer fluid may flow between them. Alternative MCE regenerator bed constructions include pressed spheres and parallel rods. There is a wide range of potential materials for MCE refrigeration, including $La_{0.67}Ca_{0.33-x}Sr_xMn_{1.05}O_3$, and alloys such as Gd-Si-Ge. However, one of the major challenges in implementation of MCE coolers is the construction of the MCE regenerator bed themselves.

In the first place, the MCE elements must have high surface area, to allow maximal cooling capacity for a given volume. This concept is treated by the modeling of heat diffusion through the depth of an MCE element:

$$L_d = \left(\frac{k}{\pi \nu \rho C}\right)^{\frac{1}{2}} \tag{156}$$

L_d is the penetration depth of heat, ν is the regenerator's operating frequency, and k, ρ and C are the material's thermal conductivity, density and heat capacity, respectively. Given a maximum material depth, Equation 157 can be rewritten to give a maximum frequency at which the full volume of

magnetocaloric material experiences the temperature change:

$$\nu = \frac{k}{\pi L_d^2 \rho C} \tag{157}$$

The frequency at which the regenerative action can occur directly determines not only the cooling capacity of the system, but also the achievable end-to-end temperature difference. [48]

As well, the flow of the heat conduction fluid inside an AMR tends to induce mechanical damage in the MCE elements, Debris from this damage accelerates wear on seals, and presents a dominant limitation on the reliability of MCE cooling systems.

Multilayer structures produced by the microfabrication process explored in this chapter could offer distinct advantages if used to fabricate parallel-plate regenerator beds similar to the one shown in Figure 144. In the first place, arbitrary surface area to volume ratios could be achieved, allowing the selection of much higher operating frequencies than the typical values of 0.5-10 Hz [48] while still sweeping the entire volume of regenerator material through the full magnetic phase transition. As well, the interlayer gaps could be adjusted as desired to give fluidic channels of desired geometry and flow resistance, allowing precise control over the implicit tradeoff between magnetic volume, fluidic flow rate, and thermal frequency limits. Meanwhile, the electrodeposition process that deposits the magnetocaloric material could be modified to allow for the deposition of significantly tougher materials, such as nickel, in the interior regions of the magnetocaloric layers, imparting improved toughness and stiffness to the regenerator structure.

There are also very practical aspects that further augment the fit of sacrificial multilayer plated structures into the field of MCE cooling. First, as progress is made towards general-purpose, near-room-temperature magnetic refrigerators, strong interest in miniaturization will probably follow. Highly compact MCE coolers, given their potential for high density and efficiency combined with their lack of moving parts, may find application not only in electronics cooling but also at the lab-on-chip scale. Second, many of the materials of interest for MCE cooling, including those targeted for room-temperature use, have well-studied electrochemical behavior and are compatible with electrodeposition: erbium, yttrium, nickel, dysprosium, tin, cobalt, germanium, lead, silver, rhodium, gadolinium, and neodymium. This implies that, contrary to methods used to produce conventional MCE elements, one process – sequential multilayer electroplating – could give access

to a broad spectrum of materials. This is of particular interest for room-temperature development, where expected constructions require an array of MCE elements of varying composition, to present a monotonically-arranged set of magnetic phase transition temperatures. [62].

7.3.3 Magnetohydrodynamic pumps

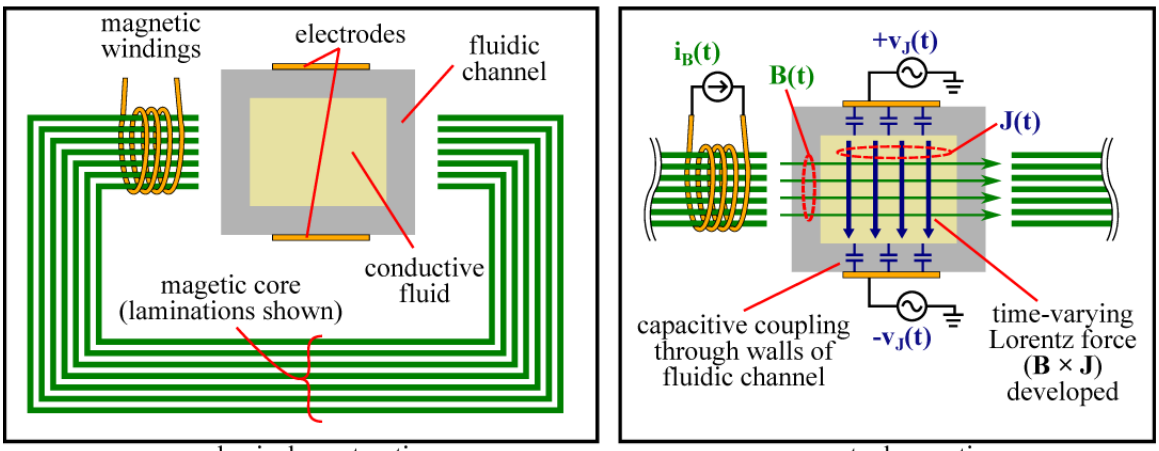

Figure 145: Schematic of high-speed magnetohydrodynamic pump. The laminated cores enable the use of high-frequency magnetic fields, while the use of high-frequency electrical drives will allow the use of capacitive coupling through the channel walls.

Magnetohydrodynamic (MHD) pumping is a technique for pumping liquids along closed channels that makes use of Lorentz forces. The construction and operation of an MHD pump based on a laminated magnetic core is illustrated in Figure 145. Voltage is applied across two electrodes, located along opposing boundaries of the fluidic channel, and this creates a transverse-directed electrical current. Note that this requires the fluid being pumped to have some degree of electrical conductivity. At the same time, a magnetic field is applied in a direction perpendicular to both the fluidic channel and the electrical current. The interaction of the magnetic field with the motion of the charge carriers within the fluid imparts a force on the charge carriers that is parallel to the fluidic channel.

MHD pumps are extremely attractive for microfluidic applications, particularly microscale total analysis systems (μ-TAS), as they are capable of creating flows without involving any moving parts. MHD pumps can employ either constant (DC) voltages and fields or time-varying (AC) voltages and fields to generate the necessary Lorentz forces. Each variety has advantages and disadvantages.

DC MHD pumps offer simpler operation than AC pumps, and allow the use of permanent magnets to generate the magnetic field. However, DC MHD pumps generally require that the electrodes be in direct contact with the fluid being pumped. This requires that careful attention be paid to the chemical compatibility between the electrode material and the fluid, and also creates the possibility of electrochemical interactions with the fluid's components. An example of the latter is the hydrolysis of water at 2.1 V, which imposes an identical limit on the maximum voltage that can be applied to aqueous solutions. This in turn either translates to a maximum pumping current achievable for a given fluid, or requires additional chemical species in the fluid that undergo reversible reduction-oxidation cycles at the electrodes, greatly reducing the apparent resistivity of the liquid. AC MHD pumps allow much greater voltages to be applied, as the pumped fluid's electrolysis reaction does not occur in appreciable amounts under AC conditions. As well, by creating a $\pi/2$ phase shift between the current and magnetic fields crossing the fluid, an AC MHD pump could potentially be employed as an electrically-controlled valve.

AC MHD pumps are widely reported, but have operated at frequencies no higher than 5 kHz. However, at substantially higher frequencies, the displacement currents (*i.e.* capacitive coupling) through a nonconducting fluidic wall should be large enough that most of the applied pump voltage is transferred across the fluidic wall. This would obviate the requirement that the MHD electrodes be in contact with the pumped fluid, a factor that could not only greatly simplify microfabrication of MHD pumps, but also, given the availability of highly-durable polymers such as PTFE, allow pumping of chemically aggressive fluids.

7.4 Final Conclusions

In alignment with the sets of conclusions presented at the end of each chapter, a set of broad conclusions may be drawn.

Most generally, it can be inferred from this work that the fabrication technique of sacrificial multilayer electroplating has many potential uses in the construction of power electronics systems. Successful application of the process to selected components – capacitors and inductors – has been shown, but the potential can be clearly seen for application to additional device types found in power electronics, such as transformers, batteries, and saturable-core reactors.

As well, the results presented suggest that there is an enormous benefit to be reaped from improved integration of these microfabricated power electronics components into power electronics systems. Tighter integration would further facilitate the operation of SMPSs at higher switching speeds, and enable reduced unit costs in a commercial production environment.

As well, a well-understood integration scheme would open an avenue of power electronics systems aimed at developing system architectures and system topologies that exploit the unique capabilities of fully-integrated systems. For instance, a high-performance fabrication platform allowing co-fabrication of inductors and capacitors, in batch fashion, might enable new techniques for employing quasi-resonant behavior in SMPSs, or perhaps new topologies for distributed power generation at the microscale.

In each portion of this work, there have been several immediately adjacent bases of knowledge that could provide substantial benefits if they could be leveraged in this research. Access to such related fields is an implicit premise in multidisciplinary fields such as MEMS. However, the breadth, depth, and maturity of the fields adajcent to this research, including commercial electroplating, fundamental electrochemistry knowledge, and silicon-based circuitry (*i.e.* analog CMOS design) seem remarkable. This suggests that efforts of any scale aimed at gathering and applying advances from these neighboring fields might be highly amplified in their impact.

Finally, a near absolute in this work has been the critical role played by microfabrication processes in enabling the devices and systems presented. This not only validates the research already performed in this area, but also clearly creates demand for continued and expanded work. This also implies that the discipline of microfabrication is likely to, in the future, be a significant source of high-impact advances in the traditionally system-oriented field of power electronics.

APPENDIX A

CONVENTIONS AND DERIVATIONS FOR EIKONAL GEOMETRIES

A.1 Background

Amongst the various projects presented in this thesis, a largely self-consistent group of conventions and derivations has evolved. These conventions deal with sets of closed paths, where the paths are eikonally interrelated. An eikonal relationship is one where there is a *constant minimum separation* between the two paths, as illustrated in Figure 146.

Eikonal curve sets model certain physical phenomena such as the advance of electromagnetic and acoustic wavefronts. In this work, they are used to model the advance of isotropic wet etches as well as magnetic flux distributions inside closed magnetic circuits.

These are well-explored topics within the field of mathematics, and this appendix serves only to develop ready-to-use expressions that are tightly tuned to the situations being considered. More rigorous and more complete treatment of these situations can be found in three disciplines:

- analytical geometry, in which paths may be described by a variety of curve constructions (*e.g.* as Bezier curves, B-splines, non-uniform rational B-splines (NURBS)), and robust algorithms for generating eikonal sets are known;
- image processing, in which the usage of dilation and erosion filters can be used to generate eikonal curves; and

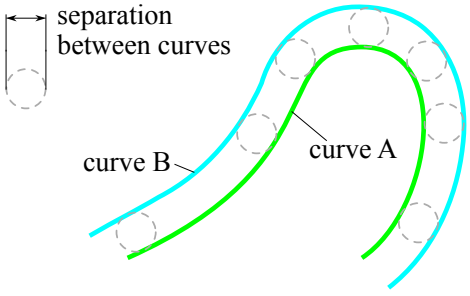

Figure 146: Eikonally-related curves are separated by a constant distance; in this case, curve A and curve B are eikonally related, and the constant minimum distance separating them is illustrated as the diameter of the circles

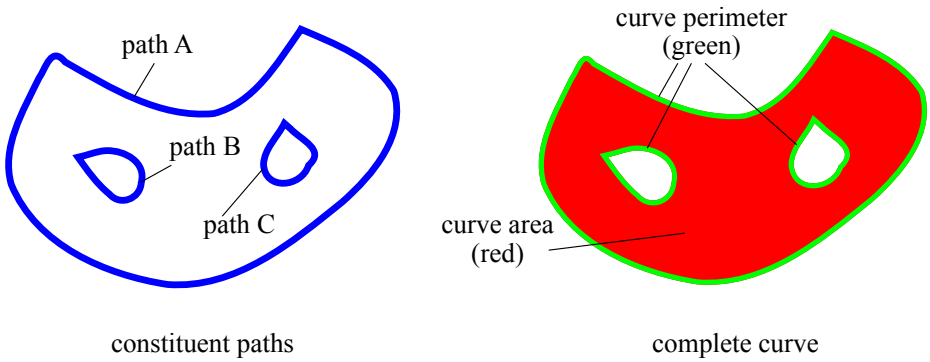

Figure 147: A curve (shown in green) consists of some geometric combination of one or more closed paths (shown in blue). In this case, the areas enclosed by the two inner paths are subtracted from the area enclosed by the outer path. The area (red) and perimeter (green) are defined only for the curve.

- differential geometry, in which differential equations in standard coordinate systems are developed that describe eikonal relationships.

A.2 Definitions

Curve is taken to mean a collection of one or more closed paths that, taken together, define a finite enclosed area. Where multiple closed paths are involved, the geometric operation used to combine their enclosed areas (*e.g.* subtraction, intersection, union), as well as the role of each path's area in that operation, will be clear from context. The area of the curve is defined as simply the area of the enclosed area, and the perimeter of the curve is the total perimeter of this enclosed area, including the perimeters of any interior regions. In general, a given curve could be specified by any number of different constituent path sets. However, since the curve's useful properties are related to the combination of the paths, and not the paths themselves, it is unimportant which set of paths is used to describe the curve.

In all cases, the sets of curves under consideration have a common structure. There is a single "anchor curve," along with some subset of all of its eikonally-related curves. The anchor curve is typically equal to the nominal lateral geometry of the microfabricated device under analysis. The remainder of the set of curves is spanned by a single continuous variable s, so that each value for s corresponds to a single curve. For a given curve, s also represents the eikonal offset (minimum distance) between that curve and the anchor curve. Therefore, $s = 0$ is taken as the anchor curve. Note that, in cases where the curve conists of multiple closed paths, the direction of the offset, as

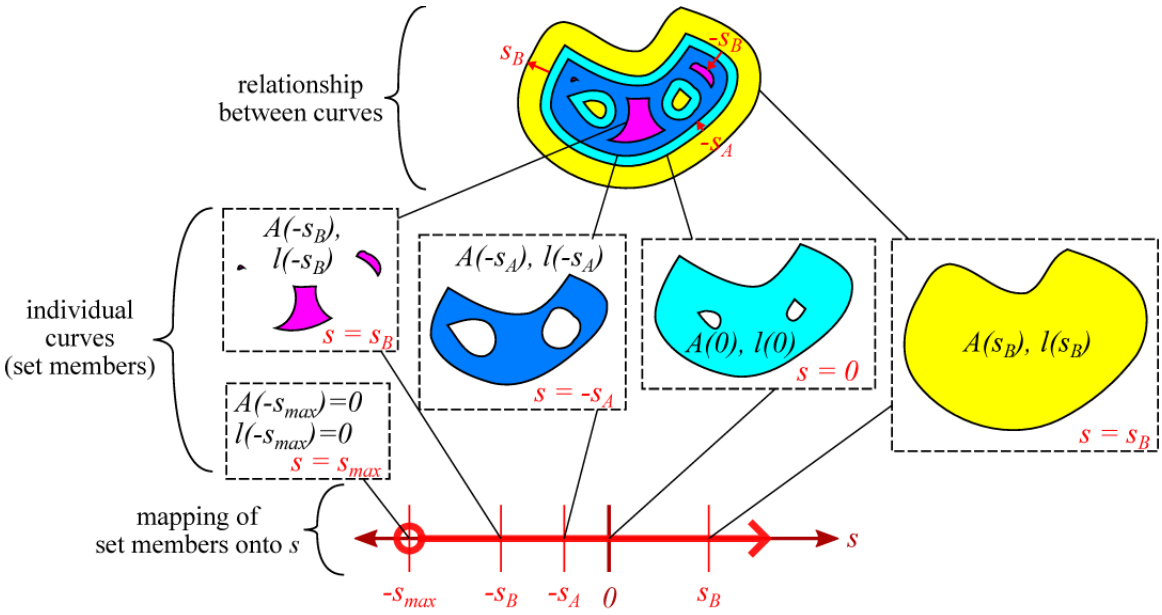

Figure 148: An example of an eikonal curve set, showing the use of the scalar variable s to span all members of the set.

applied to each path, will be apparent from context.

This structure is illustrated in Figure 148. For any given value of s, the area of the curve is written as $A(s)$ and the perimeter of the curve is written as $l(s)$. Note that for some particular value of s, the area and the perimeter of the curve will reach 0. In Figure 148, this value is shown as $-s_{max}$.

A.3 Perimeter and Area

In most applications considered, the perimeter and area of a given curve are the quantities of primary interest. Fortunately, the actual device geometries under consideration – that is, the anchor curves of the eikonal curve sets – have simple constructions. So, it is often feasible, for certain values or ranges of s, to construct closed-form, though *ad hoc*, expressions for both the area and perimeter of a curve.

A.3.1 General Behavior

Given an eikonal curve set, the curve area $A(s)$ will always vary continuously with s. Further, area is always differentiable with respect to s, and that derivative is precisely the perimeter, or

$$\frac{\partial A(s)}{\partial s} = l(s). \tag{158}$$

Equivalently, from an integrable expression for $l(s)$, the area can be defined as:

$$A(s) = A(s_0) + \int_{-s_0}^{s} l(\xi)\, d\xi \tag{159}$$

In contrast, the perimeter will not necessarily be either continuous or differentiable with respect to s.

A.3.2 Perimeter Cancellation Discontinuities

Curves with parallel or congruent edges are subject to sizeable perimeter cancellation discontinuities. For each value of s, any two paths or path segments bear exactly one of the following relationships:

- they are fully disjunct, and do not intersect at all,
- they intersect each other at a finite number of points, or
- they intersect at an infinite number of points (*i.e.* finite lengths of each path are fully coincident).

Perimeter cancellation discontinuities occur under the latter circumstance. Perimeter cancellation discontinuities are noteworthy because, in this work, their magnitude was frequently a signficant fraction of the total perimeter.

A.3.3 Quadratic Perimeter Approximation

Particularly for multi-path curves, the interaction between the constituent path's areas can become quite complex. However, to a coarse level of accuracy, curve area will grow as the square of s, while curve perimeter will grow as s. Formalizing this coarse approximation, as long as perimeter cancellation discontinuities are accounted for, can give usable approximations to otherwise intractable expressions for area and perimeter.

Employing this approximation amounts to choosing a set of values for s at which both $A(s)$ and $l(s)$ are known, assuming a quadratic form for $l(s)$ in between the known-value points. A linear approximation for $l(s)$

$$l_{lin}(s) = l(s_0)\frac{s_1 - s}{s_1 - s_0} + l(s_1)\frac{s - s_0}{s_1 - s_0} \tag{160}$$

would be simpler, and reasonably accurate, but would allow for a mismatch between the integral of the perimeter and the area:

$$\int_{s_0}^{s_1} l_{lin}(s)\,ds = \frac{[s_1 - s_0]\,[l(s_1) + l(s_0)]}{2} \neq A(s_1) - A(s_0) \tag{161}$$

Augmenting the assumed form for $l(s)$ with a degree of freedom about its midpoint, that is at $l(0.5s_0 + 0.5s_1)$, with symmetric quadratic support terminating at the interval endpoints, can exactly nullify this shortcoming of the linear approximation. The support function providing the extra degree of freedom, $\phi_{corr}(s)$ would be given as

$$\phi_{corr}(s) = \frac{(s - s_0)(s_1 - s)}{(s_1 - s_0)^2} \tag{162}$$

Determining the amplitude of the correction term K_{corr} is done by first observing that

$$\int_{s_0}^{s_1} \phi_{corr}(s)\,ds = \frac{s_1 - s_0}{6} \tag{163}$$

and introducing a temporary constant K_{corr} to ensure that the area change due to $\phi_{corr}(s)$ accommodate the discrepancy between the real area change $A(s_1) - A(s_0)$ and the area under the curve of $l_{lin}(s)$:

$$A(s_1) - A(s_0) = K_{corr}\int_{s_0}^{s_1} \phi_{corr}(s)\,ds + \int_{s_0}^{s_1} l_{lin}(s)\,ds \tag{164}$$

$$A(s_1) - A(s_0) = K_{corr}\frac{s_1 - s_0}{6} + \frac{(l_1 + l_0)(s_1 - s_0)}{2} \tag{165}$$

$$K_{corr} = \frac{6}{s_1 - s_0}\left[A(s_1) - A(s_0) - \frac{(l_1 + l_0)(s_1 - s_0)}{2}\right] \tag{166}$$

$$K_{corr} = 6\left[\frac{A(s_1) - A(s_0)}{s_1 - s_0} - \frac{l_1 + l_0}{2}\right] \tag{167}$$

Then, the assumed form for $l(s)$ over the open interval (s_0, s_1) can be written as the sum of a linear term $l_{lin}(s)$ and an appropriate scaled quadratic area-correction term $K_{corr}\phi_{corr}(s)$:

$$l(s) = l_{lin}(s) + K_{corr}\phi_{corr}(s) \tag{168}$$

$$= l(s_0)\frac{s_1 - s}{s_1 - s_0} + l(s_1)\frac{s - s_0}{s_1 - s_0} + 6\left[\frac{A(s_1) - A(s_0)}{s_1 - s_0} - \frac{(l_1 + l_0)}{2}\right]\left[\frac{(s - s_0)(s_1 - s)}{(s_1 - s_0)^2}\right] \tag{169}$$

$$= \begin{bmatrix} l(s_0) \\ 6\left[\frac{A(s_1)-A(s_0)}{s_1-s_0} - \frac{(l_1+l_0)}{2}\right] \\ l(s_1) \end{bmatrix} \cdot \begin{bmatrix} \frac{s_1-s}{s_1-s_0} \\ \frac{(s-s_0)(s_1-s)}{(s_1-s_0)^2} \\ \frac{s-s_0}{s_1-s_0} \end{bmatrix} \tag{170}$$

APPENDIX B

CORE LOSS MEASUREMENT

B.1 Motivation

Magnetic cores provide the indispensable capability to contain magnetic fields. But, this comes at the price of parasitic energy dissipation within the cores. These parasitic losses inside magnetic cores, or simply "core losses," can very easily become prohibitive to the desired operation of the magnetic component or system. Therefore it is important to understand these losses, both to be able to predict their impact on a system, and also to help guide efforts to minimize them.

B.2 Loss Types

Core losses can be attributed to a wide range of physical phenomena occurring within the magnetic material. The origins and basic analyses of the predominant loss types are presented in this section.

B.2.1 Eddy Current Losses

Eddy currents, and the associated losses, are part of a a parasitic effect that occurs with time-varying magnetic fields. They are accounted for in Maxwell's Equations. Faraday's Law

$$\oint E\, ds = -\frac{\partial \Phi}{\partial t} \tag{171}$$

dictates that an electromotive force will arise wherever there is a changing magnetic field. Ohm's Law

$$\vec{J} = \sigma \vec{E} \tag{172}$$

in turn dictates that in the presence of such electromotive force ($\vec{E}$) an electrical current will flow, and, moreover, that power will be dissipated as P_v:

$$P_v = \vec{J} \cdot \vec{E} = \sigma \left| \vec{E} \right|^2 \tag{173}$$

. Therefore, a magnetic core *will* experience nonzero eddy current losses *unless* it either has zero conductivity ($\sigma = 0$) or carries a constant magnetic field ($\frac{\partial \vec{B}}{\partial t} = 0$).

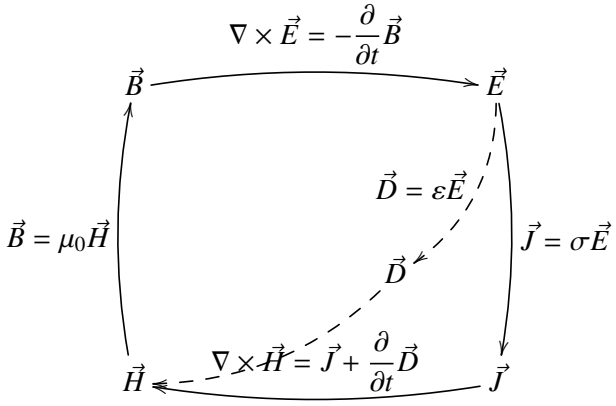

Figure 149: Transduction pathways involved in creating eddy currents, along with the associated analytical relations

The complete closed transduction loop that gives rise to eddy currents is shown in figure Figure 149. The electrical induction current pathway ($\vec{E} \rightarrow \vec{D} \rightarrow \vec{H}$) is shown for general completeness; it is relevant when electric displacement currents ($\frac{\partial}{\partial t}D$) are significantly large, but can be neglected otherwise. In the latter case, the governing relation can be developed:

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \tag{174}$$

$$\nabla \times \nabla \times \vec{E} = -\nabla \times \frac{\partial \vec{B}}{\partial t} \tag{175}$$

$$\nabla\left(\nabla \cdot \vec{E}\right) - \nabla^2\vec{E} = -\frac{\partial}{\partial t}\left[\nabla \times \left(\mu\vec{H}\right)\right] \tag{176}$$

$$\nabla^2\vec{E} = \mu\frac{\partial \vec{J}}{\partial t} \tag{177}$$

$$\nabla^2\vec{E} = \mu\sigma\frac{\partial \vec{E}}{\partial t} \tag{178}$$

$$\tag{179}$$

With the phasor convention for all fields and currents, this relation becomes

$$\nabla^2\vec{E} = \mu\sigma\omega\vec{E} \tag{180}$$

a relation that is formally equivalent to the Helmholtz Equation

$$\nabla^2 w + \lambda w = 0 \tag{181}$$

with $\lambda = -\mu\sigma\omega$ and $w = \vec{E}$. The characteristic polynomial for one-dimensional solutions to this

equation are found as

$$\left[\frac{\partial^2}{\partial x^2} - \lambda\right] w = 0 \tag{182}$$

$$D^2 - \lambda = 0 \tag{183}$$

$$D^2 = \lambda = \mu\sigma\omega \tag{184}$$

$$D = \pm\sqrt{\mu\sigma\omega} \tag{185}$$

giving a general solution of the form

$$w(x) = A\cosh\frac{x}{\delta} + B\sinh\frac{x}{\delta} \tag{186}$$

where A and B are free constants, to be determined by boundary conditions, and

$$\delta = \frac{1}{\sqrt{-\lambda}} = \sqrt{\frac{1}{\mu\sigma\omega}} \tag{187}$$

The quantity δ is known as *skin depth*, and and plays a pivotal role in analysis and modeling of eddy current effects. Further, it can be seen that for an invariant excitation or boundary condition, the total power lost to eddy currents increases as the square of frequency.

B.2.2 Hysteresis Losses

B.2.3 Anomalous Losses

Eddy current losses and hysteresis losses are both clearly predicted by theory and generally tractable in analysis. However, real magnetic systems and real magnetic materials invariably experience losses that are not accounted for by eddy current or *B*-*H* hysteresis models. These losses can be attributed to a number of physical phenomena, as well as to inadequacies of the standard eddy current and hysteresis models.

What are the mechanisms?

What if anomalous losses == losses not addressed?

Why is 1.5 such a popular exponent?

B.3 Analytical Modeling of Losses

B.3.1 Losses Within a Core Lamination

The differential relations presented in the previous section offer insight into the nature and behavior of eddy and hysteresis losses. Their utility can be further extended by considering their solutions

in typical magnetic core constructions. Fortunately, laminated cores possess a generally universal geometry: an orderly combination of identical laminations. The cross-section of each lamination – that is, the height and width of the lamination, measured perpendicular to magnetic flux – is constant within each lamination and uniform between laminations. Further, the height of the lamination is much less than either its width or length.

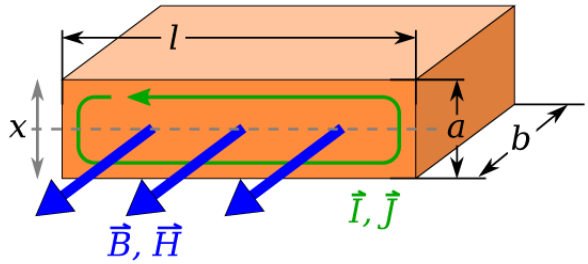

Figure 150: Single Lamination Geometry

These geometric assumptions lead to closed-form solutions to the eddy current equations, as derived in [4]. Figure 150 illustrates the assumed lamination geometry and the orientation of the associated quantities. The magnetic flux density $|B|$ is shown to be a function of height x and total flux Φ through the lamination:

$$|B(x)| = \frac{\sqrt{2}\,|\Phi|}{2l\delta}\sqrt{\frac{\sinh^2\frac{x}{\delta} + \cos^2\frac{x}{\delta}}{\sinh^2\frac{a}{2\delta} + \sin^2\frac{a}{2\delta}}} \tag{188}$$

where δ is the skin depth

$$\delta = \sqrt{\frac{2}{\omega\mu\sigma}} = \sqrt{\frac{1}{\pi f\mu\sigma}} \tag{189}$$

The current density $|J|$ is also found to be a function of height x:

$$|J(x)| = \frac{\omega\sigma\,|\Phi|}{2l\delta}\sqrt{\frac{\sinh^2\frac{x}{\delta} + \sin^2\frac{x}{\delta}}{\sinh^2\frac{a}{2\delta} + \sin^2\frac{a}{2\delta}}} \tag{190}$$

Utilizing these current and field distributions, the total eddy current losses P_e and hysteresis losses P_h within a given lamination can be calculated from the following integrals:

$$P_e = \int_V \frac{|J|^2}{2\sigma} dV \tag{191}$$

$$P_h = \int_V \left| \oint \vec{H} \cdot d\vec{B} \right| dV \tag{192}$$

giving the following expressions

$$P_e = \left(|\Phi|^2 \omega \frac{b}{l}\right) \frac{\sigma\delta}{8} \left(\frac{\sinh\frac{a}{d} - \sin\frac{a}{d}}{\cosh\frac{a}{d} - \cos\frac{a}{d}}\right) \tag{193}$$

$$P_h = \left(|\Phi|^2 \omega \frac{b}{l}\right) \frac{S}{2\pi\mu_0\mu_R\delta} \left(\frac{\sinh\frac{a}{d} + \sin\frac{a}{d}}{\cosh\frac{a}{d} - \cos\frac{a}{d}}\right) \tag{194}$$

B.3.2 Physical Interpretation

B.3.2.1 Volumetric Loss Expressions

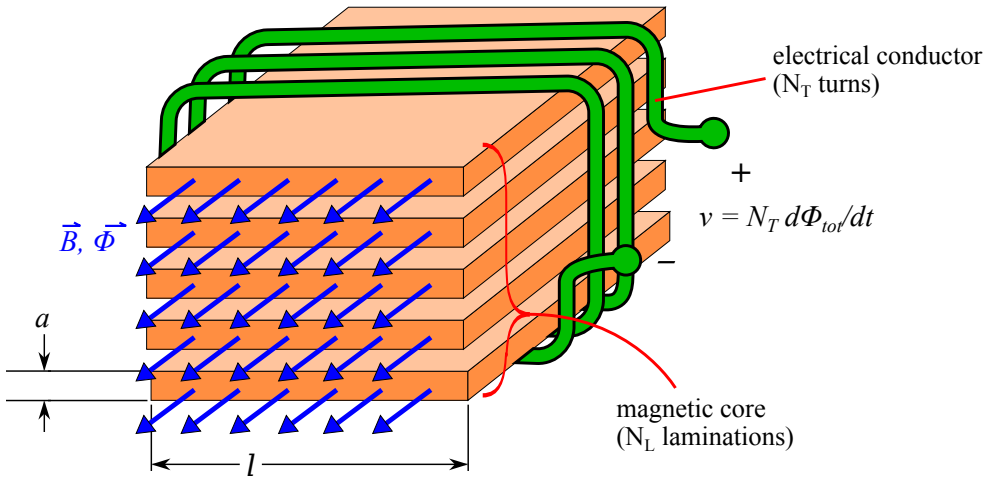

Figure 151: Geometric Loss Scaling Constants

Equation 193 and Equation 194 allow useful conclusions to be drawn regarding macroscopic behavior of core losses. Figure 151 shows the generic arrangment that should be considered in drawing these conclusions. In this arrangement, which matches the device implementations presented in Chapter 5, a set of N_L parallel laminations is linked by an electrical conductor with N_T virtually identical turns. As Φ_{tot}, the magnetic flux in the core changes, a voltage v is developed across the conductor.

$$v = N_T \frac{d}{dt}\Phi_{tot} \tag{195}$$

It is assumed that all of the core's laminations behave identically and that they therefore each carry

precisely the same flux Φ, giving the relations

$$\Phi_{tot} = N_L \Phi \tag{196}$$

$$v = N_T N_L \frac{d}{dt} \Phi \tag{197}$$

where Φ is the total flux through each lamination, as treated in Equation 193 and Equation 194. From observations of v, with the conductor's self-inductance neglected, Φ_{tot} can be estimated, either ballistically:

$$\Phi_{tot} = \frac{1}{N_T} \int v \frac{d}{dt} \tag{198}$$

or, for sinusoidal signals of frequency ω:

$$v = N_T \frac{d}{dt} \Phi_{tot} \tag{199}$$

$$|v| = N_T \omega \, |\Phi_{tot}| \tag{200}$$

$$|\Phi_{tot}| = \frac{1}{N_T} \frac{1}{\omega} \, |v| \tag{201}$$

This allows a single apparent average flux density B_a is defined such that

$$\Phi = \quad B_a ab \tag{202}$$

$$\Phi_{tot} = N_L B_a ab \tag{203}$$

$$B_a = \frac{\Phi_{tot}}{N_L ab} \tag{204}$$

Likewise, the total losses throughout the entire core, P_{e-tot} and P_{h-tot}, associated with eddy currents and hysteresis respectively, are given as:

$$P_{e-tot} = N_L P_e \tag{205}$$

$$P_{h-tot} = N_L P_h \tag{206}$$

where P_e and P_h are as treated in Equation 193 and Equation 194. Finally, the total core volume V and total cross-sectional area A are by

$$V = N_L abl \tag{207}$$

$$A = N_L ab \tag{208}$$

where a, b, and l are as shown in Figure 150.

From this perspective, the loss magnitudes P_e and P_h can be converted into respective volumetric quantities P_{Ve} and P_{Vh}:

$$\begin{bmatrix} P_{Ve} \\ P_{Vh} \end{bmatrix} = \frac{1}{V} \begin{bmatrix} P_{e-tot} \\ P_{h-tot} \end{bmatrix} = \frac{1}{N_L abl} \begin{bmatrix} N_L P_e \\ N_L P_h \end{bmatrix} = \frac{1}{abl} \begin{bmatrix} P_e \\ P_h \end{bmatrix} \tag{209}$$

Further pursuing the right-most expressions in Equation 209, while making use of Equation 202 reveals strong similarity between P_{Ve} and P_{Vh}:

$$P_{Ve} = \frac{1}{abl} P_e \tag{210}$$

$$= \frac{1}{abl} \frac{\pi}{4} \left(f \frac{2 |B_a|^2}{\mu_R \mu_0} \right) \left(\frac{a}{\delta} \cdot \frac{\sinh \frac{a}{\delta} - \sin \frac{a}{\delta}}{\cosh \frac{a}{\delta} - \cos \frac{a}{\delta}} \right) \tag{211}$$

$$P_{Vh} = \frac{1}{abl} P_h \tag{212}$$

$$= \frac{1}{abl} \frac{S}{2} \left(f \frac{2 |B_a|^2}{\mu_R \mu_0} \right) \left(\frac{a}{\delta} \cdot \frac{\sinh \frac{a}{\delta} + \sin \frac{a}{\delta}}{\cosh \frac{a}{\delta} - \cos \frac{a}{\delta}} \right) \tag{213}$$

By introducing the interpretive terms

$$P_{magV}(f, B) = f \frac{2 |B|^2}{\mu_R \mu_0} \tag{214}$$

$$G_{Ve}(a_S) = a_S \cdot \frac{\sinh a_S - \sin a_S}{\cosh a_S - \cos a_S} \tag{215}$$

$$G_{Vh}(a_S) = a_S \cdot \frac{\sinh a_S + \sin a_S}{\cosh a_S - \cos a_S} \tag{216}$$

the volumetric loss expressions become

$$P_{Ve} = \frac{\pi}{4} P_{magV}(f, B_a) G_{Ve} \left(\frac{a}{\delta} \right) \tag{217}$$

$$P_{Vh} = \frac{S}{2} P_{magV}(f, B_a) G_{Vh} \left(\frac{a}{\delta} \right) \tag{218}$$

These equations embody the entire dependence of core losses on the magnitude of the magnetization in the P_{magV} term. They also embody the entire dependence of losses on lamination thickness and on skin depth in the respective G_{Ve} and G_{Vh} terms. These terms are explored in more depth below.

B.3.2.2 Geometric Factors G_{Ve} and G_{Vh}

G_{Ve} and G_{Vh} in Equation 217 and Equation 218, respectively, can be seen as geometric scaling factors that capture the influence of lamination thickness on the magnitude of the associated losses.

By employing the Taylor series expansions

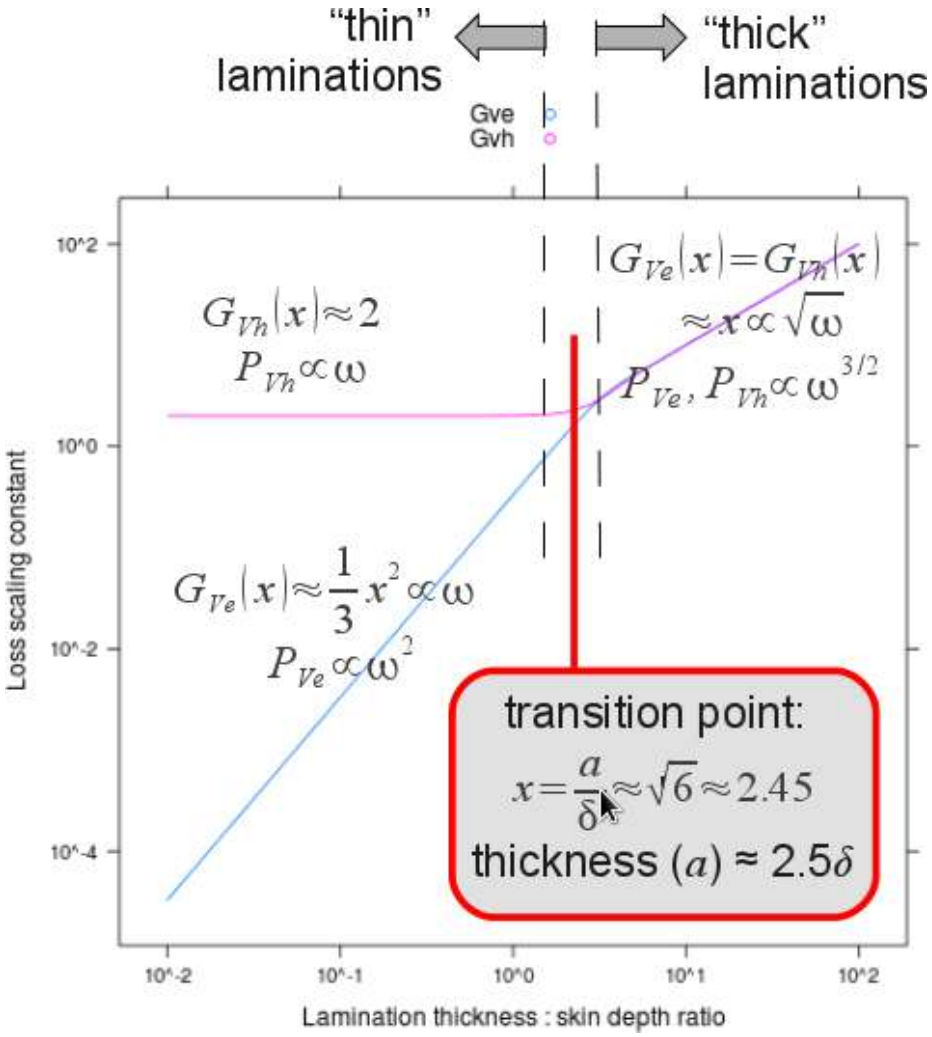

Figure 152: Geometric Loss Scaling Constants

$$\cos x = \frac{e^{jx} + e^{-jx}}{2} \approx 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \frac{x^6}{6!} + \frac{x^8}{8!} \quad \ldots \tag{219}$$

$$\cosh x = \frac{e^{x} + e^{-x}}{2} \approx 1 + \frac{x^2}{2!} + \frac{x^4}{4!} + \frac{x^6}{6!} + \frac{x^8}{8!} \quad \ldots \tag{220}$$

$$\sin x = \frac{e^{jx} - e^{-fx}}{2} \approx x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \frac{x^9}{9!} \quad \ldots \tag{221}$$

$$\sinh x = \frac{e^{x} - e^{-x}}{2} \approx x + \frac{x^3}{3!} + \frac{x^5}{5!} + \frac{x^7}{7!} + \frac{x^9}{9!} \quad \ldots \tag{222}$$

the expressions for each geometric scaling quantity become:

$$G_{Ve}(a_S) = a_S \frac{\frac{x^3}{3!} + \frac{x^7}{7!} \quad \cdots}{\frac{x^2}{2!} + \frac{x^6}{6!} \quad \cdots} \approx \begin{cases} \frac{1}{3}a_S^2 & a_S << 1 \\ a_S & a_S >> 1 \end{cases} \tag{223}$$

$$G_{Vh}(a_S) = a_S \frac{\frac{x^1}{1!} + \frac{x^5}{5!} \quad \cdots}{\frac{x^2}{2!} + \frac{x^6}{6!} \quad \cdots} \approx \begin{cases} 2 & a_S << 1 \\ a_S & a_S >> 1 \end{cases} \tag{224}$$

showing that two distinct regimes exist. These are illustrated in Figure 152. In thinner laminations, where the thickness a is less than the skin depth δ, a_S is less than one and G_{Ve} and G_{Vh} diverge from each other as the square of the specific lamination thickness. As the lamination thickness is increased substantially beyond the skin depth, such that a_S becomes much larger than one, G_{Ve} and G_{Vh} become equivalent. The loss factors G_{Ve} and G_{Vh} reach close parity when a_S is $\sqrt{6}$, or approximately 2.45.

B.3.2.3 Power of Magnetization P_{magV}

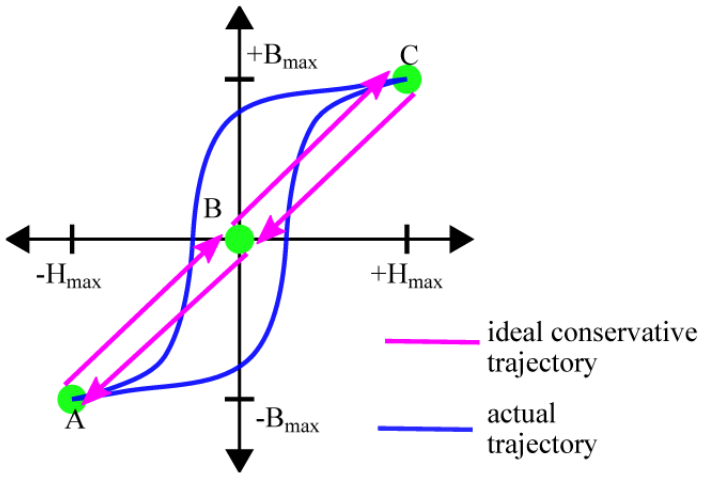

Figure 153: Interpretation of P_{magV}

Figure 153 shows a typical $B-H$ curve for a ferromagnetic material, in blue, as well as the direct path, in purple, connecting the curve extrema A and C. Note that the point B is an inferred waypoint along the direct path, and its location at the origin is a consequence of the symmetry of the $B-H$ curve about the origin. The point B divides the direct $B-H$ path into four equally-sized segments along which the magnetic stored energy W_v

$$W_v = \frac{1}{2}BH \tag{225}$$

varies monotically. If the magnitudes of the change in energy of each segment of the direct $B-H$ path are summed, and this energy sum is multiplied by the operating frequency f, the result is $P_{magV}(B_{max})$.

$$f\sum|\Delta W_v| = f\sum\left|\frac{1}{2}B_eH_e - \frac{1}{2}B_bH_b\right| \tag{226}$$

$$= f\frac{1}{2}\{|B_{max}H_{max}| + |-B_{max}H_{max}| + |B_{max}H_{max}| + |-B_{max}H_{max}|\} \tag{227}$$

$$= 2fB_{max}H_{max} \tag{228}$$

$$= 2fB_{max}\left(\frac{H_{max}}{\mu_R\mu_0}\right) \tag{229}$$

$$= f\frac{2B_{max}^2}{\mu_R\mu_0} \tag{230}$$

$$= P_{magV}(B_{max}) \tag{231}$$

Thus, the apparent average flux B_a from Equation 218 and Equation 217 is equivalent to B_{max} in Figure 153.

From this equivalence between P_{magV} and the absolute-value summation of all magnetic energy changes per unit time, it can be inferred that the total losses over a given volume are similarly proportional to the absolute-value summation of the changes in magnetic energy over the entire volume.

$$P_{magV} = \frac{P_{Ve}}{\frac{\pi}{4}G_{Ve}} \tag{232}$$

$$\int P_{magV}\,dV = \int\left[\frac{P_{Ve}}{\frac{\pi}{4}G_{Ve}}\right]dV = \frac{\int P_{Ve}\,dV}{\frac{\pi}{4}G_{Ve}} \tag{233}$$

$$= \frac{P_e}{\frac{\pi}{4}G_{Ve}} \tag{234}$$

$$P_{magV} = \frac{P_{Ve}}{\frac{pi}{4}G_{Ve}} = \frac{P_{Vh}}{\frac{S}{2}G_{Vh}} \tag{235}$$

$$P_{mag} = \frac{P_e}{\frac{S}{2}G_{Vh}} \tag{236}$$

$$= \frac{P_e}{\frac{S}{2}G_{Vh}} \tag{237}$$

$$= \frac{P_h}{\frac{S}{2}G_{Vh}} \tag{238}$$

Integrating the magnetic energy storage W_v over an inductor's magnetic core

In practical usage, it is difficult to directly measure the magnetic flux $|B|$, and this physical interpretation of P_{magV} as the length of the direct $B-H$ path is utilized to justify an alternative expression for P_{magV}. Considering an inductor with measured inductance L and sinusoidal current of magnitude I_{pk}, the total flux Φ in the inductor core can be expressed as

$$|\Phi| = \frac{LI_{pk}}{N_t} \tag{239}$$

where N_t is the number of turns in the inductor's windings. This same Φ, however, is assumed to be equal to the flux, summed across the entire cross-section of all of the core's laminations

$$|\Phi| = N_L ab\,|B| \tag{240}$$

Assuming that the mechanisms modeled in Equation 217 and Equation 218 act the same throughout the core volume, summing those across laminations gives loss figures for the total core volume – P_{Te} and P_{Th}, respectively –

$$P_{Te} = (N_L abl)\,P_{Ve} \tag{241}$$

$$= \frac{\pi}{4} P_{magT}\,(f,B)\,G_{Ve}\left(\frac{a}{\delta}\right) \tag{242}$$

$$P_{Th} = (N_L abl)\,P_{Vh} \tag{243}$$

$$= \frac{\pi}{4} P_{magT}\,(f,B)\,G_{Vh}\left(\frac{a}{\delta}\right) \tag{244}$$

with the use of a totalized form of P_{mag}, P_{magT}

$$P_{magT} = N_L abl P_{mag} \tag{245}$$

Finally, we arrive at

$$P_{magT} = f\frac{2LI_{pk}^2}{N_t} \tag{246}$$

$$= f\frac{LI_{rms}^2}{N_t} \tag{247}$$

Dividing by the core volume V, an experimental figure for P_{mag} can be achieved

$$P_{mag} = \frac{1}{N_L abl} f\frac{\Phi_{tot} I_{rms}}{N_t} \tag{248}$$

$$= \frac{1}{N_L abl} f\frac{LI_{rms}^2}{N_t} \tag{249}$$

$$P_{mag} = \frac{1}{V} f \frac{L I_{rms}^2}{N_t} \tag{250}$$

B.3.3 Hysteresis Loop Coefficient S

The hysteresis losses are seen to depend directly on a coefficient S. S is illustrated in Figure 154,

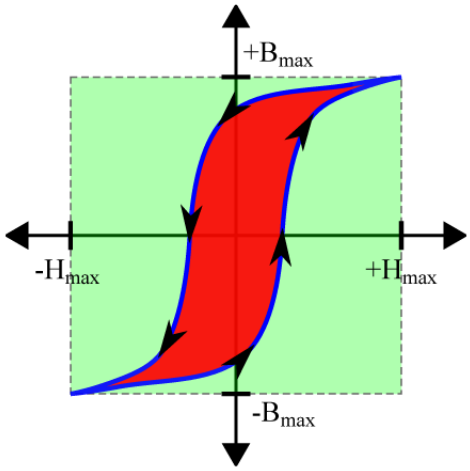

Figure 154: Interpretation of S

as the ratio of the interior area of the $B-H$ curve to the area of the rectangle containing the two extrema of the $B-H$ curve.

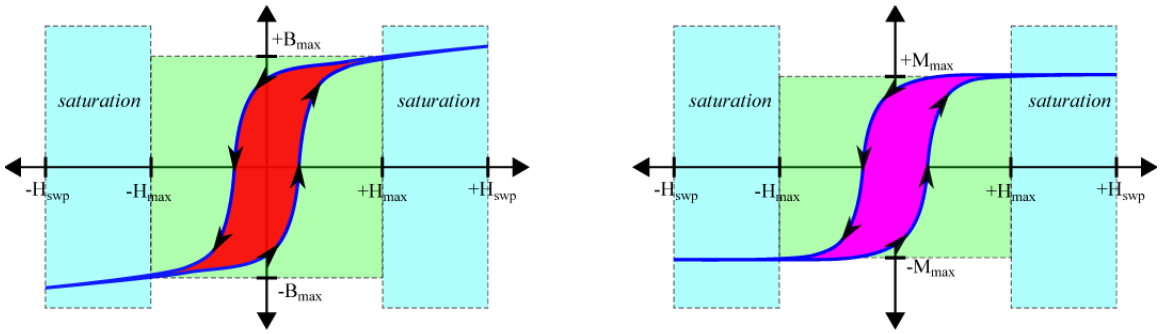

B vs H trajectory M vs H trajectory

Figure 155: Observation and Discrimination of S from VSM Trace

A vibrating sample magnetometer (VSM) was used to estimate S for loss estimation. Figure 155 shows typical VSM traces, and the orientation of the loop area S within these traces. Note that typical VSM traces include H fields well in excess of saturation field densities, so real VSM traces will include regimes, labeled as "saturation," that will require manual exclusion from calculations. Although the $B-H$ loop contains minimal area in these saturation regimes, correct determination

of S requires correct determination of H_{max}, which is identical to manually identifying saturation.

Also, although the VSM only directly measures magnetization M, and not flux density B, the $M - H$ curve can be directly used for determination of S. This is because the conversion from M to B involves the addition of an invariant quantity $\mu_0 H$ to each segment of the $M - H$ curve, and the subtraction implicit in area measurement will effectively nullify this addition. Consequently, calculatng the area of the loop as $\int B\, dH$ will give precisely the same result as calculating the are as $\int M\, dH$.

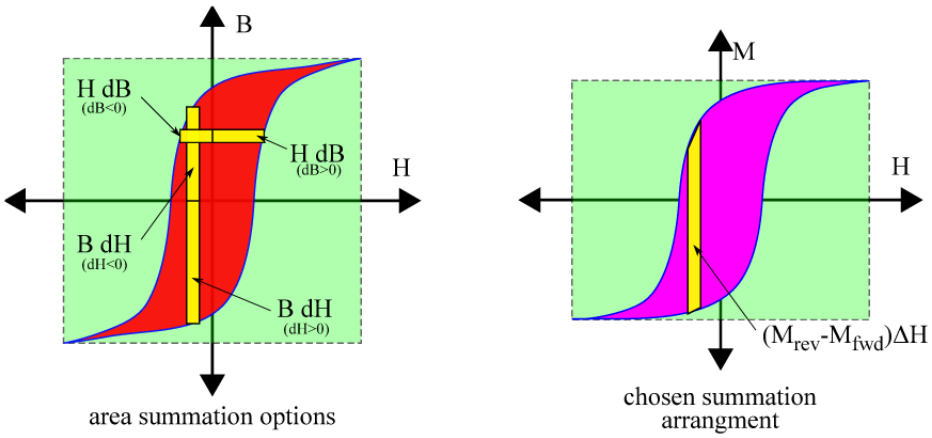

Figure 156: Numerical Determination of S

The numerical addition approach is shown in Figure 156, along with an illustration of possble approaches. iIt was chosen to sum ΔM over H.

B.4 Loss Mechanism Separation

B.4.1 Statement of Problem

There are a number of physical phenomema that create losses within an inductor:

- winding conduction losses,
- eddy current losses within the core,
- hysteresis losses within the core, and
- anomalous losses within the core,

as well as other potential, less easily modeled phenomena such as parasitic lamination actuation and the magnetocaloric effect. However, all core loss measurement techniques discussed in this thesis,

including the balanced-reactivity approach used in the experimental characterization, only reveal the total losses experienced by the DUT. Although an accurate measure of the winding current enables a faithful estimation of the losses due to winding impedance, loss measurement techniques offer little, if any, insight into the distribution of the remaining device losses amonst phenomena occuring internal to the core.

Constructing a plausible attribution of DUT losses to the constituent loss phenomena is the essential definition of the loss separation problem.

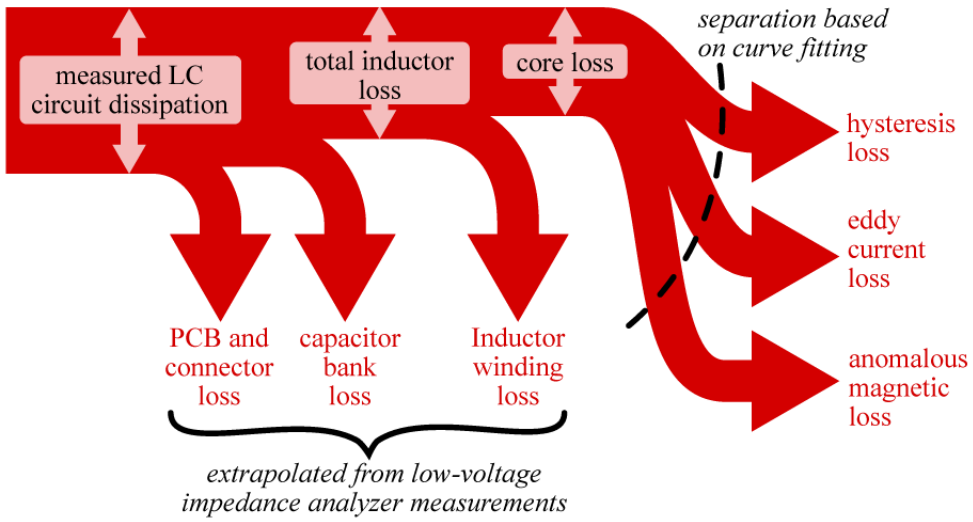

Figure 157: Assumed Power Flow for Inductor Characterization

The assumed decomposition of power losses is shown in Figure 157. Note that this model is tailored to the HFHF loss measurement system described in Chapter 5, with explicit allocation of losses to PCB, connector, and capacitor bank losses. With these losses compensated out, the remaining losses may be attributed to the inductor alone, comprising the winding and the core. Within the inductor, the winding losses are then removed, based on impedance measurements taken at low currents. Finally, only the core losses remain, at which point curve fitting must be used to further divide the losses. This final step is the focus of the remainder of this section.

B.4.2 Simple Linear Regression

B.4.2.1 Mathematical Theory

The problem of simple linear regression can be stated as fitting a set of datapoints $\{(x_i, y_i)\}$, where $i \in \mathbb{N}$, to a relation of the form

$$y(x) = mx + b \tag{251}$$

The process of fitting is taken to mean, qualitatively, minimizing the mismatch, across all points x_i, between the so-called "predicted" values $y(x_i)$ and the "actual" values y_i. This mismatch is formalized as the mean squared error E

$$E = \frac{1}{T}\sum_{i=1}^{T}\left[y_i - y(x_i)\right]^2 \tag{252}$$

where T is the number of datapoints in the set $\{(x_i, y_i)\}$.

Identifying the best fit, then, becomes mathematically defined as choosing the values for m and b that minimize the mean squared error. Elementary calculus establishes that, without limits on the domains of m and b, this minimum will occur at a point where one or both of the partial first derivatives of E with respect to m and b are equal to zero.

$$\frac{\partial E}{\partial m} = 0 \tag{253}$$

$$\frac{\partial E}{\partial b} = 0 \tag{254}$$

Expanding the first equation gives

$$\frac{\partial E}{\partial m} = \frac{\partial}{\partial m}\left[\frac{1}{T}\sum_{i=1}^{T}\left[y_i - y(x_i)\right]^2\right] \tag{255}$$

$$= \frac{1}{T}\sum_{i=1}^{T}\frac{\partial}{\partial m}\left[\left[y_i - y(x_i)\right]^2\right] \tag{256}$$

$$= \frac{2}{T}\sum_{i=1}^{T}\left[y(x_i) - y_i\right]\frac{\partial y(x_i)}{\partial m} \tag{257}$$

$$= \frac{2}{T}\left\{\sum_{i=1}^{T} y(x_i)\frac{\partial y(x_i)}{\partial m} - \sum_{i=1}^{T} y_i\frac{\partial y(x_i)}{\partial m}\right\} \tag{258}$$

$$= \frac{2}{T}\left\{\sum_{i=1}^{T} mx_i\frac{\partial y(x_i)}{\partial m} + \sum_{i=1}^{T} b\frac{\partial y(x_i)}{\partial m} - \sum_{i=1}^{T} y_i\frac{\partial y(x_i)}{\partial m}\right\} \tag{259}$$

$$= \frac{2}{T}\left\{\sum_{i=1}^{T} mx_i^2 + \sum_{i=1}^{T} bx_i - \sum_{i=1}^{T} y_i x_i\right\} \tag{260}$$

A similar derivation gives, for the second equation,

$$\frac{\partial E}{\partial b} = \frac{2}{T}\sum_{i=1}^{T}\left[y_i - y\left(x_i\right)\right]\frac{\partial y\left(x_i\right)}{\partial b} \tag{261}$$

$$= \frac{2}{T}\left\{\sum_{i=1}^{T} m x_i \frac{\partial y\left(x_i\right)}{\partial b} + \sum_{i=1}^{T} b\frac{\partial y\left(x_i\right)}{\partial b} - \sum_{i=1}^{T} y_i \frac{\partial y\left(x_i\right)}{\partial b}\right\} \tag{262}$$

$$= \frac{2}{T}\left\{\sum_{i=1}^{T} m x_i + \sum_{i=1}^{T} b - \sum_{i=1}^{T} y_i\right\} \tag{263}$$

Setting the two first derivatives to zero gives two equations for the two unknowns m and b:

$$m\sum_{i=1}^{T} x_i + b\sum_{i=1}^{T} 1 = \sum_{i=1}^{T} y_i \tag{264}$$

$$m\sum_{i=1}^{T} x_i^2 + b\sum_{i=1}^{T} x_i = \sum_{i=1}^{T} y_i x_i \tag{265}$$

which can be put into matrix form:

$$\begin{bmatrix} \sum_{i=1}^{T} x_i & T \\ \sum_{i=1}^{T} x_i^2 & \sum_{i=1}^{T} x_i \end{bmatrix}\begin{bmatrix} m \\ b \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^{T} y_i \\ \sum_{i=1}^{T} x_i y_i \end{bmatrix} \tag{266}$$

Having found a unique zero of the first derivatives of E, the only remaining task, in order to ensure that this zero is in fact a global minimum for E, is to establish that the second derivatives with respect to both parameters m and b are positive. It can be seen by inspection that

$$\frac{\partial^2 E}{\partial m^2} = \frac{2}{T}\sum_{i=1}^{T} x_i^2 \geq 0 \tag{267}$$

$$\frac{\partial^2 E}{\partial b^2} = \frac{2}{T}T = 2 \geq 0 \tag{268}$$

$$\tag{269}$$

B.4.2.2 Application

The application of simple linear regression is quite straightforward. Closed-form expressions for m and b are easily attainable.

Assessing the suitability of a linear model for a given data set is usually done by inspecting a plot of y_i vs x_i. Numerical indicators of the quality of the linear fit are also available. Certainly, the residual error figure itself, E, is useful. Clearly, even though E is expressed on a per-data-point basis, its magnitude must be compared relative to those of the data points themselves.

Another useful numerical measure for quality of fit is the coefficient of determination, usually written as R^2. The R^2 indicator takes into account both the mean squared error and the magnitude of the y_i, as well as the intrinsic variance of the y_i. Two terms are introduced to measure the variance within the data:

$$SS_{tot} = \sum_{i=1}^{T} [y_i - \bar{y}_i]^2 \tag{270}$$

$$SS_{err} = \sum_{i=1}^{T} [y_i - y(x_i)]^2 = T \cdot E \tag{271}$$

where $\bar{y}_i$ is the mean of the y_i. SS_{tot} measures the intrinsic variance of the y_i, while SS_{err} measures the residual variance of y_i relative to the linear fit. R^2 is a measure of how much the fit reduces (or "explains") the variance; presumably, an ideal model would fit the data perfectly, and the variance SS_{err} would be zero.

$$R^2 = \frac{SS_{tot} - SS_{err}}{SS_{tot}} \tag{272}$$

$$= 1 - \frac{SS_{err}}{SS_{tot}} \tag{273}$$

R^2 values close to 1 indicate good fits.

It is also possible to use linear regression for situations where nonlinear models are expected. This can be accomplished if the nonlinear model can be transformed to a linear eqation. For example, consider a data set that is expected to fit the Nernst equation for electrochemical half-cell potentials

$$E \approx E_0 - \frac{0.05916\text{V}}{z} \log_{10} Q \tag{274}$$

where E is the measured half-cell potential, E_0 is the "reference potential," a constant, z is a dimensionless constant for the given half-cell reaction, and Q is the reaction quotient. z is nominally the valence of the ion undergoing reduction or oxidation. But in cases where competing reactions occur or, generally, Faradaic efficiency is less than 100%, z can be substantially reduced and becomes generally unknown. Q is a quantity constructed from reactant activities, with each activity raised to a particular exponent. Where reactant concentrations are low, the activity is approximately equal to the concentration. So, if only one reactant's concentration, c, is varied at low levels, Q can be

written as

$$Q = Kc^{\alpha} \tag{275}$$

where K and α can be considered as empirical constants.

The experimental data would consist of a number of data points (c_i, E_i), where c_i is varied and the half-cell potential E_i is recorded. The assumed Nernst equation can be manipulated into a linear form:

$$E_i = E_0 - \frac{0.05916\text{V}}{z} \log_{10} Q_i \tag{276}$$

$$E_i = E_0 - \frac{0.05916\text{V}}{z} \log_{10} K c_i^{\alpha} \tag{277}$$

$$E_i = E_0 - \frac{0.05916\text{V}}{z} \left[\log_{10} K + \alpha \log_{10} c_i\right] \tag{278}$$

$$E_i = E_0 - \frac{0.05916\text{V}}{z} \log_{10} K - \frac{0.05916\text{V}}{z} \alpha \log_{10} c_i \tag{279}$$

$$\underbrace{E_i}_{y_i} = \underbrace{\left[\frac{0.05916\text{V}}{z}\alpha\right]}_{m} \underbrace{\log_{10} c_i}_{x_i} + \underbrace{\left[E_0 - \frac{0.05916\text{V}}{z} \log_{10} K\right]}_{b} \tag{280}$$

From a given set of data, an m and b can be calculated. This would yield a predictive model that could be used to predict E_i for any given value of c_i. As well, from at least two different extracted values for m and b (*i.e.* from two data sets), the values for z, K, and α could be determined. With appropriate manipulation and transformation, it is possible to use simple linear regression in a large number of situations where the governing relations are not strictly linear equations.

B.4.3 Linear Least-Squares Regression

The algebraic derivation that was used to achieve the matrix-form 2-equation linear system for simple linear regression can be generalized to accommodate more sophisticated functions than a simple line. This allows the decomposition of a given set of data into any number of arbitrary basis functions.

B.4.3.1 Mathematical Theory

In simple linear regression, the linear relation

$$y(x_i) = mx_i + b$$

established $y(x)$ as a *predictor* of y_i, based on a single quantity x_i. This predictor relation is generalized in two ways for linear least-square regression. Instead of a single independent variable x, a set of l independent variables, $x_1 \ldots x_l$, collectively referred to as the vector $\hat{\mathbf{x}}$, is incorporated into the predictor relation, and specified at each data point (as $\hat{\mathbf{x}}_i$), The predictor function is expressed as the linear combination of an arbitrary number k of arbitrarily-valued basis functions $g_j(\hat{\mathbf{x}})$:

$$y(\hat{\mathbf{x}}) = \sum_{j=1}^{k} \beta_j g_j(\hat{\mathbf{x}}) \tag{281}$$

where β_j are the coefficients used in the linear combination. The coefficients β_j are, like the individual x_i values, organized into a vector $\hat{\beta}$ for convenience. It is also assumed that functions g_j, referred to as the *basis functions*, are independent of all β_m, m in $1 \ldots k$.

The problem of linear least-squares regression, then, is to find the values for $\beta_1 \ldots \beta_k$ that minimize the average squared error $E\left(\hat{\beta}\right)$

$$E\left(\hat{\beta}\right) = \sum_{i=1}^{T} \left(y_i - y(\hat{\mathbf{x}}_i)\right)^2 \tag{282}$$

Just as in simple linear regression, this minimum is found by equating the first derivative of $E\left(\hat{\beta}\right)$ with respect to each β_i to zero. One key observation toward this end depends on the independence of the $g_j(\hat{\mathbf{x}})$ on $\hat{\beta}$:

$$\frac{\partial y(\hat{\mathbf{x}})}{\partial \beta_j} = \frac{\partial}{\partial \beta_j} \sum_{i=1}^{T} \beta_i g_i(\hat{\mathbf{x}}) \tag{283}$$

$$= g_j(\hat{\mathbf{x}}) \tag{284}$$

Then, the derivatives of the squared error become

$$\frac{\partial E\left(\hat{\beta}\right)}{\partial \beta_j} = \frac{\partial}{\partial \beta_j} \sum_{i=1}^{T} \left[y_i - y\left(\hat{\mathbf{x}}_i\right)\right]^2 \tag{285}$$

$$= 2 \sum_{i=1}^{T} \left[y_i - y\left(\hat{\mathbf{x}}_i\right)\right] \frac{\partial y\left(\hat{\mathbf{x}}_i\right)}{\partial \beta_j} \tag{286}$$

$$= 2 \sum_{i=1}^{T} \left[y_i - y\left(\hat{\mathbf{x}}_i\right)\right] g_j\left(\hat{\mathbf{x}}_i\right) \tag{287}$$

$$= 2 \sum_{i=1}^{T} \left[y_i - \sum_{n=1}^{l} \beta_n g_n\left(\hat{\mathbf{x}}_i\right)\right] g_j\left(\hat{\mathbf{x}}_i\right) \tag{288}$$

$$= 2 \sum_{i=1}^{T} y_i g_j\left(\hat{\mathbf{x}}_i\right) - 2 \sum_{i=1}^{T} \sum_{n=1}^{l} \beta_n g_n\left(\hat{\mathbf{x}}_i\right) g_j\left(\hat{\mathbf{x}}_i\right) \tag{289}$$

$$= 2 \sum_{i=1}^{T} y_i g_j\left(\hat{\mathbf{x}}_i\right) - 2 \sum_{n=1}^{l} \beta_n \left(\sum_{i=1}^{T} g_n\left(\hat{\mathbf{x}}_i\right) g_j\left(\hat{\mathbf{x}}_i\right) \right) \tag{290}$$

Each first derivative, then, is seen to be, with all $\hat{\mathbf{x}}_i$ and y_i known, a linear combination of the β_j. Equating all first derivatives to zero creates k linear equations involving the k coefficients β_j, enabling the expression of the system in standard, square-matrix linear form:

$$\begin{bmatrix} \sum_{i=1}^{T} g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) & \cdots & \sum_{i=1}^{T} g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) \\ \vdots & \ddots & \vdots \\ \sum_{i=1}^{T} g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) & \cdots & \sum_{i=1}^{T} g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) \end{bmatrix} \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_k \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^{T} y_i g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) \\ \vdots \\ \sum_{i=1}^{T} y_i g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) \end{bmatrix} \tag{291}$$

so that

$$\begin{bmatrix} \beta_1 \\ \vdots \\ \beta_k \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^{T} g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) & \cdots & \sum_{i=1}^{T} g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) \\ \vdots & \ddots & \vdots \\ \sum_{i=1}^{T} g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) & \cdots & \sum_{i=1}^{T} g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) \end{bmatrix}^{-1} \begin{bmatrix} \sum_{i=1}^{T} y_i g_1\left(\hat{\mathbf{x}}_\mathbf{i}\right) \\ \vdots \\ \sum_{i=1}^{T} y_i g_k\left(\hat{\mathbf{x}}_\mathbf{i}\right) \end{bmatrix} \tag{292}$$

From these formulations, matrix inversion algorithms can be used to find $\hat{\beta}$ that makes all first derivatives zero.

Checking that all second derivatives are always non-negative is straightforward:

$$\frac{\partial^2 E\left(\hat{\beta}\right)}{\partial \beta_j^2} = \frac{\partial}{\partial \beta_j} \frac{\partial E\left(\hat{\beta}\right)}{\partial \beta_j} \tag{293}$$

$$= \frac{\partial}{\partial \beta_j} \left[2 \sum_{i=1}^{T} y_i g_j\left(\hat{\mathbf{x}}_i\right) - 2 \sum_{n=1}^{l} \beta_n \left(\sum_{i=1}^{T} g_n\left(\hat{\mathbf{x}}_i\right) g_j\left(\hat{\mathbf{x}}_i\right) \right) \right] \tag{294}$$

$$= 2 \frac{\partial}{\partial \beta_j} \sum_{n=1}^{l} \beta_n \left(\sum_{i=1}^{T} g_n\left(\hat{\mathbf{x}}_i\right) g_j\left(\hat{\mathbf{x}}_i\right) \right) \tag{295}$$

$$= 2 \frac{\partial}{\partial \beta_j} \sum_{n=1}^{l} \delta_{jn} \left(\sum_{i=1}^{T} g_n\left(\hat{\mathbf{x}}_i\right) g_j\left(\hat{\mathbf{x}}_i\right) \right) \tag{296}$$

$$= 2 \sum_{i=1}^{T} g_j\left(\hat{\mathbf{x}}_i\right)^2 \tag{297}$$

indicating that any first-derivative zeroes will indeed be global minima points.

B.4.3.2 Application

Linear least-squares regression allows much more flexibility than simple linear regression in arranging the expectation of a fit between a dataset and the model. For instance, by choosing

$$g_i(x) = e^{i\omega x} \tag{298}$$

with ω set to the fundamental frequency, x_i interpreted as time, and y_i interpreted as a signal's value at time x_i, linear least-squares regression becomes equivalent to a discrete Fourier series calculation. However, there are two major practical considerations in using linear least-squares regression.

First, with the low analytical and computational cost of adding basis functions, there is a tendency to incorporate a large number of basis functions into the model. This presents a problem primarily in qualitatively interpretating the least squares fit result. Various approaches exist for identifying a minimum "suitable" set of basis functions.

The second major concern with least-squares fitting is the relative degree to which individual samples influence the overall fit. It is the total squared error function E that is minimized, its

sensitivity to a particular data point y_p is (without considering y_p's impact on $\hat{\beta}$):

$$\frac{\partial E}{\partial y_p} = \frac{\partial}{\partial y_p}\left[\sum_{i=1}^{T}\left[y_i - \sum_{n=1}^{l}\beta_n g_n\left(\hat{\mathbf{x}}_i\right)\right]^2\right] \tag{299}$$

$$= 2\sum_{i=1}^{T}\left[y_i - \sum_{n=1}^{l}\beta_n g_n\left(\hat{\mathbf{x}}_i\right)\right]\delta_{ip} \tag{300}$$

$$= 2\left[y_p - \sum_{n=1}^{l}\beta_n g_n\left(\hat{\mathbf{x}}_p\right)\right] \tag{301}$$

Though this is quantitatively inconclusive, due to the assumption that $\frac{\partial\hat{\beta}}{\partial y_p} = 0$ it shows that, qualitatively, data points where either y_p or the functions $g_n\left(\hat{\mathbf{x}}\right)$ are large in magnitude, have a large effect on E. For basis functions with large dynamic ranges (over the $\hat{\mathbf{x}}_i$), this can lead to observations at the small end of the dynamic range being effectively ignored.

One way to address this concern is to transform the y_i and the $g_j\left(\hat{\mathbf{x}}\right)$ to reduce their dynamic range. Note that, for the intended model fitting to retain its predictive value, both y_i and the $g_j\left(\hat{\mathbf{x}}\right)$ must undergo precisely the same transformation.

Another way to address this concern is through the use of weighted least squares. Weighted least squares adds, to each data point y_i, a weight w_i, which is then used to weight that data point's contribution to E:

$$E = \frac{1}{T}\sum_{i=1}^{T} w_i\left[y_i - y\left(x_i\right)\right]^2 \tag{302}$$

Weighted least squares allows not only compensation for variation in magnitude of basis functions and samples, but also for accommodation of non-uniform variance (standard deviation) among subsets of the data points.

B.4.4 Nonlinear Least-Squares Regression

Linear least-squares regression offers a compelling triplet of advantages:

- it always identifies *the* global minimum for $E\left(\hat{\beta}\right)$ (unless the matrix in Equation 292 is singular, in which case the method will yield no solution at all),
- it is numerically and analytically simple, and
- arbitrary basis functions $g_j\left(\hat{\mathbf{x}}\right)$ may be accommodated.

However, linear least-squares suffers from two equally compelling limitations. First, in its identification of an absolute global minimum solution, it disallows the introduction of constraints on the solution space. For instance, in allocating core losses among different loss mechanisms, linear least squares is liable to identify negative coefficients for some basis functions. This obviously violates the same physical principles that were used to formulate the basis functions, meaning that linear least squares can ultimately introduce inconsistencies into the analysis produced.

Second, linear least squares only allows solutions that are linear combinations of the basis functions. In many cases, the unknowns that must be identified by curve fitting are not simple linear coefficients. For instance, in many cases, hysteresis losses are empirically found to vary with magnetic flux density raised to an empirically-discovered power other than 2. In this case, the loss is a nonlinear function of the unknown, and in the absence of very strong experimental assumptions (*ex.* that only hysteresis losses are present), linear least squares is not able to be of assistance.

The method of nonlinear least squares alleviates both of these issues with linear least squares.

B.4.4.1 Mathematical Theory

Nonlinear least squares can be presented using the same formalism presented in subsubsection B.4.3.1. The formulation of the predictor y Equation 281 is generalized to include $\hat{\beta}$ as an argument, rather than as a coefficient vector:

$$y = y\left(\hat{\beta}, \hat{\mathbf{x}}\right) \tag{303}$$

With this, the average squared error $E\left(\hat{\beta}\right)$ becomes

$$E\left(\hat{\beta}\right) = \sum_{i=1}^{T} \left(y_i - y\left(\hat{\beta}, \hat{\mathbf{x}}_i\right)\right)^2 \tag{304}$$

From this point, the method of nonlinear least squares amounts to a general minimization problem: the average squared error $E\left(\hat{\mathbf{beta}}\right)$ is to be minimzed with respect to the multidimensional vector $\hat{\mathbf{beta}}$. Any of a wide number of minimization or optimization algorithms may be used, including gradient descent and Newton-Raphson.

B.4.4.2 Application

One common step in many of the optimization algorithms used for nonlinear least squares is the evaluation of the derivatives $\frac{\partial E}{\partial \beta_i}$. These derivatives may be evaluated either numerically or symbolically. Symbolic evaluation requires that the predictor y be differentiable with respect to β_i while numeric evaluation entails a separate evaulation of the predictor to determine each β_i's partial derivative.

Identifying a global minimum with typical optimization algorithms is a matter of probability and, to a lesser extent, tuning of the initial conditions and iteration parameters of the optimization algorithm. This stands in contrast to the provably global minimum that is identified by a single iteration of a least-squares algorithm.

In summary, applying nonlinear least squares is a nontrivial task that involves a substantial amount of oversight. The MATLAB/Octave functions *lsqnonneg* and *fminunc* present well-tuned implementations of nonlinear least squares. These were relied on for all nonlinear least squares-based results presented in this thesis.

APPENDIX C

DESIGN AND OPERATION OF ELECTROPLATING ROBOT

C.1 System Architecture

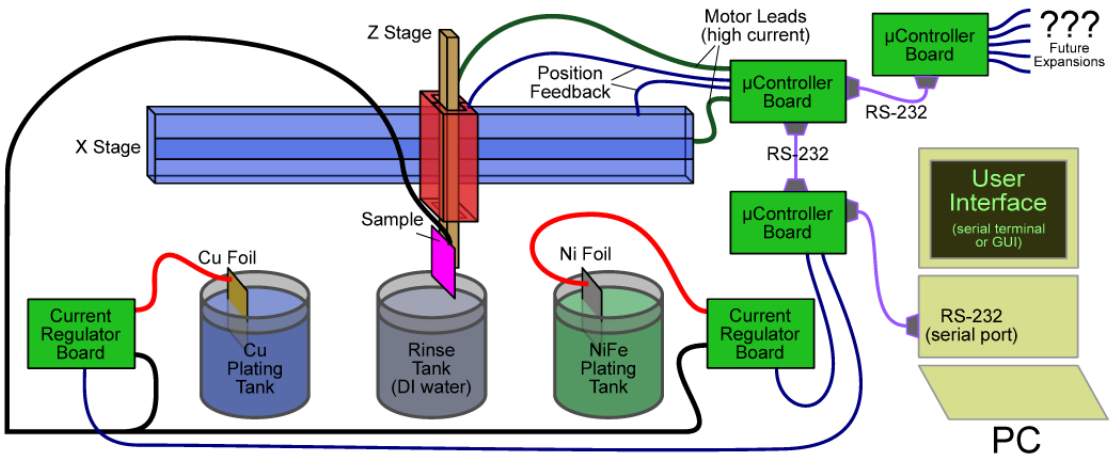

Figure 158: Automated Multilayer Electroplating System Architecture

The electroplating robot was developed to perform the well-defined task of electrodepositing layers onto a single wafer from two or more different plating baths. This directly established any number of straightforward requirements for the system, such as:

- provide for two separate plating tanks, and at least one rinse tank,
- provide a means to drive a controlled plating current through the plating baths,
- allow adjustment of plating times, currents, and layer counts, and
- transport the wafer between the various tanks while preserving the electrical connections required to carry out the electrodeposition.

These requirements were largely unchanged from the robot's prototype development, presented in [56]. They determined the basic design of the robot, such as its two-axis design and horizontal arrangement of tanks.

However, beyond these requirements, the system was charged with optimizing for two major performance parameters, and these heavily influenced its architecture.

C.1.1 Reliability

The electroplating robot was designed for extremely robust operation, allowing process runs that might take several days to complete. The anticipated need to produce high-volume cores, well over 100 μm in total height, was the principal driver for this need, in light of the relatively immutable rate of magnetic alloy deposition (approximately 10 μm per hour). As well, the large amount of time and resource investment required to produce a multilayer plated structure, and the numerous ways in which a malfunction of the plating system could render these structures unusable, easily justified a strong upfront emphasis on reliable system operation.

This directive for high reliability was manifested in several ways. First, all motion control was performed with closed-loop control systems. The hazard of wafer breakage due to mislocation was an obvious driver for this decision, but the potentials for the mechanical systems to damage themselves due to mislocation, or for the controller software to enter an unknown state were also definitively avoided. Second, all real-time control was performed with microcontrollers. The use of a PC, particularly one running Windows XP, as well as the use of a large, license-bound and memory leak-prone application such as LabView or MATLAB were judged to be prohibitive risks to the system's achievable uptime. Third, all high-frequency circuits, such as the microcontrollers or the PWM motor drivers, were implemented on custom PCBs with carefully designed ground planes. This decision was made chiefly because the possibility of electrical noise from one subsystem, such as the z stage motor driver interfering with another system, such as the plating controller's current sensor or the microcontroller's reset line was judged to be a substantial risk.

C.1.2 Flexibility

Another overarching goal for the development of the plating system was that it be as flexible as possible. This general concept of flexibility was realized in a few major areas.

C.1.2.1 Control Flexibility

The use of the command line interface (CLI) framework, discussed later in this appendix, for the microcontroller coding allowed for arbitrary changes to the controller's operation to be effortlessly accommodated in the user interface. As new features were found to be desirable in the different subsystems, this software architecture greatly reduced the development and debug time associated with the addition of each of these new features.

To allow an even greater degree of flexibility in the system's high-level operation, the sequencer subsystem was physically separated from the rest of the control electronics. The interconnection between the sequencer and the system was fully encapsulated in a standard 25-pin "D-Sub" (DB-25) cable and made exclusive use of quasi-static 5 Volt logic signals. This arrangement allows for the replacement of the system's original sequencer with *any* system capable of 5 Volt digital input and output. Examples of such systems include a PC running LabView or MATLAB, equipped with a digital I/O card, another microcontroller system, or even a panel consisting of manual switches and dials.

C.1.2.2 Tank Array Flexibility

One of the most often-used degrees of freedom in the plating system was its ability to accommodate an arbitrary set of plating and/or rinse tanks. This was designed into the system from the beginning. Not only was the travel range of the x axis extended to its absolute maximum, but the x axis controller software was designed to allow any number of tank locations, each with an arbitrary location along the axis of travel.

C.1.2.3 User Interface Flexibility

The plating system was designed for use with a PC-based GUI. The use of serial ports, which are well-supported by an enormous range of programming platforms in all major PC operating systems, was the key feature to enable GUI control. The commonality of the command-line syntax, due to the uniform use of the CLI framework, also facilitated the rapid development and maintenance of PC-based control programs.

C.2 Design Approaches

C.2.1 Fluidic

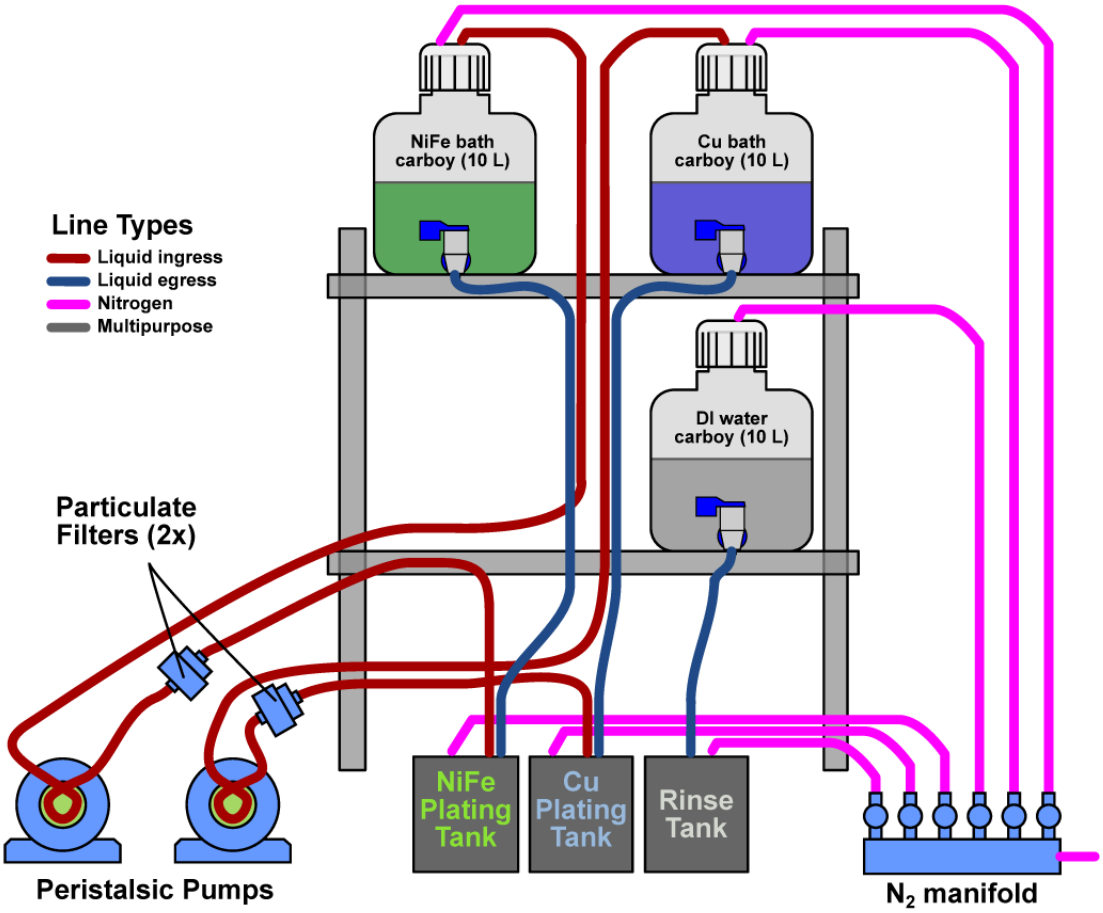

Figure 159: Fluidic System Design for Electroplating Robot

The fluidic system for the plating robot was developed incrementally, and served the following roles:

- provision of nitrogen-passivated storage of plating baths while not in use,
- bidirectional transport of plating baths between storage containers and plating tanks,
- precise control over fluid levels in plating and rinse tanks,
- regular filtration of plating baths,

- transportation of rinse water between rinse tank(s) and rinse water storage containers, and
- low-flow nitrogen aeration of plating tanks during plating.

The complete implementation of the fluidic system is shown in Figure 159. Carboys with caps modified to accept nitrogen line pass-throughs were used to provide long-term storage of the plating baths. Gravtiy-driven flow, along with spigots installed on the plating bath carboys, was used to move the plating bath from the carboys into the plating tanks. Peristaltic tanks were used to move the plating bath from the tanks back into the carboys. The filtration task was achieved with two inline filters, one for each plating bath carboy; this filtration was driven with the peristaltic pumps and therefore performed every time the plating baths were moved back into their carboys.

The nitrogen aeration of the plating tanks was implemented with lengths of polypropylene tubing, coupled with needle valves to achieve precise flow control at low rates. A carboy was used to transport rinse water into the robot's area, in lieu of installing a new water faucet. This carboy was also outfitted with a nitrogen line to allow deoxygenation prior to plating runs, as it was found that remanent oxygen in the rinse water could lead to small amounts of iron oxide in the rinse water. The peristaltic pumps were plumbed so that either pump, along with its dedicated tubing, could be used to move used rinse water from the rinse tank(s) to a waste container.

C.2.2 Electronics

C.2.2.1 Implementation and Packaging

All electronics were implemented using printed circuit boards (PCBs) wherever possible, not only to provide robust interconnect and mechanical stability, but also to minimize noise-inducing electrical parasitics. In particular, the electronic communication between boards needed to be free of spurious events caused by noise. As well, all subsystems had the potential to be adversely affected by internal noise.

Toward the same end, all electronics were housed in walled, modestly-sealed enclosures. The principal enclosures, as well as a view of the electronics implementations, is shown in Figure 160. Also shown in Figure 160, industry-standard connectors were used to make connections between the electronics enclosures and other system elements. The use of such enclosures not only allowed

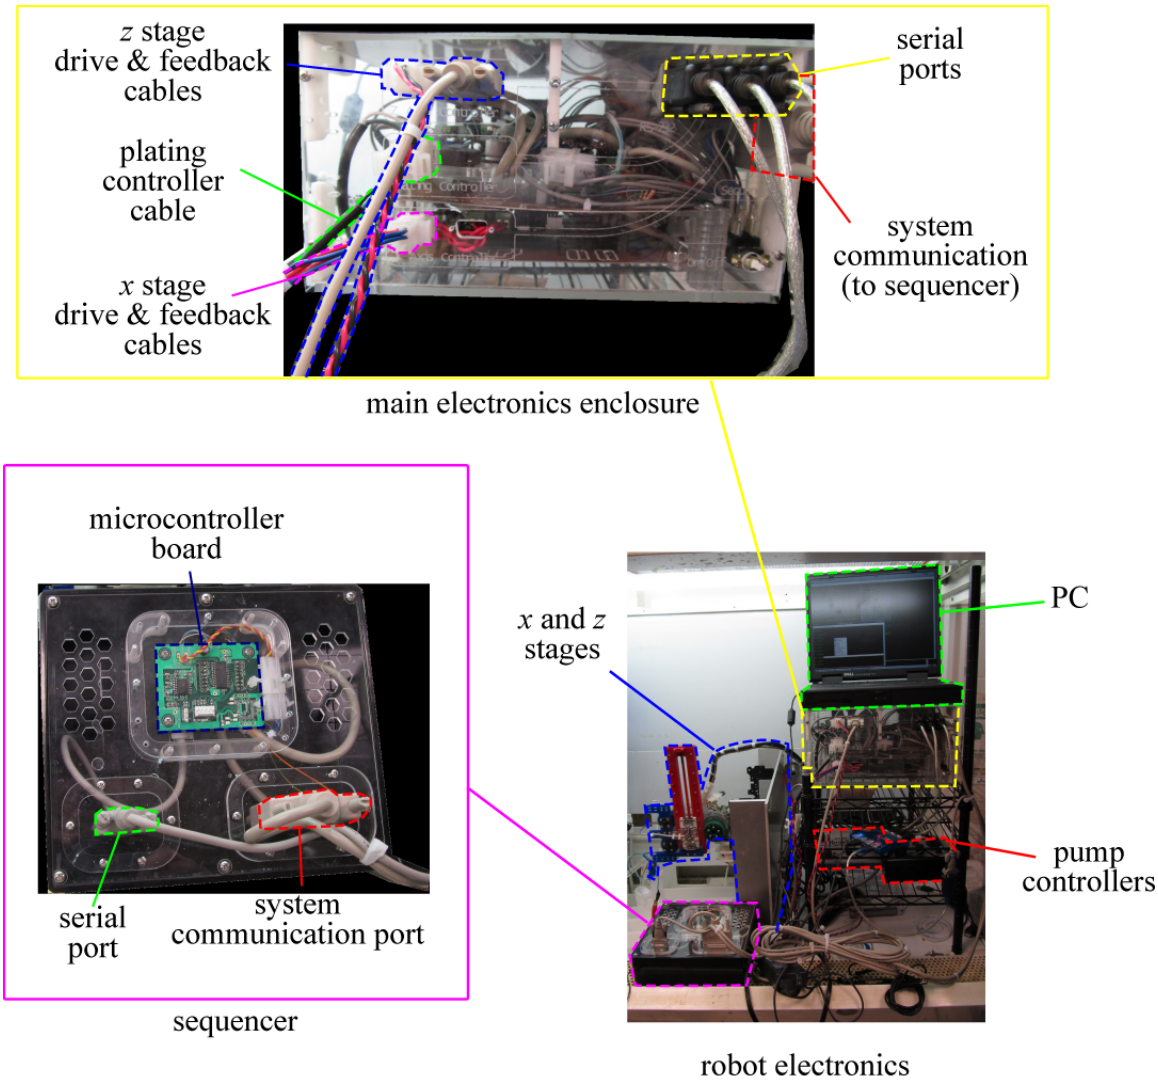

Figure 160: Electronics System for Electroplating Robot. The robot's electronics were divided between two separate enclosures, to allow modularized high-level control.

easy and repeatable connection between components, but also protected the electronics from the potentially destructive chemicals used in the plating processes.

C.2.2.2 Communication

Electronic communication was split into three domains:

Programming communication consisted solely of the task of downloading compiled programs into the microcontrollers' on-board flash memory. Each board had a manufacturer-specified 5-pin header for programming.

Static communication consisted of all user interaction, and principally provided for configuration and parameterization of the various tasks. Each microcontroller had its own RS-232 port for static communication; the host PC, then, needed to have 4 available serial ports to simultaneous communicate with all 4 processors. USB-to-RS-232 converters were used to provide the needed number of serial ports.

Dynamic communication occurred between the sequencer and the other boards, and provided for the automatic execution of the complete multilayer plating process. A low-overhead digital protocol, the dynamic control system (DCS), was developed for this purpose, with the sequencer acting as the master for the other three processors.

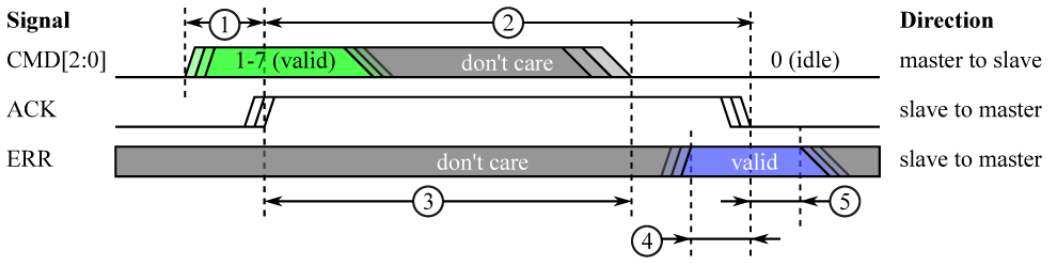

Timing Parameter	Description	Min/Max Value	Comments
①	valid command to ACK assert	0 ms / 500 ms	specified as very long time, to allow for slow polling rates at slaves
②	ACK assert to ACK deassert	100 ms / 5 min.	actual DCS command (ex. x stage move) executes in full during this time
③	ACK assert to command idle	0 ms / 100 ms	short enough to allow rapid execution of fast-executing slave commands
④	ERR to ACK deassertion setup	30 ms / ∞	long enough to allow substantial debounce of ERR by master
⑤	ERR to ACK deassertion hold	150 ms / ∞	long enough to allow extensive debounce of ERR by master

Figure 161: Dynamic Communication System (DCS) protocol

Programming communication was handled entirely by the microcontrollers. Static communication was implemented using the automatically-generated CLI interfaces, described below.

The DCS system was designed to place minimal requirements on all parties, mostly to allow the substitution and interworking of a very wide variety of components. The DCS system was organized as a strict, single-master system, wherein one component, called the “master,” issued commands to

all other components, called “slaves.”

There were five dedicated digital signals connecting each slave to the master: three lines for the master to specify a command, and two lines for the slave to provide status to the master. Figure 161 illustrates the DCS protocol, as well as the major timing requirements. The idle state of the DCS link is where the master is driving all zeros (“idle” command) onto the CMD lines, and the slave has its acknowledge (ACK) signal deasserted. The master initiates a command by driving a nonzero value onto the CMD lines. The slave, in response, asserts its ACK line, and begins executing the command corresponding to the value on the CMD lines. The slave deasserts its ACK line upon completion of the command, and holds its error indicator (ERR) at the appropriate value to signal to the master whether or not an error was encountered.

All DCS lines are positive-high 5 Volt CMOS logic. This, in combination with the extremely wide timing windows, makes the DCS protocol something with which a wide variety of systems can interoperate. In particular, the DCS protocol should be very compatible with systems where the digital inputs and outputs are handled solely by software (*i.e.* no hardware assistance).

C.2.3 Software

C.2.3.1 Standard Platform

All embedded software was implemented on Cypress PSoC CY8C29466 microcontrollers. This decision was made to leverage previous experience with this platform, as well as a preponderance of assembled general-purpose PCBs within the means of this research. The PSoC microcontroller includes an on-board array of programmable, very flexible analog and digital blocks, such as op-amps, analog-digital and digital-analog converters, serial-to-parallel converters, PWM generators, and timers. These on-board functions were heavily utilized in the implementation of the electro-plating system.

Use of the same platform across all embedded subsystems was a key aspect of the rapid rate at which the embedded software was developed. All development was done in assembly language, partly to sidestep the cost of a C compiler, but also to allow seamless integration with the CLI framework. This required intimate articulation with the PSoC’s memory architecture, as well as with various other critical aspects of the PSoC’s operation, and all associated learning was leveraged

across all software development tasks.

C.2.3.2 CLI Framework

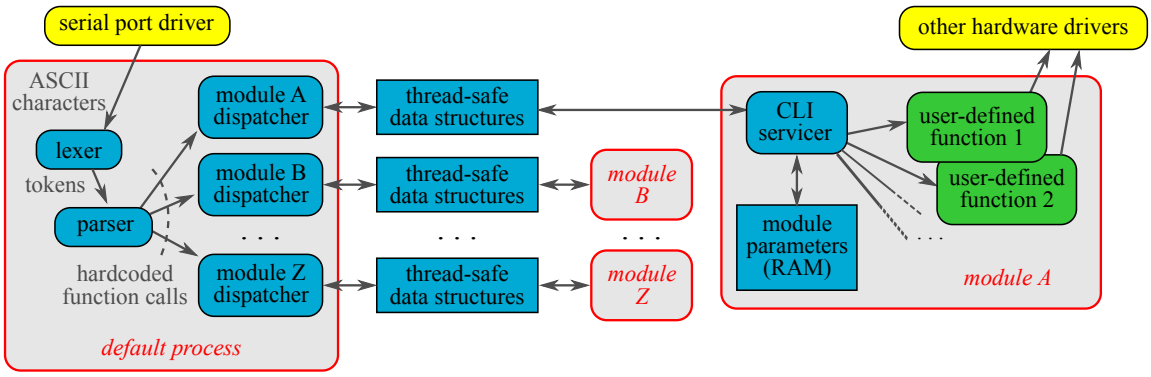

Figure 162: CLI Framework Architecture. The serial port interface is the principal point of user interaction. Functionality is divided into a number of modules, each of which run asynchronous to each other. All blue blocks are generated by a Java program; yellow blocks represent hardware interaction functions; green blocks require manual authoring and debug.

The command line interface (CLI) framework was a principal component of the plating system's software development. The CLI framework, diagrammed in Figure 162 automates the process of developing user interface code. It also implements thread-safe communication between a user interface process and all "business" processes (*i.e.* processes that actually perform the microcontroller's externally-relevant tasks). As a result, business process code could be implemented in a relatively naive and independent manner without negatively impacting either the software's usability or its stability. This gain is in addition to the avoidance of needing to manually implement user interface routines.

The CLI interface allows the specification of any number of "parameters" and/or "functions." These parameters and functions are all directly accessible via a command-line (*i.e.* DOS-like) interface that is automatically attached to the microcontroller's built-in serial port. From the CLI, the user can query or set any parameter's value, as well as invoke any function.

An extensible markup language (XML) file specifies all parameters and functions, as well as their organization into any number of "modules." The CLI framework, when given this XML file, automatically generates all source code necessary to implement the CLI. For each parameter, the CLI framework generates a variable name that the user's code can access. For each function, the

user specifies the name of the function, and the CLI framework-generated code calls this function. All CLI framework-generated code is in assembly, and is typically many thousands of lines of code.

This includes, as shown in Figure 162, not only the functionality for accessing parameters and functions, but also a highly-optimized set of building blocks for the CLI interface. The latter include the lexer, which translates a disorganized stream of input characters into a stream of lexical elements, known as tokens. The parser takes the stream of tokens, as distinguished and indicated by the lexer, and recognizes valid syntactic streams, and dispatches all associated commands and returns appropriate output to the user.

C.3 Subsystems

C.3.1 *x* Stage

The *x* stage was taken in entirety from a decomissioned pen plotter. It consisted of a boss that was mounted, via linear slide bearings, on two parallel unthreaded shafts. The mechanism was driven by a two-phase stepper motor, which was coupled to the boss via a steel cable and pulleys. It was found that the *x* stage motor was not prone to missing steps, without the application of very large forces. Accordingly, precise position feedback was not implemented, although end-of-travel sensing was implemented using optointerrupters.

C.3.2 *z* Stage

The *z* stage was a fully-custom mechanism designed around an automotive power window motor. The power window motor provided an extremely high amount of torque but very limited angular travel. A schematic of the *z* stage is shown in Figure 163. The mechanism was a double-sided rack and pinion system, built in planar fashion atop a large backplate (stator). Two idler gears were included opposite the drive gear (pinion) to act as in-plane bearings. A small plate with teflon bushings attached acted as an out-of-plane bearing. This small plate incorporated two optointerrupters which were used to detect end-of-travel conditions for the rotor. Not shown in Figure 163 is an optical quadrature feedback wheel which was used to sense the drive gear's position.

This design of the *z* stage was chosen because it fit well with in-house mechanical fabrication capabilities, *viz* laser cutting of plastic blanks, and allowed for a very low-cost construction. This approach also minimized the motion that would be experienced by the associated wiring. The only

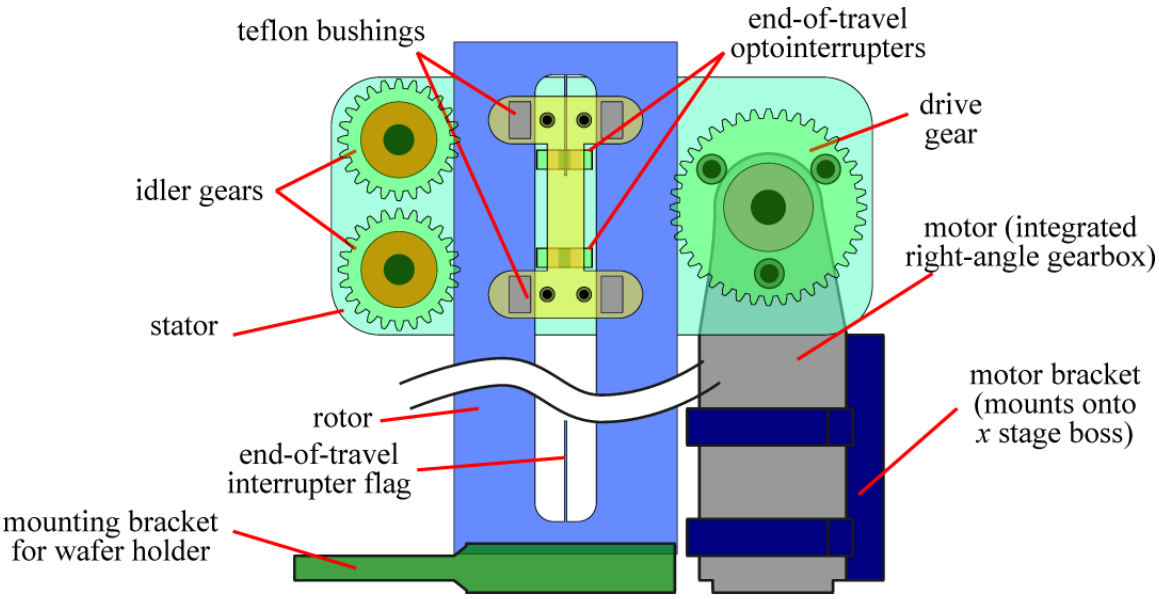

Figure 163: z Stage Mechanism

wire that would be required to flex as the z stage was actuated is the plating lead that connects to the wafer. All other drive and sense connections are stationary relative to the x stage boss.

This design performed well, although it was found that the resolution of the laser machining process, as well as the durability of the acrylic plastic used to fabricate all components, was not sufficient to achieve maintenance-free operation. The idler gears and rotor were found to need replacement approximately after approximately 80 hours of system operation. A mechanically tougher polymer would have been superior in this regard. However, such a polymer would also have to be easily laser-machinable. Delrin, one such polymer, was used in early z stage components, but it was found that the acidic plating baths corroded the Delrin to the point of being unusable after only a few runs.

C.3.3 Plating Controller

The 2-channel plating controller was built with full-bridge and half-bridge MOSFET drivers, to allow reverse pulse plating, and employs closed-loop control. A schematic of the plating controller, including its connection to the sample and plating foils, is shown in Figure 164. During plating, a PWM signal is applied to the half-bridge connected to the appropriate foil. At all times, the wafer's half-bridge is held at DC output, to avoid stray electrochemical interactions between the wafer and the plating bath.

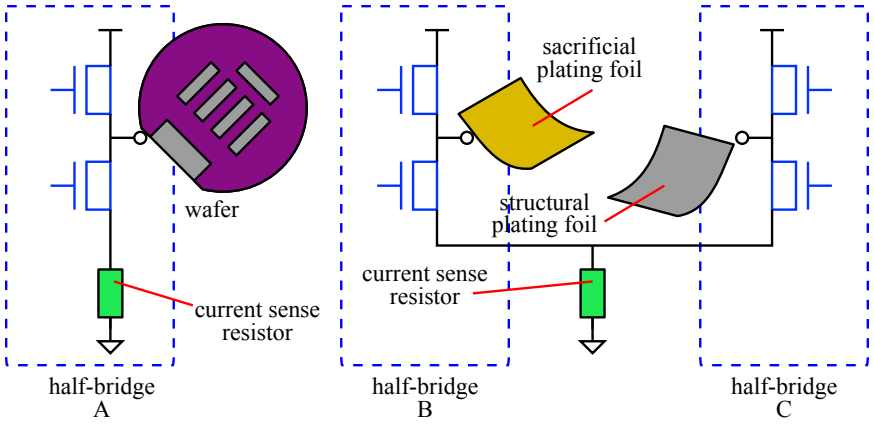

Figure 164: Plating Controller Schematic

This design allows for forward or reverse plating to occur in either plating bath, assuming the wafer is positioned properly. For instance, to perform forward plating (deposition) in the sacrificial plating bath, the wafer's half bridge (A) will be driven to a low output, tying the wafer to ground through the current sense resistor. The microcontroller's on-board analog-to-digital converter was used, in combination with a sense resistor, to monitor the plating current. A PWM signal is applied to the sacrificial foil's half-bridge (B) to achieve the desired current level. Note that the two foils' half-bridges are able to share a current sense resistor, leading to their implementation in a single full-bridge

Inside the microcontroller, a low-bandwidth proportional-integral (PI) controller was used. All associated electronics and software were implemente on board the microcontroller. The voltage from each sense resistor was independently amplified by a programmable-gain op-amp stage. A 14-bit D/A converter, with its input multiplexed between the two sense resistors, was used to read these current sense voltages. A 100 Hz interrupt was used to update the control output.

C.3.4 Sequencer

The sequencer subsystem was responsible for issuing dynamic commands, using the DCS protocol described earlier in this appendix, to the other three subsystems, as needed to orchestrate the desired fabrication process. The sequencer was implemented on a dedicated PSoC microcontroller. It was encoded entirely in assembly language, although a lightweight multitasking system was incorporated.

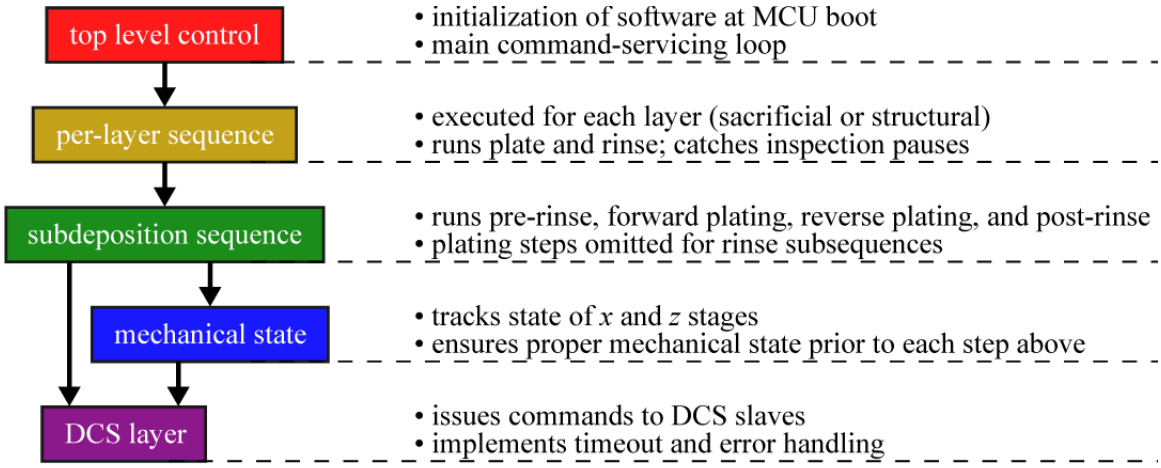

Figure 165: Sequencer Abstraction Layers

This multitasking system allowed for re-entrant execution of assembly routines. This enabled the use of procedure-based programming, rather than a state machine-based implementation of the machine's operation. This multitasking facility also provided for an arbitrary number of software abstraction layers to model the system's operation. This led to the layered implementation shown in Figure 165.

C.4 Operating Procedure

C.4.1 Before the Run

C.4.1.1 Boric Acid Preparation

Dissolve 25g of boric acid (dry powder) in approximately 1 L of water. This will be enough to condition 10 L of rinse water.

C.4.1.2 Tank Setup

Mechanical

All tanks should be clean, freshly rinsed with DI water, and in-place, with the tank-top fixture installed, prior to the introduction of any liquids.

Foils

The foils should be thoroughly cleaned, deoxidized, and inspected prior to operating the robot. Poorly-conditioned foils can introduce chemical and particulate contamination into the plating baths, as well as significantly altered surface areas.

Cleaning should be both mechanical and chemical. Mechanical cleaning can be done with Texwipes moistened with various solvents. At the least, a DI water-soaked should Texwipe be used on the entire surface of the foil before every run. In cases where non-water-soluble contamination is apparent on the foils (*ex.* grease), a standard organic solvent clean sequence, such as acetone-methanol-DI water, should be followed, with mechanical assistance from swabs or Texwipes as needed.

Deoxidation is very important for the copper foil, but can be generally skipped for the nickel foil. The copper foil can be deoxidized with a 1:4 HCl:DI solution, requiring no more than 5 minutes total immersion time, followed by a thorough DI rinse. Mechanical assistance should not be necessary for copper foil deoxidization.

Finally, the foils should be thoroughly inspected. Certainly, any contamination or mechanical damage to the foils should be checked for. But, it is also very important to check the foils for excessive topology or pores. As the foils dissolve during normal plating operation, they naturally roughen and, particularly in the case of nickel foils, form pores. As this roughness and the pores deepen, a number of undesirable effects become more prevalent. The foils' larger surface area can affect the plating process and the electrolyte balance of the bath, leading to overloaded baths. The foils' extreme relief makes it difficult or, in the case of deep pores, impossible to adequately clean the foils, leading to checmical contamination of the baths. Finally, foils with rough surfaces were found to produce much higher levels of particulate contamination, presumably due to the disintegration of the foil in the baths. In cases where nickel foils had pores extending through the entire thickness of the foil, unacceptably high levels of particulate contamination were found to be imminent.

Once the foils have been properly resurfaced, they should be clamped into place using the tank top assembly.

Tubing

Each plating tank requires three tubes that will reach approximately to the bottom of the tank: fluid ingress, fluid egress, and nitrogen agitation. The fluid ingress and egress tubes were used to, respecitvely, add and remove the appropriate fluid to and from the tank. All of these tubes should be situated in the tanks, and held in place by the tank top assembly, prior to the addition of any plating baths.

Rinse Tank Deoxygenation

It was found that iron oxide can develop in the nickel-iron rinse tank, and that not only the addition of boric acid into the rinse water but also the prior removal of dissolved oxygen from the water is required. Deoxygenation of the rinse water is achieved by filling the DI carboy with water and flowing nitrogen gas through it at a high rate, enough to produce violent bubbling, for approximately two hours.

C.4.1.3 Controller Configuration

Plating Controller

The plating controller should be configured to produce the appropriate currents for each type of layer – sacrificial and structural. Setting the plating controller's output current for a given layer type requires setting two parameters: current sense gain and target analog-to-digital converter (ADC) setting. Both of these are exposed as CLI parameters.

Since there is not a well-calibrated expression relating these two parameters to a desired current setting, it is necessary to use an ammeter (*i.e.* digital multimeter) to observe the current produced while these parameters are manually adjusted. Note that the foil lead should be electrically connected to the wafer lead while monitoring the controller's output current during this step. Although the plating controller acts as a constant-current source, maximally-accurate settings can be achieved by including a 1.5 Ω resistor between the wafer lead and the appropriate foil lead. This rather closely matches the actual loading that the plating controller will experience during plating.

x *Axis Controller*

The *x* axis controller requires only that the tank positions be set. There are four CLI parameters, one corresponding to each tank. These parameters are in units of motor micro-steps. In general, the default values, which are loaded at power-up of the microcontroller, are very close to the optimal values. However, some trial and error in adjusting these values is needed, at the least to accommodate the necessarily varied alignment between the tanks and the *x* stage assembly itself. In other words, the *x* stage has no means for sensing the positions of the tanks, and this variability can only be compensated out by the user, at this step.

z *Axis Controller*

The z axis controller requires no configuration. However, prior to starting a run, it should be cycled up and down a few times. During its motion, particular attention should be paid to the articulation of the gear teeth. If necessary, and dust visible on the gear teeth should be removed with a brush, and/or a small amount of lubricant (spray-on solid teflon) should be added.

Sequencer

The sequencer requires that the process parameters be configured. These parameters include the number of structural layers, whether the first layer is a sacrificial or structural layer, the time spent for each type of layer, and the number and duration of rinse steps. These are all directly exposed as CLI parameters.

C.4.2 Starting the Run

C.4.2.1 Critical Delays

When starting the run, there are a few timing-sensitive transactions, particularly segments of the process where time delays must be minimized.

Sample Preparation

One of the most critical parasitic processes is the oxidation of the copper seed layer on the sample. After the topmost titanium layer is removed, exposing the copper seed layer, this oxidation begins. Substantial oxidation of the areas to be plated presents a number of problems. Most importantly, the presence of excessive oxide at the beginning of the plating process can very easily result in nonuniform plating, as well as the generation of particulate contamination. Such oxides can typically be removed with a very brief (15s or less) dip in hydrochloric acid (HCl), but it was found that even small amounts of residue from this step, if left on the exposed areas, can be as harmful to the quality of the plated film as the oxide itself. Regardless of the risk of contamination from a HCl dip, the vigorous rinsing required after the dip presents further opportunity for oxidation, as well as for damage to the photoresist structures. This is especially true for high aspect-ratio structures, where longer rinses with more agitation are required, while the photoresist structures are more vulnerable to damage and delamination. Finally, despite the high purity of the copper in the seed layer, the oxide parasitically formed can not be readily removed in some cases, particularly in areas where the DI rinse water has fully evaporated, leaving regions of dry (and highly oxidized) copper.

This oxidation proceeds while the sample is fully submerged in the post-HF rinse, as well as while the sample is exposed to air. The post-HF rinse oxidation proceeds at a modest rate, requiring at least ten minutes or longer to become problematic, so it is not a major constraint on the operation of the robot. However, the exposure of the sample to air, such as during its transport from the wet bench to the robot, or while it is being clamed into the robot, must be carefully managed. While the sample is exposed to air, a few very damaging processes can take place. Certainly, oxidation was found to proceed at a much faster rate than while the sample was submerged in DI water. It was found that unacceptable amounts of oxidation could occur during less than one minute of exposure to air. As well, it is possible that the copper surfaces can become completely dry, and badly oxidized, as mentioned above. Finally, it was found on several occasions that very large airborne contaminants – visible to the naked eye – had adhered to the wafer during its exposure to air. Minimizing the time the sample spends in air is the key to reducing all of these risks, and is the single most important timing constraint during the robot's operation.

C.4.3 During the Run

During the run, very little user action should be required. However, the state of the rinse tanks should be monitored. As needed, the robot's operation can be paused using the sequencer's hold and inspect CLI commands. Hold commands will pause with the wafer still immersed in its most recent process tank, be it a plating tank or rinse tank. Inspect commands will pause immediately after raising the wafer out of its most recent process tank. Both hold and inspect commands will pause indefinitely, until a resume command is issued.

C.4.4 After the Run

Once the run is complete, the most important task is to carefully remove the wafer from its bracket and place it in an environment that is both chemically and physically safe, such as a beaker of freshly-dispensed DI water.

C.4.4.1 Plating bath return

Once the sample has been extricated, the first task is to return the plating baths to their long-term storage carboys. This is done using the peristaltic pumps. Prior to turning the pumps on, ensure

that there is fresh filter paper loaded into the filter modules. As well, it is important to have approximately 1L of fresh DI water on hand before beginning the pumping process.

This DI will be used to cleanse the pumps' tubing. It was found that not only will the plating baths corrode the silicone tubing over time, but that the crystals produced as residual bath dries will mechanically puncture the tubing. As the pumping process nears the end of the liquid in the tank, the ingress end of the tube should be quickly transferred into the 1L container of fresh DI.

As the pumping continues, the DI will progressively replace the plating bath along the length of the tube. The color of the liquid entering the carboy can be monitored at the carboy lid. The pump should be stopped only when the clear DI reaches the carboy's entrance. At this point, not only is the tube filled with putatively nonreactive DI water, but also has the vast majority of the plating bath been returned to the carboy.

C.4.4.2 Removing the rinse water

The peristaltic pumps can be used to pump the rinse water out of the rinse tanks into a waste container. Note that the rinse water might contain both chemical and particulate contamination, so it is very important to flush the pump's tubing with copious amounts of DI water after using it to move used rinse water.

C.4.4.3 Cleaning the foils

The penultimate task in post-run maintenance is to clean the plating foils. This consists of a simple, yet thorough, DI rinse followed by a very complete drying with nitrogen. This can be done while the pumps evacuate the plating tanks.

C.4.4.4 Cleaning the tanks

The final task in finishing off a robot run is the cleaning of the plating tanks. A thorough DI rinse and subsequent nitrogen drying is generally sufficient. However, if any chemical contamination is visible, particularly iron oxide, this should be addressed immediately after the run, to minimize the potential of the contamination forming strong adhesion to the tank's inner surface. Phosphoric acid is very effective at removing iron oxide, while hydrochloric acid can remove other types of visible ionic contamination (presumably metal salts).

The tank walls should also be inspected for organic or particulate contamination. A dry Texwipe can be used to test for the presence of these contaminants, and, with the proper solvent, can also be an optimal method for removing the contaminants.

APPENDIX D

OPERATING PROCEDURE FOR INDUCTOR LOSS CHARACTERIZATION SYSTEM

D.1 Test Procedure

Before doing any testing, the following initial setup procedure must be followed:

1. Ensure that the function generator, RF amplifier, and oscilloscope are turned off.
2. Disconnect the function generator from the RF amplifier.
3. Turn the function generator on, and ensure that it is configured to output sine waves into a 50 Ω load, and that its amplitude is less than 0 dBm (632 mV peak-to-peak). Driving more than 0 dBm into the RF amplifier's input will damage the equipment.
4. Connect the function generator's output to the RF amplifier's input, optionally including an attenuator. Note that including an attenuator will allow the function generator to safely output an accordingly higher voltage.
5. Turn on the oscilloscope.
6. Optional: connect the function generator's sync output to an oscilloscope channel (or external trigger input), and configure the oscilloscope to trigger off of that channel (or external trigger).
7. Ensure that the oscilloscope probes are properly connected to the motherboard, and (unless function generator's sync output is being utilized for triggering) configure the scope to trigger off of one of the channels (preferably $v_{out}(t)$).
8. Ensure that the RF amplifier's output is connected to the drive input of the motherboard.
9. Ensure that the RF amplifier's output is disabled.
10. Turn on the RF amplifier.

To execute a single frequency sweep with a given DUT, follow the following procedure.

1. Ensure that the RF amplifier's output is disabled.
2. Ensure that the DUT is soldered onto the motherboard.
3. Select and install the desired capacitance bank (daughter board), with the amplifier output disabled
4. Connect the RF amplifier and oscilloscope probes to the motherboard.
5. If applicable, position the cooling jet so that it is blowing directly onto the DUT.
6. Estimate the approximate resonance frequency using Equation 103. Set the function generator's frequency accordingly.
7. Set the function generator's amplitude to the minimum desired value.

At each datapoint within the frequency sweep, the following procedure should be followed:

1. Ensure that the oscilloscope is configured correctly and connected to the test circuit.
2. Ensure that the RF amplifier output is enabled.
3. Adjust the amplitude on the function generator to the desired level.
4. Adjust the frequency on the function generator until the waveforms on the oscilloscope are 90 degrees out of phase, with $v_{out}(t)$ lagging $v_{in}(t)$. This is the point of apparent resonance.
5. Measure and record the amplitudes of $v_{in}(t)$ and $v_{out}(t)$.
6. If not at ideal resonance, measure the time delay between homologous zero crossings in the two waveforms.
7. Record the frequency at which resonance was found and/or at which the measurements were taken.
8. If desired, capture the waveform data for $v_{in}(t)$ and $v_{out}(t)$.

When testing is finished, the following steps should be taken:

1. Set the function generator amplitude to the lowest value.

2. If applicable, turn off the cooling jet.

3. Turn off the RF amplifier.

REFERENCES

[1] Arpaia, P., Avallone, F., Baccigalupi, A., de Capua, C., and Landi, C., *The Measurement, Instrumentation, and Sensors Handbook*, chapter 39, pp. 39.1–39.34. Reading, Massachusetts: CRC Press, second ed., 10 January 2000. 5.4.2.2, 5.4.2.2, 5.4.3.1

[2] Bahl, C., Engelbrecht, K., Bjork, R., Eriksen, D., Smith, A., Nielsen, K., , and Pryds, N., "Design concepts for a continuously rotating active magnetic refrigerator," *International Journal of Refrigeration*, vol. 34, pp. 1792–1796, 2011. 7.3.2

[3] Baliga, J., *Silicon Carbide Power Devices*. Singapore: World Scientific, 2005. 2

[4] Barranger, J., *Hysteresis and Eddy Current Losses of a Transformer Lamination Viewed as an Application of the Poynting Theorem (Technical Note D-3114)*. Washington, D.C.: NASA, 1965. 5.5.5.4, 5.5.6.1, 5.5.6.3, 5.5.6.4, B.3.1

[5] Beidaghi, M. and Wang, C., "On-chip micro-power: Three dimensional structures for micro-batteries and super-capacitors," *Proceedings of the SPIE*, vol. 7679, no. G, pp. 1–7, 2010. 3.1.1.4, 38

[6] Benazzi, A., Brunet, M., Dubreuil, P., Mauran, N., Bary, L., Laur, J.-P., Sanchez., J.-L., and Isoird, K., "Performance of 3d capacitors integrated on silicon for dc-dc converter applications," in *2007 European Conference on Power Electronics and Applications*, EPE, (Aalborg, Denmark), pp. 1–8, September 2007. 3.1.1.2, 35, 4.2.2

[7] Blanter, M. S., Golovin, I. S., Neuha user, H., and Sinning, H.-R., eds., *Internal Friction in Metallic Materials*. Berlin: Springer-Verlag, 2007. 5.5.6.5

[8] Boon, C. and Robey, J., "Effect of domain wall motion on power loss in grain-oriented silicon steel," *Proceedings of the IEE*, vol. 115, no. 10, pp. 1535–1540, 1968. 5.5.6.1, 113, 5.5.6.1

[9] Brand, O. and Fedder, G. K., *Advanced Microsystems Volume 2: CMOS-MEMS*. Weinheim, Germany: Wiley, 2005. 1.3.1.4

[10] Broerson, P. M. T., "Estimation of the accuracy of mean and variance of correlated data," *IEEE Transactions on Instrumentation and Measurement*, vol. 47, pp. 1085–1091, October 1988. 5.4.3.1, 5.4.3.1

[11] Canham, L., ed., *Properties of Porous Silicon*. IET, 1997. 3.1.1.1, 33

[12] Chandrasekaran, S. and Mehrotra, M., "Matrix integrated magnetics for low voltage interleaved dc-dc converters," in *Eighteenth Annual IEEE Applied Power Electronics Conference and Exposition*, vol. 1 of *APEC 2003*, pp. 103–108, 2003. 12

[13] Chen, C.-H., Tsai, D.-S., Chung, W.-H., Lee, K.-Y., Chen, Y.-M., and Huang, Y.-S., "Electrochemical capacitors of miniature size with patterned carbon nanotubes and cobalt hydroxide," *Journal of Power Sources*, vol. 205, pp. 510–515, 2012. 3.1.1.5, 39

[14] Chen, H., Wei, B., and Ma, D., "Energy storage and management system with carbon nanotube supercapacitor and multidirectional power delivery capability for autonomous wireless sensor nodes," *IEEE Transactions on Power Electronics*, vol. 25, no. 12, pp. 2897–2899, 2010. 1.3.2, 7

[15] Chen, X., Gerapolous, K., Guo, J., Brown, A., Weng, C., Ghodssi, R., and Culver, J. N., "A patterned 3d silicon anode fabricated by electrodeposition onto a virus-structured current collector," *Advanced Functional Materials*, vol. 21, pp. 380–387, 2011. 3.1.1.6, 41

[16] Conway, B. E., *Electrochemical Supercapacitors: Scientific Fundamentals and Technological Applications*. Kluwer-Plenum, 1999. 4.1.1, 4.2.1.2

[17] Errien, N., Vellutini, L., Louarn, G., and Froyer, G., "Surface characterization of porous silicon after pore opening processes including chemical modifications," *Applied Surface Science*, vol. 253, pp. 7265–7271, 2007. 3.1.1.1

[18] Futurrex, Franklin, New Jersey, *Futurrex NR-21 Series Photoresist Datasheet*. 1.3.1.1

[19] Gad-el-Hak, M., ed., *MEMS Design and Fabrication*. Taylor and Francis, 2006. 1.3.1.1, 5, 1.3.1.3

[20] Gamburg, Y. D. and Zangari, G., *Theory and Practice of Electrodeposition*. Springer, 2011. 1.3.1.2

[21] Gardner, J. W., Varadan, V. K., and Awadelkarim, O. O., *Microsensors, MEMS and Smart Devices*. Wiley, 2001. 1.3.1.3

[22] Gerasopoulos, K., McCarthy, M., Royston, E., Culver, J. N., and Ghodssi, R., "Microbatteries with tobacco mosaic virus templated electrodes," in *IEEE MEMS*, IEEE MEMS, (Tucson, Arizona), pp. 960–963, January 2008. 3.1.1.6, 41

[23] Greenhouse, H. M., "Design of planar rectangular microelectronic inductors," *IEEE Transactions on Parts, Hybrids, and Packaging*, vol. PHP-10, no. 2, pp. 101–109, 1974. 2.1.1.2, 2.1.1.2, 2.1.1.3

[24] Han, Y., Cheung, G., Li, A., Sullivan, C. R., and Perreault, D. J., "Evaluation of magnetic materials for very high frequency power conversion," in *2008 IEEE Power Electronics Specialists Conference*, PESC, (Rhodes, Greece), pp. 4269–4276, June 2008. 5.4.3.2, 3

[25] Harris, R. D., Herbert, P. C., Knieser, M. J., Kretschmann, R. J., Lucak, M. A., Pond, R. J., Szabo, L. F., and Discenzo, F. M., "Microelectromechanical isolating power converter." Patent, 08 2007. US 7262522. 8, 1.3.2

[26] Hart, D., *Power Electronics*. New York, NY: McGraw Hill, 2011. 17, 18

[27] Hilt, T., Boutry, H., Franiatte, R., Rothan, F., Sillon, N., Stam, F., Mathewson, A., Wang, N., O'Mathuna, C., and Rodgers, K., "Dc/dc converter 3d assembly for autonomous sensor nodes," in *Proceedings 60th Electronic Components and Technology Conference*, ECTC, (September), pp. 834–839, September 2010. 6, 1.3.2

[28] Howard I. Becker, P. B. F., "Low voltage electrolytic capacitor." Patent, 7 1957. US 2800616. 4.2.1.2

[29] Hwang, Y., Gao, F., Hong, A. J., and Candler, R. N., "Porous silicon resonators for improved vapor detection," *JMEMS*, vol. 21, no. 2, pp. 235–242, 2012. 3.1.1.1, 34

[30] Jain, P. and Rymaszewski, E. J., "Embedded thin film capacitors – theoretical limits," *IEEE-TAdvPkg*, vol. 25, pp. 454–458, August 2002. 54

[31] Johari, H. and Ayazi, F., "High-densty embedded deep trench capacitors in silicon with enhanced breakdown voltage," *IEEE Transactions on Components and Packaging Technology*, vol. 32, pp. 808–815, December 2009. 3.1.1.2, 4.2.2

[32] J.T. Strydom, J., van Wyk, J., and Ferreira, J., "Some limits of integrated lct modules for resonant converters at 1 mhz," vol. 37, no. 3, pp. 820–828, 2004. 1.3.3.1, 10

[33] Kuhn, L. T., Pryds, N., Bahl, C. R. H., and Smith, A., "Magnetic refrigeration at room temperature – from magnetocaloric materials to a prototype," in *Joint European Magnetic Symposia*, JEMS, (Krakow, Poland), pp. 1–11, IOP, August 2010. 7.3.2

[34] Leblebici, D. and Leblebici, Y., *Fundamentals of High-Frequency CMOS Analog Integrated Circuits*. Cambridge University Press, 2009. 4.2.1.4

[35] Lee, K. J., Damani, M., Pucha, R. V., Bhattacharya, S. K., Tummala, R. R., and Sitaraman, S. K., "Reliability modeling and assessment of embedded *IEEE Transactions on Components and Packaging Technology*, vol. 30, no. 1, pp. 152–162, 2007. 1.3.3.2, 13

[36] Liang, Y., Liu, W., and van Wyk, J. D., "Design of integrated passive component for a 1 mhz 1 kw half-bridge llc resonant converter," in *2005 IEEE Industry Applications Conference and Fortieth IAS Annual Meeting*, vol. 3, (Hong Kong, China), pp. 2223–2228, October 2005. 9, 1.3.3.1

[37] Mandakis, B. J., "The solid tantalum capacitor - a solid contributor to reliability," in *11th Annual Reliability Physics Symposium*, vol. 1 of *Reliability Physics Symposium*, (Las Vegas, Nevada), pp. 45–53, April 1973. 4.2.1.3, 4.4.4

[38] Maniktala, S., ed., *Troubleshooting Switching Power Converters: A Hands-On Guide*. Elsevier Newnes, 2008. 4.1.2.1

[39] Mathúna, S., Ferreira, B., van Wyk, D., and van Wyk, J., "Towards a roadmap for power systems integration," in *11th European Conference on Power Electronics and Applications (EPE 2005)*, EPE, (Dresden, Germany), pp. 1–9, September 2005. 1.2.2

[40] Mayergoyz, I., *Nonlinear Diffusion of Electromagnetic Fields*. Academic Press, 1988. 115, 5.5.6.4

[41] Microchem, Newton, Massachusetts, *Megaposit SPR 220 Series Photoresists*. 1.3.1.1

[42] Nam, K.-T., Kim, D.-W., Yoo, P. J., Chiang, C.-Y., Meethong, N., Hammond, P. T., Cheng, Y.-M., and Belcher, A. M., "Virus-enabled synthesis and assembly of nanowires for lithium ion battery electrodes," vol. 312, pp. 885–890, 2006. 3.1.1.6

[43] Neidhofer, G., "Early three-phase power," *IEEE Power and Energy Magazine*, vol. 5, pp. 88–100, September/October 2007. 1.2.1

[44] Nongaillard, M., Lallemand, F., and Allard, B., "Design for manufacturing of 3d capacitors," *Microelectronics Journal*, vol. 41, pp. 845–850, 2010. 3.1.1.2, 36

[45] Owen, E., "SCR is 50 years old," *IEEE Industry Applications Magazine*, vol. 13, pp. 6–10, November/December 2007. 1.2.1

[46] Park, J.-W., *Core Lamination Technology for Micromachined Power Inductive Components*. PhD dissertation, Georgia Institute of Technology, Department of Electrical and Computer Engineering, May 2004. (document), 43, 3.3

[47] Paul B. Koeneman, P. B. F., "Micro-electromechanical voltage converter." Patent, 12 2004. US 6833645. 1.3.2

[48] Petersen, T. F., Pryds, N., Smith, A., Hattel, J., Schmidt, H., and Knudsen, H.-J. H., "Two-dimensional mathematical model of a reciprocating room-temperature active magnetic regenerator," *International Journal of Refrigeration*, vol. 31, no. May, pp. 432–443, 2008. 144, 7.3.2, 7.3.2

[49] Pimentel, A. and Fortunato, E., "Manganese nitrate impregnation cycles optimization by addition of surfactants on ta capacitors with high charge powders manufacturing," in *The 2010 European Capacitor and Resistor Technology Symposium (CARTS Europe)*, CARTS, (Munich, Germany), pp. 891–894, ECIA, November 2010. 4.2.1.3

[50] Pry, R. and Bean, C. P., "Calculation of the energy loss in magnetic sheet materials using a domain model," *Journal of Applied Physics*, vol. 29, no. 3, pp. 532–533, 1958. 5.5.6.1, 5.5.6.2, 114, 5.5.6.5

[51] Rajaraman, S., Bragg, J. A., Ross, J. D., and Allen, M. G., "Micromachined three-dimensional electrode arrays for transcutaneous nerve tracking," *JMM*, vol. 21, no. 085014, pp. 1–14, 2011. 1.3.1.1

[52] Richardson, F. R. and Falkowski, E. C., "Relation of ac losses to hysteresis losses in electric sheet steels," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-86, no. 9, pp. 1072–1078, 1967. 5.5.6.1

[53] Schaper, L. and Thomason, C., "High density double and triple layer tantalum pentoxide decoupling capacitors," *IEEE Transactions on Components and Packaging Technology*, vol. 30, pp. 563–568, December 2007. 4.2.2

[54] Scuderi, A., Ragonese, E., Biondi, T., and Palmisano, G., *Integrated Inductors and Transformers*. CRC, 2011. 22, 2.2.1.1

[55] Seet, H., Lee, X., Hong, M., Teh, K., and Teo, H., "Electrodeposition of ni-fe micro-pillars using laser drilled templates," *Journal of Materials Processing Technology*, vol. 192-193, pp. 346–349, 2007. 3

[56] Shah, U., "Development of mems power inductors with submicron laminations using an automated electroplating system," Master's thesis, Georgia Institute of Technology, Department of Electrical and Computer Engineering, December 2007. 3.3, C.1

[57] Shepard, J., "Power electronics futures," pp. 31–34 (vol 1), February 2004. 1.2.2

[58] Stone, G., Boulter, E. A., Culbert, I., and Dhirani, H., *Electrical Insulation for Rotating Machines*. 2004. 5.5.6.3

[59] Sun, G., Hur, J. I., Zhao, X., and Kim, C., "Fabrication of very-high-aspect-ratio micro metal posts and gratings by photoelectrochemical etching and electroplating," *JMEMS*, vol. 20, no. 4, pp. 876–884, 2011. 3.1.1.3, 3.1.1.3, 37

[60] Sun, J. and Mehrotra, V., "Orthogonal winding structures and design for planar integrated magnetics," *IEEETIndElec*, vol. 55, no. 3, pp. 1463–1469, 2008. 11, 1.3.3.2

[61] SynQor, Boxborough, Massachusetts, *PQ30 PowerQor DC/DC ConverterSeries Datasheet*. 1.3.3.2

[62] Tishin, A. M. and Spichkin, Y. I., *The Magnetocaloric Effect and its Applications*. Philadelphia, PA: IOP, 2003. 7.3.2, 7.3.2

[63] Trimmer, W. S. N., "Microrobots and micromechanical systems," *Sensors and Actuators*, vol. 19, pp. 267–287, September 1989. 3.1.1

[64] Tummala, R., *Fundamentals of Microsystems Packaging*. McGraw-Hill, 2001. 4.2.1

[65] Tura, A., Nielsen, K. K., and Rowe, A., "Experimental modeling and results of a parallel plate-based active magnetic regenerator," *International Journal of Refrigeration*, vol. 35, no. June, pp. 1–10, 2012. 144, 7.3.2

[66] Van den Bossche, A. and Valchev, V. C., *Inductors and Transformers for Power Electronics*. Taylor and Francis, 2005. 2.1.1.1, 16, 2.1.1.4, 2.2.1.1, 5.1

[67] Vishay Vitramon, Shelton, Connecticut, *VJ Non-Magnetic Series Datasheet*. 55

[68] Volz, S., *Microscale and Nanoscale Heat Transfer*. Springer, 2007. 5.5.6.5

[69] Waffenschmidt, E. and Jacobs, J., "Planar resonant multi-output transformer for printed circuit board integration," in *2008 IEEE Power Electronics Specialists Conference*, PESC, (Rhodes, Greece), pp. 4222–4228, June 2008. 14

[70] Wan, Y., Sha, J., Wang, L., and Wang, Y., "Influence of ambient gas on the growth kinetics of si nanocones," *Applied Physics Letters*, vol. 97, no. 15, pp. 3128–3133, 2010. 3.1.1.5, 40

[71] Wilson, T., "The evolution of power electronics," *IEEE Transactions on Power Electronics*, vol. 15, pp. 439–446, May 2000. 1

[72] Wu, K. C., *Switched-Mode Power Converters: Design and Analysis*. Elsevier, 2006. 2.1.2.1, 142, 7.3.1.1, 143, 7.3.1.2

[73] Y. Liang, Y., Lu, B., van Wyk, J., and Lee, F. C., "Coolmos fet/sic-diode module for high performance power switching," *IEEE Transactions on Power Electronics*, vol. 20, no. 3, pp. 679–686, 2005. 9

[74] Yazdani, A. and Irivani, R., *Voltage-Sourced Converters in Power Systems*. Hoboken, NJ: John Wiley & Sons, 2010. 1.2.1

[75] Yuferev, S. and Ida, N., *Surface Impedance Boundary Conditions: A Comprehensive Approach*. CRC, 2009. 3.1.1

[76] Zielenkiewicz, W. and Margas, E., *Theory of Calorimetry*. Philadelphia, PA: IOP, 2003. 5.4.2.1

VITA

William Preston Gallé III was raised in the town of Slidell, Louisiana, across Lake Ponchartrain from New Orleans. In 1997 he received the B.S. degree in electrical engineering from Texas A&M University in College Station, Texas. He spent the following seven years working in Portland, Oregon as an electrical engineer, designing digital hardware systems for capital equipment vendors.

In August 2004 he began graduate studies at the Georgia Institute of Technology, in the Walter Cox School of Electrical and Computer Engineering. He was awarded the Walter Cox Fellowship in 2004, the NDSEG Fellowship in 2005, and the TI:GER Fellowship in 2007. The TI:GER Fellowship was part of the TI:GER program, a two-year inter-professional entrepreneurship program operated by the GT B school.

His research interests include applied electrochemistry, microscale energy harvesting and power conversion, and computer-aided MEMS design and modeling.