Design of a Novel Test Bench for Induction Heating Load Characterization

by

Lisa Fernandez del Castillo

A.B. (Hons), Physics, University of Chicago (2012)

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

Author ..... Signature redacted .....

Department of Electrical Engineering and Computer Science

May 21, 2012

Signature redacted

Certified by .....

John G. Kassakian

Professor Emeritus of Electrical Engineering and Computer Science

Thesis Advisor

Signature redacted

Accepted by .....

Leslie A. Kolodziejski

Professor of Electrical Engineering and Computer Science

Chair, Committee on Graduate Students

Design of a Novel Test Bench for Induction Heating Load Characterization

by

Lisa Fernandez del Castillo

Submitted to the Department of Electrical Engineering and Computer Science

on May 21, 2014 in partial fulfillment of the requirements

for the degree of Master of Science

Abstract

Magnetic materials used in induction heating applications have nonlinear magnetic properties with respect to field strength and frequency, which can be effectively characterized using experimental techniques. To this end, we present a test bench inverter optimized for induction heating experimentation, capable of driving an inductive load across a 1–100 kHz frequency range with up to 2 kW power. Harmonic distortion of the inverter is minimized with a novel multilevel topology and modulation scheme, thus allowing near-sinusoidal excitations to be obtained at varying field strengths and frequencies. To demonstration the capabilities of the test bench, we characterize the power dissipation of a loaded induction heating coil across a range of frequencies and power levels.

Acknowledgments

This thesis would not have been possible without the support of a great many people.

First, I would like to express my deepest gratitude to my advisor, Professor John G. Kassakian, for his unfailing patience throughout my time here. His mentorship helped me not only to write this thesis but to develop as an electrical engineer. I am also so grateful to Dave Otten and Richard Zhang for their generous help with technical problems.

I would also like to thank all the wonderful people who have shaped my experience here at MIT: all the folks in LEES who have made doing research here so much fun, and everyone in the TCC who made MIT feel like home.

Finally, I would like to thank my family —my family at home in Chicago, my new family in Boston, and especially my wonderful husband— for all their love and support.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Nonlinear magnetic properties often increase the complexity of applications involving magnetic materials. Many power applications use magnetic materials, most frequently as inductor and transformer cores. In these applications, the component is typically designed to avoid nonlinear magnetic behavior, because nonlinear behavior tends to increase power loss in the magnetic material [1].

Induction heating is an interesting application because unlike inductor and transformer cores, induction heating uses the magnetic material as the load. This is exemplified by the most common domestic application of induction heating, the induction cooktop, as shown in Fig. 1.1. During induction heating, an ac current creates an oscillating magnetic field that couples to the magnetic material —the pan— and induces an electric field. The electric field gives rise to internal ac eddy currents that generate ohmic heating in the pan. Typical induction cooktops operate under conditions for which nonlinear magnetic behavior is assumed negligible [2, 3]. However, further research is required to better understand how nonlinearities affect power dissipation in domestic induction cooktop loads.

Experimental methods have great potential to increase our knowledge of the effects of magnetic load nonlinearities in domestic induction cooktops. At present, there is no general theory that can predict the magnetic properties of any material under any applied field. Experimentation, on the other hand, can provide reliable information about magnetic materials for any condition, provided we are capable of conducting the experiment. Specifically, experiments could allow for the characterization of the impedance of nonlinear magnetic materials with respect to different induction heating operation parameters.

To that end, in this thesis I provide an inverter test bench that is specialized for conducting domestic induction cooktop load characterization experiments. In structure, the setup is comparable to that of a commercial induction cooktop (see Fig. 1.1); it consists of a power electronics/control circuit that drives an ac current through the a planar, spiral-shaped coil which heats the applied pan via induction heating. However, the nature of the proposed load characterization experiments place requirements on the inverter that require custom design features.

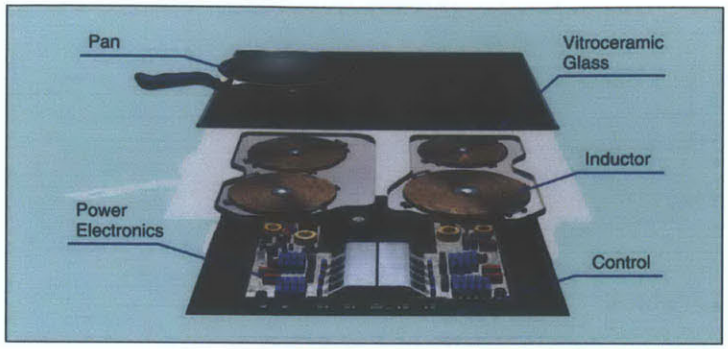

Figure 1.1: Typical domestic induction cooktop [4]. Our test bench has a similar structure, but our bench lacks the glass separator and consists of one coil instead of four.

1.1 Magnetic Nonlinearities

A material's magnetic properties are described by its magnetization curve, such as the magnetization curve of pure iron shown in Fig. 1.2, and different materials have different magnetization curves [5]. The magnetization curve can also be described in electrical terms; for an induction cooktop, the field intensity H is proportional to the applied current through the induction coil I, and the magnetic field B is proportional to the flux linkage λ that gives rise to power loss in the magnetic material.

Using either of these metrics, the degree of nonlinearity is expressed by the slope of the curves: the magnetic permeability $\mu = \frac{\mathrm{d}B}{\mathrm{d}H}$ in the B-H space, and the inductance $L = \frac{\mathrm{d}\lambda}{\mathrm{d}I}$ in the λ-I space. In linear materials, μ and L are constant, and in nonlinear materials μ and L are dependent on H and I (respectively).

The magnetization curve illustrates four significant behaviors. First, the material saturates: for large applied H, the change in the induced B approaches zero. Second, the material has saturation remanence: after being saturated by a strong applied H, a non-zero B remains when H is no longer applied. Third, the magnetization curve exhibits hysteresis: the B for a given H depends on the history of the material. Fourth, the shape of the magnetization curve is frequency dependent.

In domestic induction cooktops, the most common load material is cast iron, which is an alloy constituted of mostly iron with small (<5% by weight) amounts of silicon and carbon. Compared to pure iron, cast iron saturates at a lower B and a higher magnetic permeability $\mu = \frac{B}{H}$ (using a linear approximation) [1].[1]

For the purpose of load characterization experiments, it is important to note that the magnetic properties are dependent on both the field intensity and on the frequency. Since the field intensity is proportional to the supplied current, load characterization can be achieved by exciting the load with a range of currents and frequencies.

[1] Another difference is that cast iron has a higher resistivity, which increases the ohmic heating in the material due to eddy currents [1]. This is beneficial for induction heating.

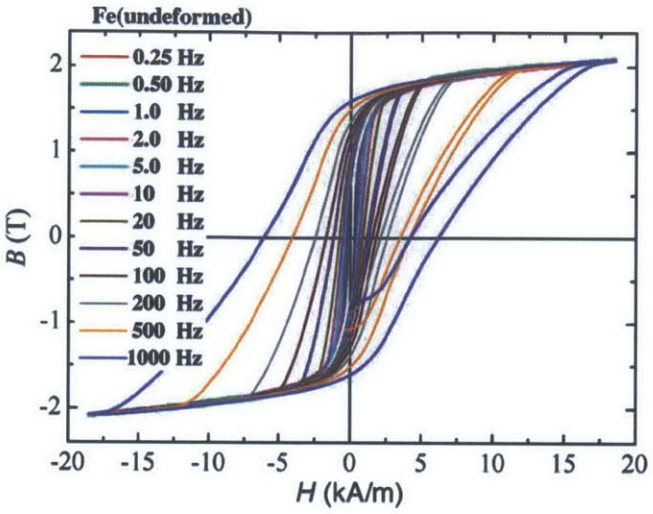

Figure 1.2: Magnetization curves of pure iron at a variety of frequencies [6].

1.2 Thesis Objective & Approach

Based on the above discussion, the test bench required for the purpose of domestic induction cooktop load characterization must be able to excite inductive loads with a range of currents and frequencies. This places three requirements on the inverter. First, it must provide current excitations at a single frequency —current excitations that are near-sinusoidal, having low harmonic distortion. Generating a sinusoidal current waveform is achieved by generating a sinusoidal voltage waveform; the inductance in the load provides low-pass filtering on the current, reducing the higher harmonic content. Second, the inverter must provide these near-sinusoidal voltage and current excitations across a range of frequencies. Third, the inverter must provide a range of current levels at each frequency. This is achieved by providing a range of voltages.

To summarize, the test bench needs to drive inductive loads with near-sinusoidal voltage and current excitations across frequency and power level ranges broad enough for characterization of induction heating load nonlinearities. Additionally, the frequency and power level ranges should include the standard operating conditions of domestic induction cooktops, which are approximately 20–100 kHz and from 50 W–3.5 kW [7, 8].

Meeting these objectives requires the design of a custom inverter test bench. Specifically, the requirement that the inverter provide both high power and frequency flexibility prevents the use of commercial inverters that satisfy only one of these parameters. The inverters typically used to drive induction cooktops are inadequate for this application; commercial induction cooktops use resonant inverters, which achieve very high efficiency when operating near the resonant frequency but are not designed to operate across a range of frequencies [7, 8]. Inverters that are optimized for frequency flexibilities, such as function generators and impedance analyzers, do not output the high level of power needed to observe the nonlinearities of induction heating loads [9].

It is necessary, then, to design a new test bench in order to conduct induction heating load characterization experiments, and the purpose of this thesis is to meet that need. At the core of our

test bench is a novel inverter topology that is optimal for induction heating load characterization experiments: a cascaded half-bridge full-bridge (CHBFB) inverter, a hybrid proposed by David Otten of the topology in [10]. In order to generate the required high power, frequency flexibility, and near-sinusoidal voltage excitations, the CHBFB has been designed to generate voltage waveforms with low harmonic distortion —including cancellation of all voltage harmonics that are multiples of 2, 3, and 5 to below 3% of the fundamental— across 1–100 kHz at up to 2 kW.

The organization of this thesis is as follows: Chapter 2 overviews the background knowledge relevant to this research. Relevant inverter topologies and control methods are discussed, and an overview of power electronics circuit design is provided. Chapter 3 covers the circuit design and implementation. The novel CHBFB topology is introduced, and specifics concerning the circuit design and component selection are described. Chapter 4 covers the control design and implementation. The basic program requirements and information flow are given, followed by the specifics of the control program. Chapter 5 presents the circuit performance. The test bench's ability to produce near-sinusoidal waveforms across the stated frequency and power range is assessed. Chapter 6 presents an initial load characterization conducted with this test bench. Chapter 7 concludes with the results of this research effort and suggestions for future work.

Chapter 2

Background

2.1 Power Circuit Topology Review

The inverter used in the test bench must generate near-sinusoidal waveforms across an inductive load across a wide frequency and at up to kilowatt power levels. During the topology selection process described in Chapter 3, Circuit Design & Implementation, several different types of inverters are considered. This section describes these inverters.

2.1.1 Bridges

One of the building blocks of many common inverter topologies is the bridge leg. There are four possible ways in which a bridge leg can be configured (see Fig. 2.1): high (in which the high-side switch is on, and the low-side switch is off), low (in which the high-side switch is off, and the low-side switch is on), open (in which neither switch is on), or shorted (in which both switches are on).

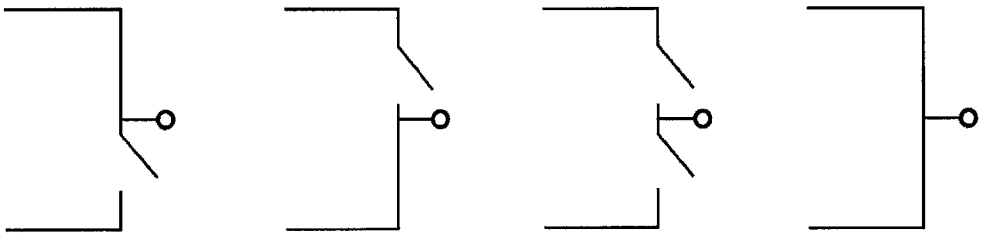

(a) High (b) Low (c) Open (d) Shorted

Figure 2.1: Possible bridge leg configurations.

Two of the simplest inverters are half-bridges, which consist of one bridge leg, and full-bridges, which consist of two bridge legs. For both half- and full-bridges, normal operation consists of switching between high and low bridge leg configurations. These topologies and their achievable output voltage levels are shown in Fig. 2.2. Because the half-bridge consists of a single bridge leg, it can only achieve two voltage levels: $+V_{dc}$ when the bridge leg is high, and 0 when the bridge is

low. The full-bridge can achieve three voltage levels: $+V_{dc}$ and $-V_{dc}$ when the bridge legs are in opposite configurations (one high and the other low), and 0 when the bridge legs are in the same configuration (both high or both low).

Typically, the shorted and open bridge leg configurations are not used. The shorted bridge leg configuration is prohibited because that shorts the power supply, a phenomenon called shoot through. Additionally, the open configuration is avoided, because that leaves the output terminal floating. However, in order to avoid shoot through, it is common to implement a brief period in which the bridge leg is open during transitions between high and low configurations. If the load is reactive, it is necessary to provide a current path during this period. This is commonly done by placing diodes across the switches, such as MOSFETs' body diodes, as shown in Fig. 2.3.

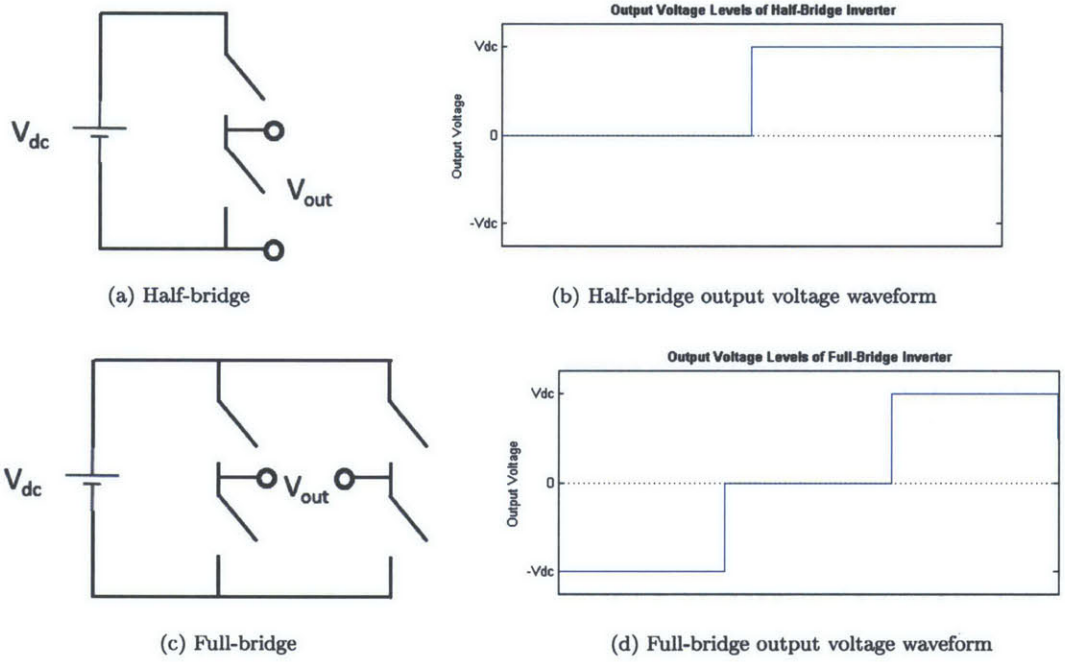

(a) Half-bridge

(b) Half-bridge output voltage waveform

(c) Full-bridge

(d) Full-bridge output voltage waveform

Figure 2.2: Half- and full-bridges and their associated waveforms.

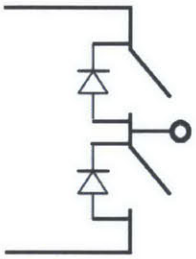

Figure 2.3: Diode placement used to provide a current path for reactive loads.

2.1.2 Resonant Inverters

Resonant inverters are a type of inverter in which the power output and efficiency are controlled by the frequency [1]. Although there are many different types of resonant inverter, the basic topology consists of a bridge leg that generates a square voltage waveform across a resonant network. When driving a reactive load, the load's reactance can be incorporated into the resonant network.

Resonant inverters are operated at maximum power output and efficiency at the undamped natural frequency, or resonant frequency, of the resonant network. Resonance is achieved at the frequency (ω_0) where the inductance (L) cancels the capacitance (C):

$$j\omega_0 L + \frac{1}{j\omega_0 C} = j\left(\omega_0 L - \frac{1}{\omega_0 C}\right) = 0 \tag{2.1}$$

$$\omega_0 = \frac{1}{\sqrt{LC}} \tag{2.2}$$

This is the operating point of maximum power output and efficiency. As the frequency increases or decreases away from the resonant frequency, the power output and efficiency decrease. Near the resonance frequency, the dependence of the power on the frequency $|\frac{dP}{d\omega}|$ (i.e. how quickly the power drops as the frequency is moved away from resonance) is determined by the quality factor (Q) of the resonant network. A high Q corresponds to a high $|\frac{dP}{d\omega}|$. Q is also the measure of the power gain at resonance; a high Q corresponds to a high current, and hence high power and efficiency, at resonance.

Resonant inverters have several key beneficial traits:

- **High efficiency.** If operated near the resonant frequency, resonant inverters are able to obtain very high efficiencies.
- **Power control.** Since the power output is sensitive to frequency, adjusting the frequency can be used as a means to control the power. The ability to make fine adjustment of the power level depends on Q; fine adjustments are easier to obtain if Q is small.
- **Low switch stress.** Resonant inverters have two characteristics that limit the stress placed on the semiconductor switches. First, because they generate sinusoidal waveforms, resonant inverters can achieve soft switching [11]. Second, the switching frequency is the carrier frequency; it is not necessary to use a high-frequency switching scheme such as PWM which causes increased switching losses in the devices. The low level of stress on the switches makes the resonant inverter topology suitable for high frequency applications.

The obvious disadvantage of resonant inverters is that they have limited frequency flexibility. Resonant inverters do not operate well outside a given frequency window surrounding the resonance frequency.

2.1.3 Multilevel Inverters

Multilevel inverters are a family of inverters that produce stepped output voltage waveforms with three or more levels.

There are three standard types of multilevel inverters: diode-clamped, flying capacitor, and cascaded H-Bridge (CHB) as shown in Fig. 2.4. For the purpose of comparison, all three circuits in Fig. 2.4 are five-level inverters: an example of an output voltage waveform is shown in Fig. 2.4d.

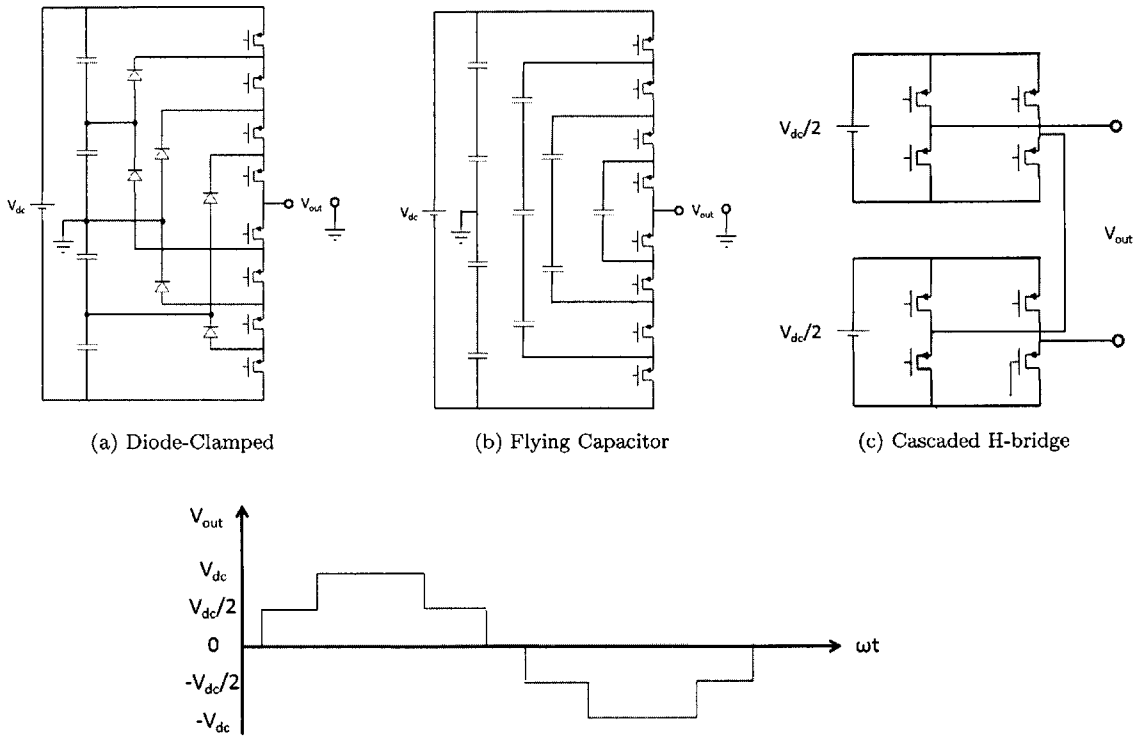

(a) Diode-Clamped (b) Flying Capacitor (c) Cascaded H-bridge

(d) Five-level Output Voltage Waveform

Figure 2.4: Five-level inverters using standard multilevel inverter topologies and an example five-level output voltage waveform [12].

The diode-clamped inverter consists of a ladder of switches held up by bulk capacitors and blocking diodes (see Fig. 2.4a). The dc voltage is distributed evenly across the stack of bulk capacitors, and the switches are connected to the stack at different voltage levels. The output voltage, then, depends on the on/off configurations of the switches. In the five-level example, the bulk capacitors across the power supply each have a voltage of $V_{dc}/4$ across them, such that the possible output voltages are $+V_{dc}$, $+V_{dc}/2$, 0, $-V_{dc}/2$, and $-V_{dc}$.

The flying capacitor inverter has a similar topology to the diode-clamped inverter (see Fig. 2.4b). However, instead of using blocking diodes, it uses additional storage capacitors to uphold the voltage level between switches. Like the diode clamped inverter, the output voltage depends on the on/off configurations of the switches and in the five-level example, the possible output voltages are $+V_{dc}$, $+V_{dc}/2$, 0, $-V_{dc}/2$, and $-V_{dc}$.

The CHB inverter is made of several single-phase, H-Bridge inverters connected in series (see Fig. 2.4c). Each H-bridge can produce at its output $+V_{dc}/2$, 0, or $-V_{dc}/2$, depending on the state of the switches. When several single-phase inverters are connected in series, the output waveform is

the sum of the outputs of each individual inverter.

All multilevel inverter topologies share certain advantageous characteristics:

- **Low harmonic content.** Multilevel inverters generate output voltages with inherently low harmonic distortion; the total amount of harmonic distortion produced by a multilevel inverter decreases with the number of voltage levels [13]. This allows for the production of near-sinusoidal voltage waveforms even before any external filtering.
- **Frequency flexibility.** The output voltage waveform is entirely determined by the on/off states of the semiconductor switches in the inverter; the number of output voltage levels, duration of levels, and waveform frequency are all controllable parameters. The frequency range, therefore, is limited only by the frequency range of the switches.
- **Switching flexibility.** Multilevel inverters can be controlled using either fundamental-frequency switching methods or high-frequency switching methods depending on the requirements of the application [14].

However, the differences between the topologies give rise to advantages and disadvantages unique to specific topologies.

The diode-clamped inverter has two major drawbacks. The first is that it requires a large number of components per output voltage level compared to the CHB. This is illustrated by the example, in which eight switches, four capacitors, and six diodes of varying voltage rating are required to generate five output voltage levels. The second problem is bulk capacitor voltage unbalance. When creating a generic voltage waveform, the bulk capacitors do not necessarily have the same charging time, which causes them to charge to different voltages [12]. Because of this, the desired output voltage waveform is not produced by the inverter. Various approaches have been proposed to solve this problem, such as replacing the bulk capacitors by a controlled constant dc voltage source [12] and using an external switching circuit to balance the voltages [15].

The flying capacitor inverter has similar disadvantages to the diode-clamped inverter. Like the diode-clamped inverter, the flying capacitor inverter requires a large number of components per output voltage level compared to the CHB. In the example, eight switches and ten capacitors are required to generate five output voltage levels. In addition, voltage unbalance is a bigger problem in this topology than in the diode-clamped inverter because there are more capacitors that need to be balanced. This is countered by the fact that the flying capacitor inverter has more flexibility in the switching scheme than the diode-clamped inverter. Therefore, it is possible to balance the capacitors by using the proper switching scheme, but the design of the switching scheme becomes very complicated for inverters with many voltage levels [12, 13].

The CHB has several advantages over flying capacitor and diode-clamped inverters. It requires the fewest components per output voltage level [16]. Additionally, it does not require voltage balancing [12]. The devices do not need to be rated for the full output voltage, and its modular structure simplifies repair and bypassing of faulty devices [16].

The main disadvantage of the CHB topology is that it requires multiple isolated voltage sources. This prohibits the use of many benchtop dc power supplies, which can exhibit problems with isolation due to parasitic capacitance across the transformers. This capacitance couples the output ground to

the supply line ground. Therefore, additional measures must be taken in order to ensure isolation, such as adding transformers between the voltage supply and inverting circuit [17, 18].

2.2 Switching Method Review

The switching method determines how a given power circuit topology is operated, and therefore it determines what output waveforms are generated. For this application, the switching method should generate a voltage waveform with low harmonic distortion. Selection of the optimal switching method depends on the circuit topology and application parameters. This section describes common switching methods for multilevel inverters that may be considered for this application.

2.2.1 Overview

There are two categories of switching methods commonly used with multilevel inverters: high-frequency switching methods and fundamental-frequency switching methods [13].

High-frequency switching methods, which are based on either pulse width modulation (PWM) and space vector algorithms, are most suitable for low-frequency applications. While these methods can produce a very accurate sinusoidal waveform, they are difficult to use in high frequency applications because they require the switching frequency to be much higher than the carrier frequency. PWM, for example, generally requires a switching frequency of at least ten times the carrier frequency. This is problematic because a switching frequency that high places additional stress on the devices that causes increased device power loss, requires additional heat sinking, and creates a higher risk for device breakdown.

On the other hand, fundamental-frequency switching methods require switching frequencies that are on the same order of magnitude as the carrier frequency, which is more favorable for high-frequency appliations. The three common varieties of fundamental-frequency switching methods are space vector control, nearest level control, and harmonic cancellation. Space vector and nearest level control methods both approximate the desired output voltage as the nearest achievable switching vector or output voltage achievable by the inverter. The harmonic cancellation method manipulates Fourier properties to minimize harmonic distortion in the output voltage waveform. Harmonic cancellation generally produces the lowest total harmonic distortion for inverters with five or fewer voltage levels, and space vector control and nearest level control work best for inverters that generate many output voltage levels [14].

2.2.2 Harmonic Cancellation

Harmonic cancellation is used to eliminate selected harmonics from the output voltage waveform of an inverter. In the stepped waveforms produced by multilevel inverters, the amplitudes of the harmonics inherently decay with increasing harmonic number. Therefore, in order to obtain an output voltage waveform from a multilevel inverter that is as sinusoidal as possible —that is, that has the lowest harmonic distorion as possible— it is necessary to cancel the lowest harmonics. Harmonic cancellation is typically only used for inverters with five or fewer levels because the number of equations that need to be solved to implement harmonic cancellation increases with the number of

voltage levels [14]. This section considers a five-level inverter, for which it is possible to cancel all harmonics that are multiples of 2, 3, and 5 (all 2n, 3n, and 5n harmonics).

The prinicipal behind harmonic cancellation is that every waveform can be decomposed into a sum of sinusoids:

$$f(t) = \frac{b_0}{2} + \sum_{n=1}^{\infty} a_n \sin(n\omega t) + \sum_{n=1}^{\infty} b_n \cos(n\omega t) \tag{2.3}$$

where

$$a_n = \frac{2}{T} \int_0^T f(t) \sin(n\omega t) \tag{2.4}$$

$$b_n = \frac{2}{T} \int_0^T f(t) \cos(n\omega t) \tag{2.5}$$

Based on this property, some useful generalities may be observed. First, if the waveform f(wt) is odd, then

$$b_n = \frac{2}{T} \int_0^T f_{odd}(t) \cos(n\omega t) = 0 \tag{2.6}$$

and so all odd functions have no cos harmonics. Similarly, all even functions have no sin harmonics. Finally, it can also be shown that half-wave symmetric waveforms have no even harmonics (a_{2k} and $b_{2k} = 0$ for all k). Therefore, in order to cancel all cos harmonics and all even sin harmonics from the output voltage, we will generate an odd half-wave symmetric waveform. The resulting waveform is composed solely of a_1, a_3, a_5, etc.

Harmonic cancellation of a given harmonic (a_n or b_n) is achieved by setting the durations of the voltage levels such that the amplitude of that harmonic across the cycle equals zero. For example, consider the three-level waveform shown in Fig. 2.5. Because it is odd and half-wave symmetric, all cos and even harmonics are cancelled. The amplitude of the nth harmonic is

$$a_n = \frac{2}{T} \int_{\delta}^{180^\circ - \delta} V \sin(n\omega t)\, \mathrm{d}t - \frac{2}{T} \int_{180^\circ + \delta}^{360^\circ - \delta} V \sin(n\omega t)\, \mathrm{d}t \tag{2.7}$$

If δ is set such that the integrals are equal to zero, then $a_n = 0$ as well as all multiples of a_n ($2a_n$, $3a_n$, etc.). Since an odd, half-wave symmetric waveform, is already free of all even harmonics, the largest present harmonic is the third. This harmonic can be cancelled by setting $\delta = 30^\circ$. In this case, all cos harmonics, even harmonics, and harmonics that are multiples of 3 (3, 6, 9, etc.) are cancelled, and we can write the waveform as:

$$f(t) = \sum_{n=1,5,7,11,\ldots}^{\infty} a_n \sin(n\omega t) \tag{2.8}$$

The two extra levels of a five-level inverter allow for the additional cancellation of the fifth

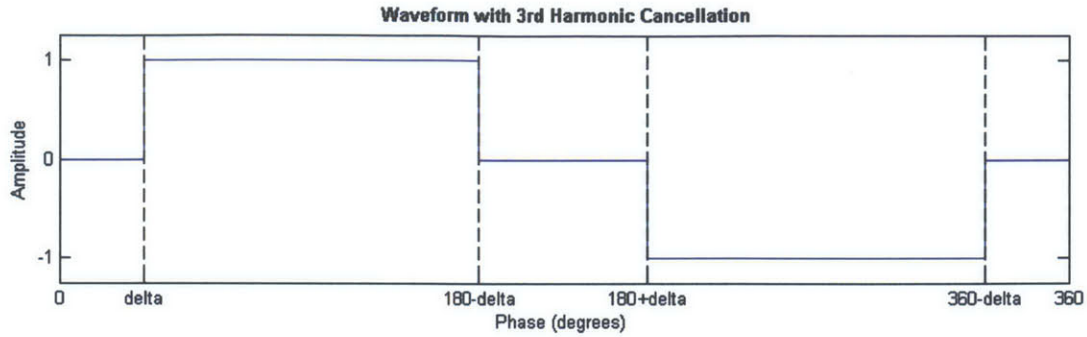

Figure 2.5: Three-level voltage waveform with cos, 2n, and 3n harmonics cancelled.

harmonics (5, 10, 15, etc.). Although the waveform that provides for cancellation of all cos, 2n, 3n, and 5n harmonics can be calculated directly using the method shown above, the extra voltage levels complicate the procedure. However, the required waveform can easily be derived by considering the five-level waveform as the sum of two time-shifted versions of the three-level waveform in Fig. 2.5. Like the original, the time-shifted waveform has no third harmonic; a linear operation (time-shifting) cannot reintroduce a harmonic that does not exist in the original waveform. This is shown in Fig. 2.6.

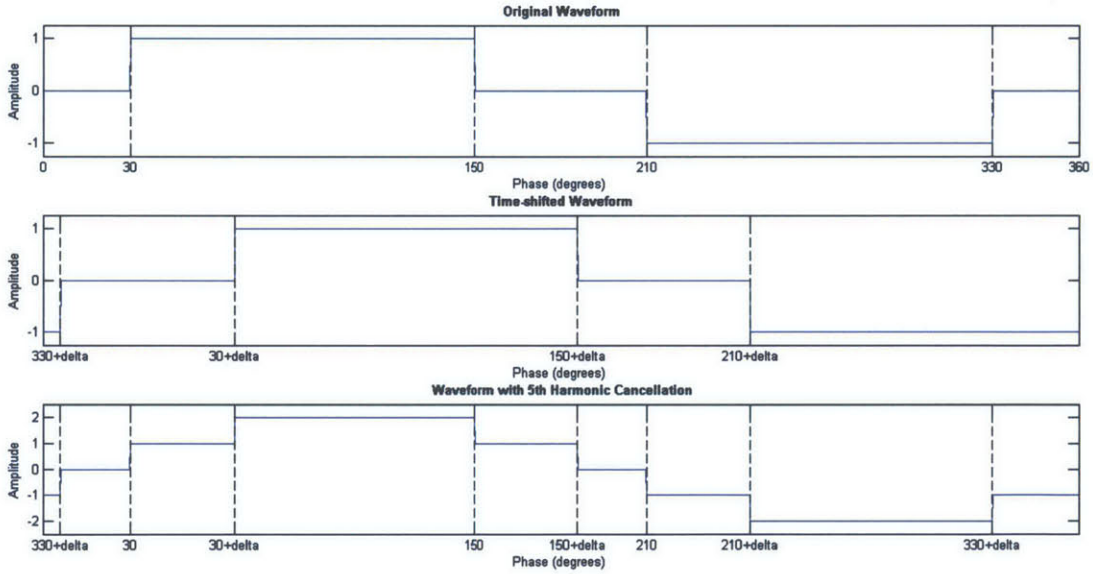

Figure 2.6: Time-shifted waveforms with cos, 2n, and 3n harmonics cancelled; and the summed waveform with cos, 2n, 3n, and 5n harmonics cancelled.

The original waveform can be written as the sum in (2.8), and therefore the time-shifted waveform can be written

$$f(t-\Delta t) = \sum_{n=1,5,7,11,\ldots}^{\infty} a_n \sin\left(n\omega(t-\Delta t)\right) = \sum_{n=1,5,7,11,\ldots}^{\infty} a_n \sin\left(n\omega t - n\omega\Delta t\right) \tag{2.9}$$

Therefore, the nth harmonic of the sum of the two waveforms is

$$f_n(t) + f_n(t - \Delta t) = a_n \sin(n\omega t) - a_n \sin(n\omega t - n\omega \Delta t) \tag{2.10}$$

Cancellation of the nth harmonics is achieved when $n\omega\Delta t = 180°$, such that

$$f_n(t) + f_n(t - \Delta t) = a_n \sin(n\omega t) - a_n \sin(n\omega t - n\omega \Delta t) = 0 \tag{2.11}$$

Therefore, cancellation of the 5th harmonics is achieved by setting $\omega\Delta t = 36°$. This waveform is shown in Fig. 2.7.

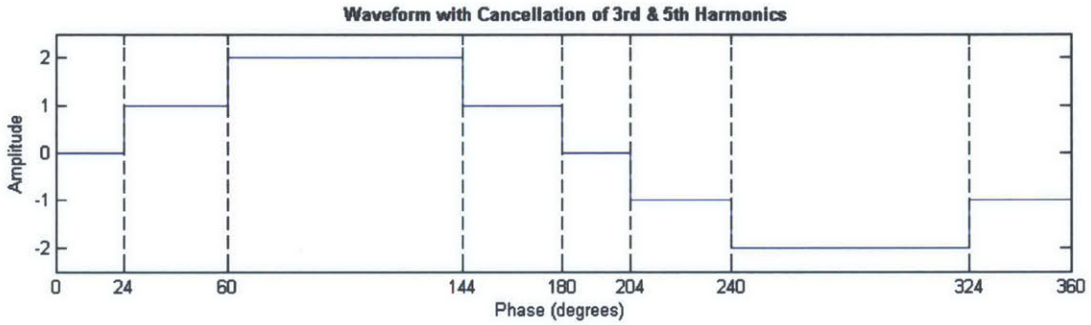

Figure 2.7: Waveform with no cos, 2n, 3n, or 5n harmonics.

The harmonic content of this waveform is shown by taking the Fourier transform, which is shown in Fig. 2.8. The 2n, 3n, and 5n harmonics are cancelled to below 10^{-16} of the fundamental, while the 7th and 11th harmonic are both about 10^{-1} of the fundamental (see Fig. 2.8). The large value of the 7th and 11th harmonics is expected because the switching scheme was not designed to cancel them. The non-zero magnitudes of the 2n, 3n, and 5n harmonics is due to the machine epsilon of the computation: 10^{-16} is the rounding error for floating point numbers. Therefore, these values are approximately zero, confirming that the ideal waveform achieves cancellation of all 2n, 3n, and 5n harmonics.

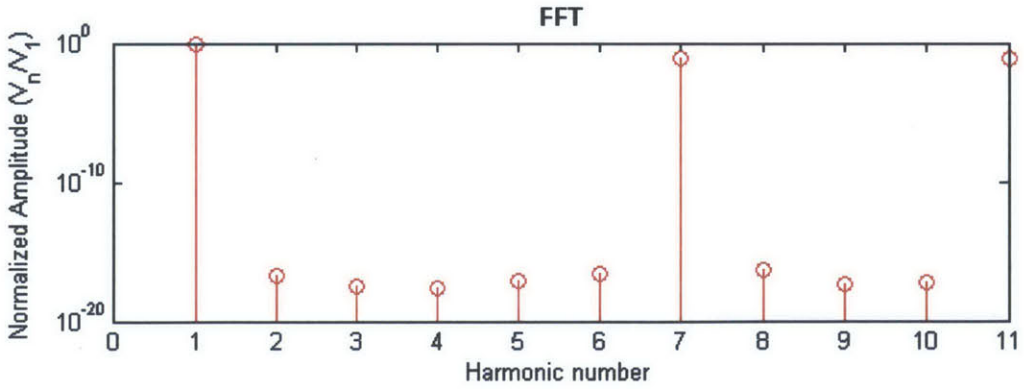

Figure 2.8: Fourier transform of the ideal voltage waveform, normalized such that the amplitude of the fundamental is 1.

2.3 Power Electronics Circuit Design

Merely choosing the correct circuit topology and switching method is not enough to create a successful inverter. Careful design of the circuit board is also crucial to building a well-functioning inverter. This section reviews two key design aspects that must be considered when designing a power electronics circuit: the switches and the driving circuit.

2.3.1 Switch Design

The switches are one of the most important components of any power electronics circuit. They are the devices that allow the circuit to operate in different modes, without which the circuit could not function as a power dc-ac inverter. Switch design consists of three processes: selection of the switch type, selection of the switch, and design of the thermal system needed to keep the switches from overheating.

2.3.1.1 Switch Type

Determining which type of switch to use is a process of comparing the specific application needs with the capabilities of the different types of switches.

There are two categories of switches: passive and active. A passive switch, such as a diode, cannot be user-controlled but is controlled by conditions in the circuit. An active switch, such as a transistor, are user-controllable. In power applications, the most commonly used transistors are the power metal–oxide–semiconductor field-effect transistor (power MOSFET) and the insulated-gate bipolar transistor (IGBT). For MOSFETs and IGBTs, the voltage at the gate determines whether the drain-source and collector-emitter channels (respectively) conduct current or block voltage.

Due to structural differences, power MOSFETs and IGBTs are preferrable under different conditions.[1] Power MOSFETs are more commonly used for applications requiring low voltages ($<\approx$200 V), low powers ($<\approx$500 W), high frequencies ($>\approx$200 kHz), long duty cycles, and big load variations. IGBTs, on the other hand, are typically used for applications requiring high voltages ($>\approx$1000 V), high powers ($>\approx$5000 W), low frequencies ($<\approx$20 kHz), short duty cycles, small load variations, and high temperatures [20].

In addition, both MOSFETs and IGBTs can be categorized by their dominant current carrier: electrons (N-type) or holes (P-type). Electrons have a higher mobility than holes. Since the resistivity of a material is inversely proportional to the carrier mobility, N-type devices tend to have lower resistivities and are more commonly used.[2]

2.3.1.2 Switch Device[3]

The switches must provide good performance while meeting the voltage and current ratings of the application. The performance of the switches can be measured with respect to three criteria: on-

[1] A thorough discussion of the physical differences between the power MOSFET and IGBT is not necessary for this work, but one can be found in [19].

[2] Again, an in-depth comparison of the n-type and p-type transistors is not necessary for this work, but one can be found in [19].

[3] N-type Power MOSFETs are considered throughout the rest of this chapter, since these are the devices used in the test bench.

state resistance, switching speed, and thermal characteristics. All of these criteria are described by parameters that are included in devices' data sheets, making it easy to select the device that is optimal for a specific application.

It is critical to have a low on-state resistance in order to minimize power losses during conduction. The on-state resistance is the resistance of the switch when it is conducting current. This device characteristic is dependent on many operation parameters such as the device's gate-source voltage, drain-source current, and temperature.

In addition, having a high switching speed is important for two reasons. First, when the switch is transitioning between the on- and off-states it is partially conducting, such that it has both a voltage and current across it, resulting in power losses in the device. By maximizing the switching speed of the device, the switching loss is minimized. Second, the devices selected must be capable of switching at the required frequency of the application. The device's rise and fall times specify how quickly its drain-source voltage can change between the on- and off-states. In addition, the device's gate charge determines how quickly it responds to control signals at its gate. A switch with a small gate charge will charge and discharge more rapidly than a switch with a large gate charge, allowing it to respond more quickly to control signals.

Finally, devices with favorable thermal characteristics place low requirements on the thermal system. The two key thermal parameters of a MOSFET are its maximum junction temperature and the thermal resistance. The maximum junction temperature is the maximum temperature that the interior of the switch device can reach without breaking down. The thermal resistance measures how much the temperature of the device increases when it generates heat. It is beneficial to use devices with high maximum junction temperatures and low thermal resistances.

2.3.1.3 Thermal Considerations

Thermal considerations play an important part in the switch design. During operation, each MOSFET dissipates power, causing its temperature to rise. If the temperature increases beyond its rating, the device breaks down. In order to avoid this, a heat sink must be attached to each device to keep it cool.

The loss in a MOSFET can be approximated as the sum of the conduction losses and the switching losses [1]:

$$P_{diss} \approx \text{conduction losses} + \text{switching losses} \approx \left(I_{rms}^2 R_{ds,on}\right) + \left(\frac{1}{2} I_{rms} V_{DS} f(t_r + t_f)\right) \qquad (2.12)$$

where I_{rms} is the rms drain current, $R_{ds,on}$ is the on-resistance, V_{DS} is the drain-source voltage, f is the frequency, and t_r and t_f are the rise and fall times of the drain-source voltage.

The heat resulting from the power dissipation in the MOSFET is removed via an attached heat sink. The heat sink is usually attached to the MOSFET package using thermally-conductive paste or a thermally-conductive pad. The necessary cooling is described by the thermal circuit in Fig. 2.9. In this figure, P_{diss} is the power dissipated by the MOSFET, T_j is the temperature of the MOSFET junction, R_{jc} is the thermal resistance between the MOSFET junction and case, R_{cs} is the thermal resistance between the MOSFET case and the heat sink, R_{sa} is the thermal resistance between the heat sink and the ambient air, and T_a is the ambient temperature. A standard design practice

is to leave a 20°C safety margin between the maximum rated device junction temperature and the maximum temperature reached during device operation. Therefore, the maximum allowable thermal resistance of the heat sink can be calculated:

$$R_{\mathrm{sa}} = \frac{(T_j - 20°\mathrm{C} - T_{\mathrm{a}})}{P_{\mathrm{diss}}} - R_{\mathrm{jc}} - R_{\mathrm{cs}} \tag{2.13}$$

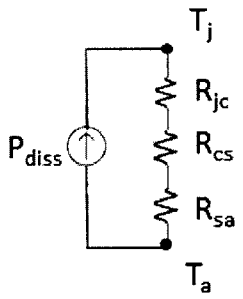

Figure 2.9: Thermal circuit describing the heat caused by power dissipation in MOSFETs.

2.3.2 Driving Circuit Design

The driving circuit consists of the devices that ensure that the switches turn on and off correctly. The key processes involved in driving circuit design are selection of the high-side driving method, driver selection, and gate circuit design.

2.3.2.1 High-side Driving Method

Special techniques must be employed to drive the high-side N-type MOSFETs. The gate voltage must be higher than the drain voltage when the MOSFET is in the linear region. This is not an issue for low-side MOSFETs because their sources are tied to circuit ground, and when in the linear region their drain voltage is low. However, the source voltage of high-side MOSFETs is elevated relative to circuit ground —it is tied to the drain of the low-side MOSFET— and therefore the drain voltage can often be too high to feasibly allow for a ground-referenced gate signal.

One of the most popular ways to drive high-side N-type MOSFETs is charge pumping. A charge pump is a type of dc-dc converter that uses energy storage elements to generate voltage pulses of higher magnitude than the gate drive supply voltage. Although there are many varieties of charge pump circuits, the most common type of charge pump used for high-side driving is the bootstrap circuit. The basic bootstrap circuit is shown in Fig. 2.10. When the driver provides a low signal, it ties the gate of the high-side MOSFET to its source. This turns off the MOSFET. During this stage of the cycle, the bootstrap capacitor charges. During the second stage of the cycle, the driver provides a high signal, tying the gate of the high-side MOSFET to the charged bootstrap capacitor. During this stage of the cycle, the bootstrap capacitor discharges to drive the high-side MOSFET, and the magnitude of the gate signal is the voltage across the bootstrap capacitor.

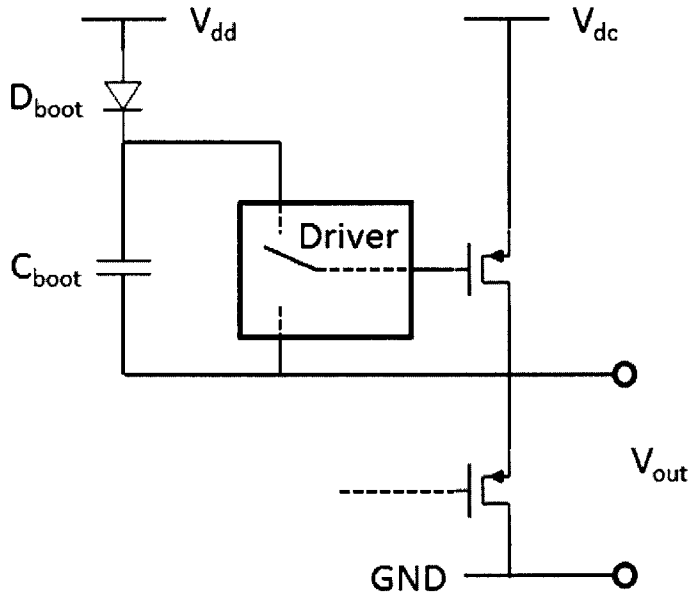

Figure 2.10: Bootstrap circuit.

The bootstrapping method of driving the high-side switch relies upon proper selection of the bootstrap diode and capacitor. The range of adequate capacitance of the bootstrap capacitor is bounded by two criteria: it must be big enough to provide the required gate voltage, and it must be small enough to fully charge when the high-side switch is turned off [21].

The minimum capacitance required to ensure that the bootstrap capacitor can drive the MOSFET gate high is determined by the relation $C = \frac{Q}{V}$:

$$C_{BS,min} = \frac{Q_G + Q_{RR} + \frac{I_{DR}+I_{QB,S}}{f}}{\triangle V} \tag{2.14}$$

where $C_{BS,min}$ is the minimum required bootstrap capacitance, Q_G is the high-side MOSFET's gate charge, Q_{RR} is the bootstrap diode's reverse recovery charge, I_{DR} is the bootstrap diode's reverse leakage current, $I_{QB,S}$ is the upper supply quiescent current, f is the minimum switching frequency, and $\triangle V$ is the allowed change in voltage across the capacitor during discharge.

On the other hand, the bootstrap capacitor must be small enough so that it can be fully charged up during the off-time of the high side MOSFET. Again, this is determined by the relation $C = \frac{Q}{V}$:

$$C_{BS,max} = \frac{I_{max} \times \triangle t_{min}}{\triangle V} \approx \frac{I_{max}}{\Delta V \times f_{max}} \tag{2.15}$$

where I_{max} is the maximum current supplied by the bootstrap diode; $\triangle t_{min}$ is the shortest amount of time the capacitor will have to charge, or $\approx \frac{1}{f_{max}}$; and again, $\triangle V$ is the allowed change in voltage across the capacitor during charge up (during steady state operation, the change in capacitor voltage during discharge equals the change in capacitor voltage during charge up).

Based on the above discussion, there are two requirements for the bootstrap diode. First, it should have low reverse recovery charge and leakage current to minimize the total charge that the capacitor needs to provide. Second, it should have a high maximum current to facilitate fast charging of the bootstrap capacitor.

2.3.2.2 Driver

The driver is the main component of the driving circuit. It receives the low-current digital control signals and delivers them to the MOSFET gates with enough current to produce clean on/off transitions.

The selected driver must be able to meet several requirements. First, it must be able to source and sink enough current to turn the MOSFETs on and off at the required speed. It must provide gate signals that are high enough to bias the MOSFET into the linear region where the on-state resistance is low. It must generate the desired gate signals based on input control signals. It must isolate the signal to the high-side switch to have a ground reference of the switch's drain. Finally, it may have a mechanism for preventing shoot through —a spike of current "shooting through" the circuit due to a shorted power supply, which usually occurs when a bridge leg is shorted briefly during the transition between high and low configurations.

Shoot through prevention is an optional feature for gate drivers because it may also be achieved in the controls. The advantages and disadvantages of preventing shoot through with the driver versus in the controls is case specific based on the characteristics of the particular driver and control method used. Many commercial drivers and control devices have built-in shoot through prevention that are reliable and easy to implement. When using a device with built-in shoot through prevention, it is often simplest to use this option rather than implementing it using external circuit devices or an additional control program.

There are many options for selecting a driver to meet these requirements. The first option is to build a custom driver. This option provides the most control over the driving circuit. The chief drawback is the additional design complexity; driver design is not a trivial task, especially if features such as shoot through prevention or undervoltage protection are incorporated. Unless the application has a non-standard driving requirement, it is generally simpler to use a commercially available driver.

Among commercial drivers, there remain many options to choose from. Particularly, there are two categories of commercial drivers: single-output drivers, for driving high-side or low-side switches, and multi-output drivers, for driving both low- and high-side switches. The main advantage of using single-output drivers is the ability to achieve more compact driving circuits, which improves switching performance. On the other hand, the benefit of using multi-output drivers is the ability to drive an entire bridge using a single driver chip. Most commercially available drivers use the bootstrapping method to drive the high-side switch.

2.3.2.3 Gate Circuit

Finally, conditioning of the gate signal is required in order to achieve clean load voltage transitions. This is because the load voltage waveform depends on the drain-source waveforms across the MOS-

FET switches, and each MOSFET's drain-source waveform depends on its gate-source waveform. Fast drain-source high/low transitions cannot be achieved without fast gate-source transitions, and the amount of drain-source overshoot and ringing is strongly dependent on the amount of gate-source overshoot and ringing. In addition, it is necessary to have fast MOSFET turn-off transitions in order to avoid shoot through.

Therefore, the gate signal is required to have acceptable gate-source rise and fall times and drain-source overshoot. The speed of the gate-source transition is typically limited by how quickly the gate current can charge or discharge the gate capacitances [22]. Although this is predominantly determined by the driving circuit, devices in the gate path of the MOSFET can alter gate-source rise and fall times. Drain-source overshoot and ringing is caused by the inherent gate-source and drain-source capacitances of the MOSFETs; resistance and inductance due to other circuit elements; and parasitic capacitance, resistance, and inductance due to layout [19].

A standard method for gate-signal conditioning is the placement of a resistor and diode between the driver and MOSFET in the configuration shown in Fig. 2.11. The gate resistor damps response to an input signal, which causes slower transition times but smaller overshoot and ringing. The additional placement of a gate diode allows for a fast gate-source voltage fall time, reducing the risk of shoot through.

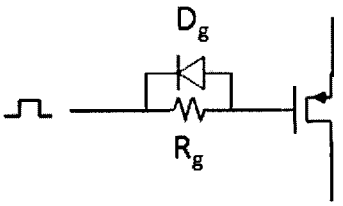

Figure 2.11: Gate resistor and anti-parallel gate diode.

Chapter 3

Circuit Design & Implementation

A test bench is needed that can provide near-sinusoidal voltage and current excitations across an inductive load across a wide frequency and at up to kilowatt power levels. The novel cascaded half-bridge full-bridge (CHBFB) inverter topology, shown in Fig. 3.1, is optimal for meeting these needs. This chapter describes the advantages of this topology, how it operates, and how it was implemented to meet the particular requirements of this work.

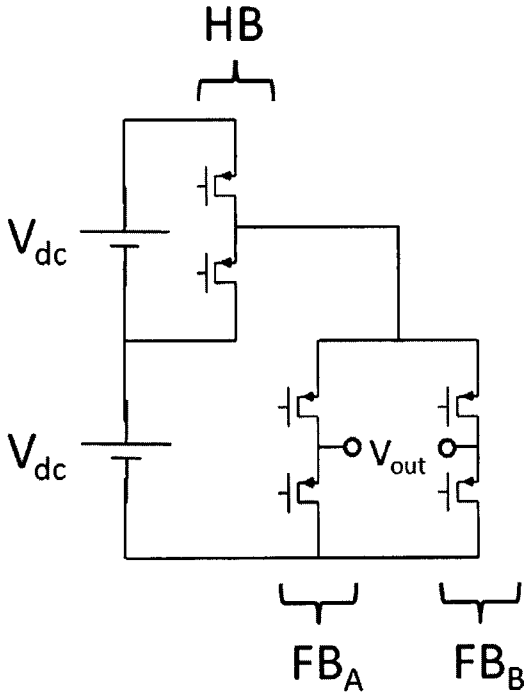

Figure 3.1: Cascaded half-bridge full-bridge inverter.

3.1 Power Circuit Topology Design

The topology of the power circuit must be optimal with respect to its ability to generate near-sinusoidal waveforms across an inductive load across a wide frequency and power range. This section describes the topology selection process.

3.1.1 Resonant Inverters

Since commercial induction cooktops use resonant inverters, it would seem reasonable that this topology could be adapted for induction heating load experimentation. However, this is not an ideal strategy because resonant inverters lack the frequency flexibility needed to conduct load characterization experiments.

Commercial induction cooktops use resonant inverters because maximizing the efficiency is critical for that application, but frequency flexibility is relatively unimportant. The inductance of induction heating loads is incorporated into the resonant network so that the system can achieve efficiencies of over 98% [8]. Peak power and efficiency are achieved by operating near the resonant frequency, and the power level can be lowered by moving the operating frequency away from resonance [7, 8, 23]. Resonant inverters are ideal for meeting the objectives of commercial induction cooktops.

However, because resonant inverters are designed at a fixed frequency with only slight variation, they are incapable of driving an inductive load across broad frequency and power ranges. Although resonant inverters are capable of driving an inductive load across a wide frequency range, doing so would result in a tremendous reduction in the output power capability of the inverter. Llorente et al. have shown that even varying the frequency by a factor of two reduces the power output to less than 40% of peak output for the most common induction cooktop resonant inverter topologies [8]. Therefore, the standard resonant inverter topology is inadequate for this work.

Although it is possible to adapt resonant inverters to operate across a range of frequencies, doing so requires significant additional circuit complexity. Puyal et al. have built a series-resonant half-bridge inverter module for the purpose of inductive load characterization [9]. Although based on a resonant inverter topology, their inverter achieves frequency flexibility by modifying the resonant frequency using a large, adjustable capacitor bank. This allows the inverter to sweep across a range of frequencies while always operating near the resonant frequency. However, the design and assembly of the capacitor bank is not trivial. Therefore, although the frequency-controllable resonant inverter topology demonstrated by Puyal et al. is a feasible method of achieving our goals, from a design perspective it is not ideal.

3.1.2 Multilevel Inverters

Multilevel inverters meet the needs of the test bench:

- **Inductive Loads.** Multilevel inverters are capable of driving inductive loads, in which the impedance angle is between 0°–90°. This is demonstrated by their common use as reactive power compensators and induction motor drives [12].

- **Near-sinusoidal voltage excitations.** The amount of harmonic distortion generated by a multilevel inverter decreases with the number of voltage levels.
- **Frequency Range.** Multilevel inverters are able to operate across a wide range of frequencies because the output voltage waveform is entirely determined by the on/off states of the semiconductor switches in the inverter; the number of output voltage levels, duration of levels, and waveform frequency are all controllable parameters. The frequency range, therefore, is limited only by the frequency range of the switches.
- **Power range.** Multilevel inverters are more than capable of operating at power levels in the kilowatt range. Multilevel inverters have been used extensively since the 1990's for medium-voltage high-power applications in the kilovolt, kilowatt range and higher [12, 24].

However, as discussed in Chapter 2, Background, each of the three standard multilevel topologies has certain disadvantages that require additional circuit complexity to overcome. Therefore, while any of the three standard multilevel topologies are adequate, none are ideal. Diode-clamped and flying capacitor multilevel inverters both require a large number of components per output voltage level and have problems with device voltage balance. While CHB inverters do not have these issues, they require two isolated dc sources. Although none of these issues render the inverters insufficient for the proposed inductive load experiments, they do require additional complexity in the circuit, the switching scheme, or both. A more optimal topology is wanted that can meet the needs of this application without a large number of components, without issues related to voltage balancing, and without the need for isolated supplies.

3.1.3 Cascaded Half-Bridge Full-Bridge

A hybrid class of multilevel inverters called the cascaded multilevel inverter (CMI) provides the advantages of the standard multilevel inverters without the problems associated with the diode-clamped, flying capacitor, and cascaded H-bridge topologies. The CMI consists of two- or three-level inverters, such as half- and full-bridges, cascaded together in series. Specifically, a CMI topology we call a cascaded half-bridge full-bridge (CHBFB) inverter is ideal for the needs of this work.

The CHBFB inverter consists of a half-bridge in series with a full-bridge, as shown in Fig. 3.1. The full-bridge is capable of achieving output voltages of zero and $\pm$ the rail voltage. The half-bridge sets the full-bridge's rail voltage to either V_{dc} or $2V_{dc}$. Therefore the possible output voltage levels are $2V_{dc}$, V_{dc}, 0, $-V_{dc}$, or $-2V_{dc}$. The bridge leg configurations that correspond to the five possible levels are shown in Fig. 3.2.

The CHBFB surpasses the performance of the standard multilevel topologies in several respects. First, it produce a five-level output voltage waveform with only six switches, compared to the eight or more required by the three standard multilevel topologies. In addition, it does not require additional blocking diodes, flying capacitors, or isolated dc sources. Although it does use two sources, they are connected in series and do not require isolation. Thus, it is capable of meeting the needs of this application with minimal circuit complexity.

In addition, the flexibility in the zero-state configuration allows for the minimization of the number of high/low bridge changes in the cycle, which reduces the demands on the controls and

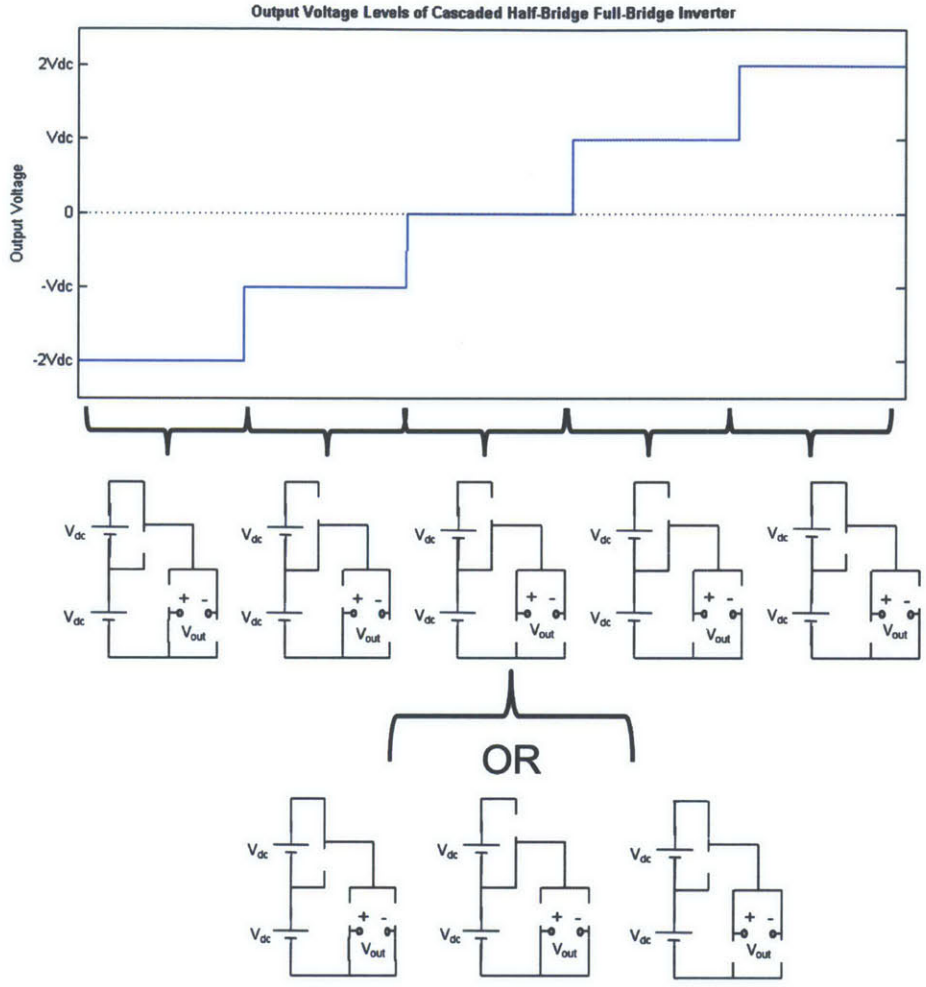

Figure 3.2: Possible output voltage levels achievable with the CHBFB inverter and the corresponding switch configurations.

allows for simpler control schemes. Unlike all of the other output voltage states —which can only be achieved by a single configuration —the zero-state can be achieved in four different ways. This allows for choosing the zero-state configuration in order to minimize the number of high/low bridge changes in the cycle. For example, the staircase-like output voltage waveform of Fig. 3.2 can be produced using any of the four zero-state configurations. However, by implementing the top configuration, only one bridge leg changes in both the transitions between the $-V_{dc}$ to 0 states and between the 0 to $+V_{dc}$ states. The ability of the CHBFB topology to achieve an output voltage of zero with four possible configurations adds flexibility to the switching scheme.

Therefore, the CHBFB is an ideal inverter topology for our application.

3.2 Circuit Implementation

Having selected the CHBFB topology, every element of the circuit was chosen and the circuit board designed in order to meet the application requirements: driving an inductive load across frequency and power level ranges broad enough for characterization of induction heating load nonlinearities, encompassing the standard operating conditions of domestic induction cooktops, which are approximately 20–100 kHz and from 50 W–3.5 kW [7, 8]. Based on these requirements, a frequency range of 1–100 kHz and a maximum power of 2 kW were selected for our test bench.

Because the impedance of the load is not known, the voltage and current range necessary to produce 2 kW power was estimated. The power is provided by two benchtop power supplies placed in series capable of producing 80 V and 50 A_{rms} (described in greater detail in Sec. 3.2.3.2). Therefore, in order to ensure operation at all operating conditions, the circuit was designed to operate at output voltages of up to 80 V (160 V_{pp}) and output currents of up to 50 A_{rms}.

3.2.1 Switching circuit

3.2.1.1 Switches

The switch type selected for this application is the N-type power MOSFET. The CHBFB requires that all of the switches be controlled, so it was necessary to choose active switches. Out of all the active switches, N-type power MOSFETs were selected because they have the lowest on-state resistance for this application's operating conditions.

The MOSFETs are required to operate across the frequency and power ranges without breaking down and to transition between on/off states with reasonable speed and drain-source ringing. This is ensured by selecting switches that meet the voltage and current ratings of this application while providing high switching speed, low on-state resistance, and favorable thermal characteristics. The Toshiba TK100E10N1 N-channel Power MOSFET was chosen because it meets these objectives. The relevant device characteristics are summarized in Table 3.1.

This MOSFET meets the device rating requirements. The MOSFETs must be able to block 80 V drain-source voltage and pass 50 A_{rms} drain current in order to operate at maximum power. The selected MOSFET exceeds this requirements with a maximum drain-source voltage of 100 V and a maximum drain current of 100 A.

Additionally, the performance characteristics are optimal for the needs of this application. Because the MOSFETs conduct large amounts of current, it was critical to select a MOSFET with low on-state resistances and favorable thermal characteristics. The on-state resistance is only 5 mΩ and the thermal resistance is 0.3 °C/W . These values are on the low end of commercially available values for MOSFETs that can operate at the current level required for this application. The maximum junction temperature of 150°C is a typical value for commercial MOSFETs. Therefore, this MOSFET has optimal on-state resistance and thermal characteristics.

The selected MOSFET has adequate switching speed for this application. The rise and fall times of the selected MOSFET are 32 ns and 45 ns, respectively, which correspond to 0.32% and 0.45% of the shortest required period. The gate charge of the selected MOSFET is 140 nC, which is mid-range in the commercially available values for MOSFETs with comparable current ratings.

Although smaller gate charges can be found for comparablte voltage and current ratings, the devices with lower gate charges also have higher on-state resistances. Maximizing the switching speed of the MOSFETs was sacrificed for the sake of minimizing the power dissipation.

Category	Characteristic	Value
Ratings	Drain-Source Voltage Rating	100 V
	Drain Current Rating	100 A
Switching Speed	Rise Time	32 ns
	Fall Time	45 ns
	Gate Charge	140 nC
On-state Resistance	On-state Resistance	5 mΩ
Thermal Characteristics	Maximum Junction Temperature	150°C
	Thermal Resistance	0.3 °C/W

Table 3.1: Toshiba TK100E10N1 N-channel Power MOSFET Characteristics

3.2.1.2 Thermal Considerations

The thermal system was designed to ensure that the MOSFETs do not break down due to overheating. Heat sinks were selected in order to remove the required amount of heat under maximum power dissipation conditions. Because of the inductive nature of the load and the high current level, maximum power dissipation occurs when the inverter is operated at the lowest frequency. Therefore, the maximum power loss in the chosen MOSFETs can be calculated using (3.1):

$$P_{diss,max} \approx \left(I^2_{rms,max} R_{ds,on,max}\right) + \left(\frac{1}{2} I_{rms,max} V_{DS,max} f_{min} (t_{r,max} + t_{f,max})\right) \tag{3.1}$$

$$P_{diss,max} \approx \left((44.5\,\mathrm{A})^2 \times 5\,\mathrm{m\Omega}\right) + \left(\frac{1}{2} \times 44.5\,\mathrm{A} \times 80\,\mathrm{V} \times 1\,\mathrm{kHz} \times 200\,\mathrm{ns}\right) = 10.3\,\mathrm{W} \tag{3.2}$$

where $R_{ds,on,max}$ is the maximum on-state resistance at the highest required operating temperature of 150 °C.

Therefore, allowing for a 20°C safety margin, the maximum allowable thermal resistance of the heat sink can be calculated using (3.3):

$$R_{\mathrm{sa}} = \frac{(T_j - 20°\mathrm{C} - T_{\mathrm{a}})}{P_{\mathrm{diss}}} - R_{\mathrm{jc}} - R_{\mathrm{cs}} \tag{3.3}$$

$$R_{\mathrm{sa}} = \frac{(150°\mathrm{C} - 20°\mathrm{C} - 20°\mathrm{C})}{10.3\,\mathrm{W}} - 0.3°\mathrm{C/W} - 0.33°\mathrm{C/W} = 10.0°\mathrm{C/W} \tag{3.4}$$

where a thermal pad with a case-sink thermal resistance of 0.33 °C/W is used to connect the switch to the heat sink.

The Assmann V7466Z heat sink was selected because, with a thermal resistance of 5.6 °C/W, it provides the necessary cooling with a sizable safety margin. Using this heatsink, the maximum junction temperature is 84°C, which is well below the maximum rated temperature of 150°C.

3.2.2 Driving Circuit

3.2.2.1 Driver

Several key requirements can be outlined for the driver for this application. The driver must be able to:

- source and sink 1.4 A current such that it can turn the selected MOSFETs on and off in 100 ns, which corresponds to 1% of the shortest period
- generate gate signals of $\geq$10 V to bias the MOSFET into the linear region where the on-state resistance is low
- receive and respond to 5 V logic control signals provided by a microcontroller (as described in Chapter 4, Controls Design & Implementation)
- drive high-side N-type MOSFETs
- prevent shoot through[1]

In order to meet these requirements with the simplest driving circuit design, a commercial multi-output driver was selected. Specifically, the Intersil HIP4081A 80 V/2.5 A Peak, High Efficiency Full Bridge FET Driver is used. This driver meets or exceeds each of the above requirements:

- It can source approximately 2.3 A_{peak} and sink approximately 2.6 A_{peak}, providing MOSFET turn-on and turn-off times of less than 100 ns.
- It can generate gate signals with magnitude up to 16 V. It provides gate signals of magnitude equal to its supply voltage, so the user may control the gate signal magnitude by setting the supply voltage (within the range 0-16 V). For this application, 12 V was selected so that the MOSFETs are well into the linear region when conducting current.
- This driver's input logic thresholds are compatible with 5 V to 15 V logic levels.
- It drives the high-side N-type MOSFETs with a bootstrap circuit. An external bootstrap diode and capacitor must be provided. This is discussed in the following section.
- This driver prevents shoot through by incorporating user-programmable dead time: the amount of time after switching off one FET in a bridge before turning on the other. The dead time is determined by the value of an external resistor that is attached to a specific pin on the driver. Because the fall time of the MOSFETs is capped at approximately 65 ns, a dead time of 90 ns dead time was selected in order to avoid shoot through. This was achieved by using 200 kΩ delay resistors.

[1] Although shoot through prevention can be accomplished either in the control system or driving circuit, we have chosen to accomplish it in the driving circuit to make the controls program simpler.

3.2.2.2 Bootstrap Circuit

The bootstrap diode and capacitor were selected in order to ensure that the bootstrap capacitor could provide the required gate voltage when the high-side switch is on and fully charge when the high-side switch is off.

The bootstrap diode is required to have low reverse recovery charge, low leakage current, and a high maximum current. Based on these requirements, a Vishay MSS1P6 Schottky Barrier Rectifier was selected. Because it is a Schottky diode, it has insignificant reverse recovery charge. In addition, this diode has insignificant leakage current and can supply an ample 25 A surge current.

The minimum capacitance can be calculated using (3.5) and the known circuit parameters:

$$C_{BS,min} = \frac{140\,\text{nC} + \frac{30\,\mu\text{A}}{1\,\text{kHz}}}{0.5\,\text{V}} = 340\,\text{nF} \tag{3.5}$$

where $C_{BS,min}$ is the minimum required bootstrap capacitance, the MOSFET's gate charge is 140 nC, the upper supply quiescent current is 30 μA, the minimum switching frequency is 1 kHz, and the allowed change in voltage across the capacitor is set to 0.5 V. Because we want the bootstrap voltage supplying the high-side MOSFET to remain relatively constant, we set $\triangle V = 0.5\,\text{V}$ out of the 12 V total gate-source voltage.

Similarly, the maximum capacitance can be calculated using (3.6) and the known circuit parameters:

$$C_{BS,max} \approx \frac{25\,\text{A}}{0.5\,\text{V} \times 100\,\text{kHz}} = 500\,\mu\text{F} \tag{3.6}$$

where $C_{BS,max}$ is the maximum allowed bootstrap capacitance, 25 A is the maximum current supplied by the bootstrap diode, 0.5 V is the allowed change in voltage across the capacitor during discharge, and 100 kHz is the maximum frequency.

Based on the limits calculated above we set the bootstrap capacitance to $C_{BS} = 10\,\mu\text{F}$.

3.2.2.3 Gate Circuit

The gate circuit was designed in order to achieve acceptable MOSFET gate-source rise and fall times and drain-source overshoot and ringing. Specifically, the overshoot was limited to below the MOSFET drain-source voltage rating of 100 V, and the gate-source rise and fall times were required to turn the selected MOSFETs on and off in 100 ns, which corresponds to 1% of the shortest period. This was achieved using a gate resistor and anti-parallel diode between each driver and MOSFET gate in the configuration shown in Fig. 2.10.

Both the transition speed and the amount of drain-source ringing are dependent on damping caused by the gate resistance: the larger the resistance, the slower the rise and fall times and the smaller the drain-source overshoot and subsequent ringing. Placement of a 10 Ω resistor in front of the gate achieves a balance between transition speed and ringing. The additional placement of an anti-parallel diode cross the resistor allows for a fast gate-source voltage fall time, reducing the risk of shoot through. A Panasonic DB2X20600L Schottky barrier diode was selected because it can pass 1 A_{avg} of current, such that it can handle the current sunk by the driver during MOSFET turn-off.

This gate circuit achieves the objectives of limiting the drain-source voltage to below 100 V and

the rise and fall times to less than 100 ns. The worst overshoot occurs under maximum power conditions, at 1 kHz and 160 V_{pp}. At this operating point, the drain-source voltage exhibits peak overshoot of ≈2.4, corresponding to a peak voltage of 97.6 V, and significant subsequent ringing that lasts approximately 200 ns. Although not ideal, the MOSFET's drain-source voltage rating is only 100 V, the overshoot does not exceed the voltage rating under any operating conditions and is therefore acceptable for our purposes. Across the frequency and power ranges, the rise and fall times are less than 55 ns and 65 ns, respectively. Sample MOSFET gate-source and drain-source transition waveforms and more detailed analysis are included in Appendix A.

3.2.2.4 Additional Considerations

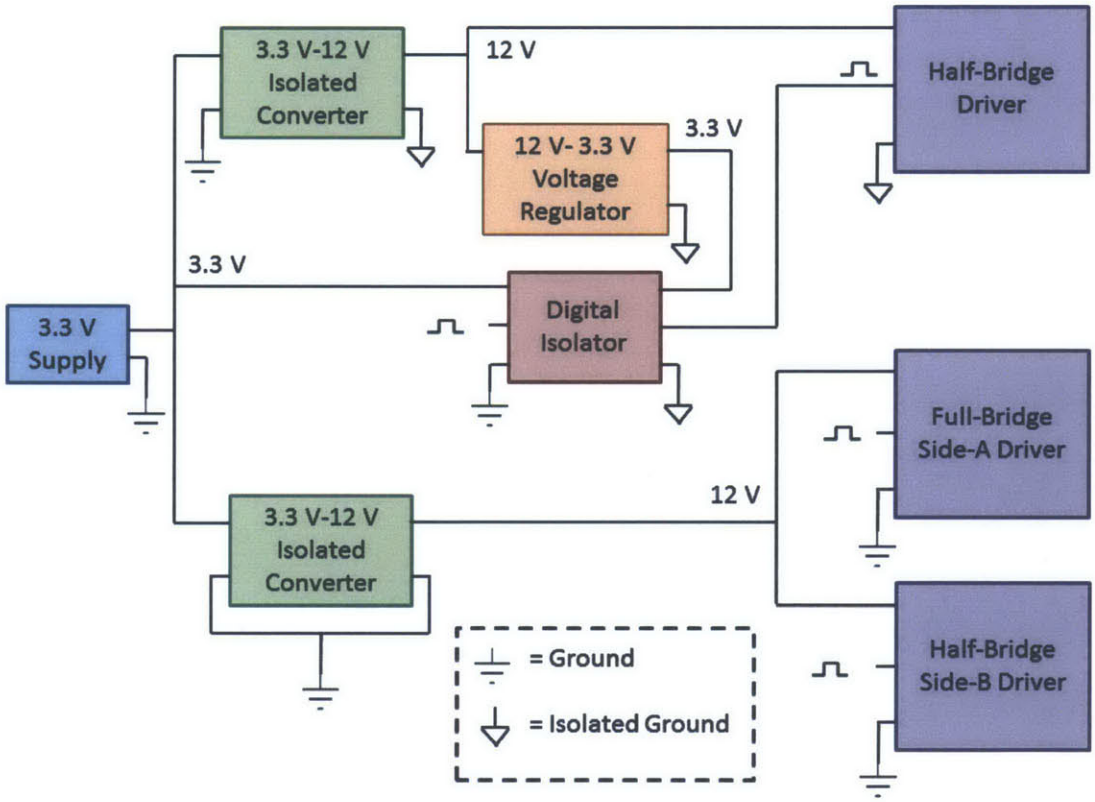

Figure 3.3: Diagram of the isolation and driving circuit power supply.

Two final considerations concerning design of the driving circuit for this test bench are isolation and the driving circuit's power supply, illustrated in Fig. 3.3.

The driving circuit for the half-bridge must be isolated from the driving circuits for the full-bridge. This is necessary because the control signals from the microcontroller are referenced to absolute ground, but the ground reference of the half-bridge driver is connected to the output of the bottom power supply (see Fig. 3.1). An isolator is used to copy the signal from the controller such that it has the same ground reference as the driver. In order to maintain syncronization between

the half- and full-bridge, it is important for the isolator to have a very fast transmission speed. An NVE IL711 digital isolator is selected because it has a typical 10 ns delay, which is fast enough to isolate the signal without causing a noticable delay even at 100 kHz. Isolation was also required for the half-bridge driver's power supply, which was achieved using an isolated dc-dc converter

A single source is used to supply power to the driving circuit (consisting of the drivers and isolation components) as well as the microcontroller providing the control signals. A separate source is used for the driving and control components because the source supplying the power circuit is very noisy and could interfere with the sensitive signals involved in these circuits. A Hewlett Packard 6235A power supply is used because it is capable of generating the maximum power that could possibly be required by the driving and control circuits. It is set to 3.3 V, which can supply the microcontroller directly, and a Recom R1S-R1D isolated dc-dc converter is used to boost the voltage to the 12 V required by the drivers. This converter provides the needed voltage conversion as well as the isolation needed by the half-bridge.

3.2.3 Additional Considerations

3.2.3.1 Layout

In addition to the design of the circuit components, careful layout is critical because a poorly-designed circuit board layout can exacerbate problems associated with noise and unclean switch transitions.

Noise minimization was accomplished using two methods. First, the separate signals —the gate signal propagated from the microcontroller to the driver, the power supplying the digital components and driving circuit, and the power through the inverter supplying the load— are kept in separate locations on the circuit board as much as possible. Placement of two different signals in the same location on adjacent layers was avoided. Additionally, to minimize noise pollution in the digital signals from the higher power signals in the inverter, the ground reference plane of the digital devices and drivers are connected to the ground plane of the bridge at one point.

Several aspects of the circuit board layout design facilitate switch transitions with minimal overshoot and ringing. In order to minimize parasitic inductances, all the current paths in the driving circuits were kept as short as possible, especially the paths from the drivers' gate signal output pins to the MOSFET gates. In order to accomplish this, it was necessary to use a separate driver for each bridge leg of the CHBFB (HB, FB_A, and FB_B), and the components in the driving circuit and switches were placed as closely together as possible. To generate consistent MOSFET gate signalling, the paths from the drivers' gate signal output pin to the MOSFET gate for all MOSFETs have equal length. Decoupling capacitors were placed close to the supply pins of the driving circuit components to provide an alternate route for the high-frequency current during switching.

3.2.3.2 Power Supply

For this application, the two dc voltage sources shown in Fig. 3.1 are supplied by benchtop dc power supplies. Two Hewlett Packard 6269b power supplies are used because they are rated to provide 0–40 V at 0–50 A_{pp}, which meets the demands of this application.

The primary requirement for the power supplies is the ability to maintain a constant voltage under maximum power conditions. As voltage sources approach their current limit, a voltage "droop" can be observed. This problem is especially significant when powering inductive loads under ac conditions, for which the maximum current is drawn at the lowest frequencies. For inductive load applications, the condition for which the power supply provides the most current is the condition for which the supply is required to maintain a constant voltage for the longest time.

The placement of large energy storage capacitors across the power supplies facilitates constant voltage output. Capacitors were added until the measured voltage droop was acceptable: 26 mF were placed across the top supply, 32 mF were placed across the bottom supply, and an additional 5.3 mF were placed across the two supplies in series. The capacitors were attached at the power terminals of the inverter circuit board in order to minimize parasitic inductance that would reduce the effect of the capacitors.

We measured the rate of the voltage droop across the power supplies to 3.6 V/ms under maximum current conditions. We measured the rms current drawn from each power supply in the maximum current condition, 1 kHz at the maximum voltage of 80 V. Then, we drew a low-frequency square wave from each power supply separately at the maximum required current. For this test, the voltage droop rate was measured at 2 V/ms for the top supply, and 1.6 V/ms for the bottom supply.The maximum total droop of 3.6 V corresponds to 4.5% of the voltage magnitude. While this percentage could be reduced by using better power supplies or still more capacitors, a maximum droop that is $<5\%$ of the voltage magnitude is an acceptable amount of voltage droop for this application.

Chapter 4

Control Design & Implementation

A control scheme is needed to implement a near-sinusoidal waveform —that is, a voltage waveform with low harmonic content— across a range of 1--100 kHz with user-controlled frequency. The optimal switching method for this work is harmonic cancellation; this method does not require lossy, high-frequency switching and generally produces the lowest total harmonic distortion for inverters with five or fewer levels, such as the CHBFB [14]. Harmonic cancellation is achieved by turning the switches on and off in such a way as to produce the waveform shown in Fig. 4.1. This chapter describes how this control scheme is implemented in the CHBFB inverter.

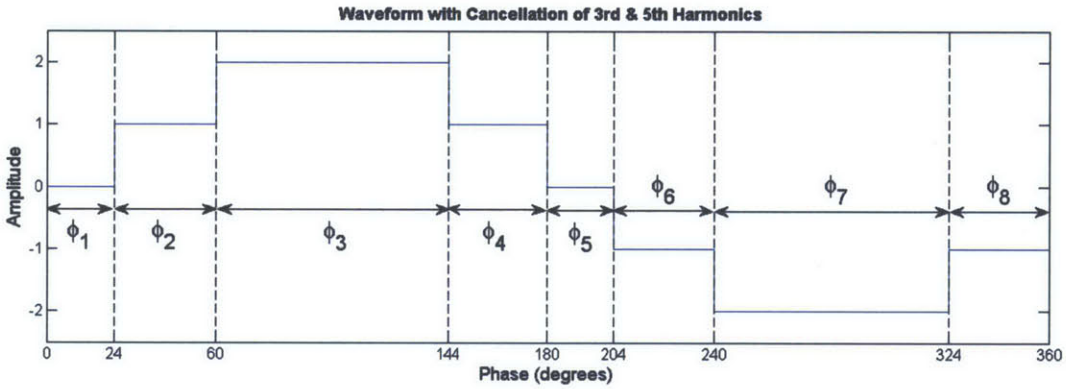

Figure 4.1: Desired near-sinusoidal output voltage waveform.

4.1 Controls Design

In a cascaded half-bridge full-bridge (CHBFB) inverter, the output voltage is determined by the on/off states of the semiconductor switches, or, equivalently, the high/low configurations of the bridge legs. The switching cycle of the CHBFB that achieves the desired voltage waveform is depicted in Fig. 4.2.

This cycle achieves the desired voltage waveform with the fewest possible bridge leg configuration changes: for each output voltage state transition, only one bridge leg changes configuration. The

small number of configuration changes simplifies the demands on the control system.

In this test bench, each bridge leg —the half bridge (HB) and both legs of the full bridge (FB_A and FB_B)— is driven by a separate driver, and a separate control signal is needed for each driver. The control signal indicates to the driver whether it should set the bridge leg into the high or low configuration. In order to generate the load voltage waveform shown in Fig. 4.1, the bridge legs must have the configuration cycle shown in Fig. 4.3. This cycle coincides with the high/low configurations of each bridge in Fig. 4.2.

To produce a given bridge configuration, the control signal to the driver must be the inverse; a low control signal yields a high bridge leg configuration, and vice-versa. This is due to the selected gate drivers: they respond to a low control signal by setting the bridge leg to the high configuration, and vice-versa. Therefore, the control signals required to produce the required bridge leg cycle are the inverse of that cycle.

For this switching cycle, the the half-bridge switches at twice the frequency of the legs of the full-bridge. As shown in Fig. 4.3, the full-bridge goes high once per cycle, while the half-bridge goes high twice. Although it is possible to write the control program to generate control signals with different frequencies, it is not a standard feature of many control devices and would require a more complicated control program. Instead, a simple hardware solution is implemented in order to obtain a control signal for the half-bridge that is twice the frequency of the control signal for the full-bridge. Two signals (HB_A and HB_B) are generated, one that sets the half-bridge high from 60°–144°, and another that sets the half-bridge high from 240°–324°. The two signals are connected using an OR digital logic gate, such that the resulting half-bridge control signal (HB) is the maximum of the two: it sets the half-bridge high from both 60°–144° and from 240°–324°.

Because of the two properties described above —the control signals are the inverse of the bridge configurations, and the half-bridge control signal is composed of two signals connected using an OR gate— the required control signals are shown in Fig. 4.4. These signals produce the desired load voltage waveform shown in Fig. 4.1. A diagram of the control circuit required to implement these signals and connect them with the CHBFB is shown in Fig. 4.5.

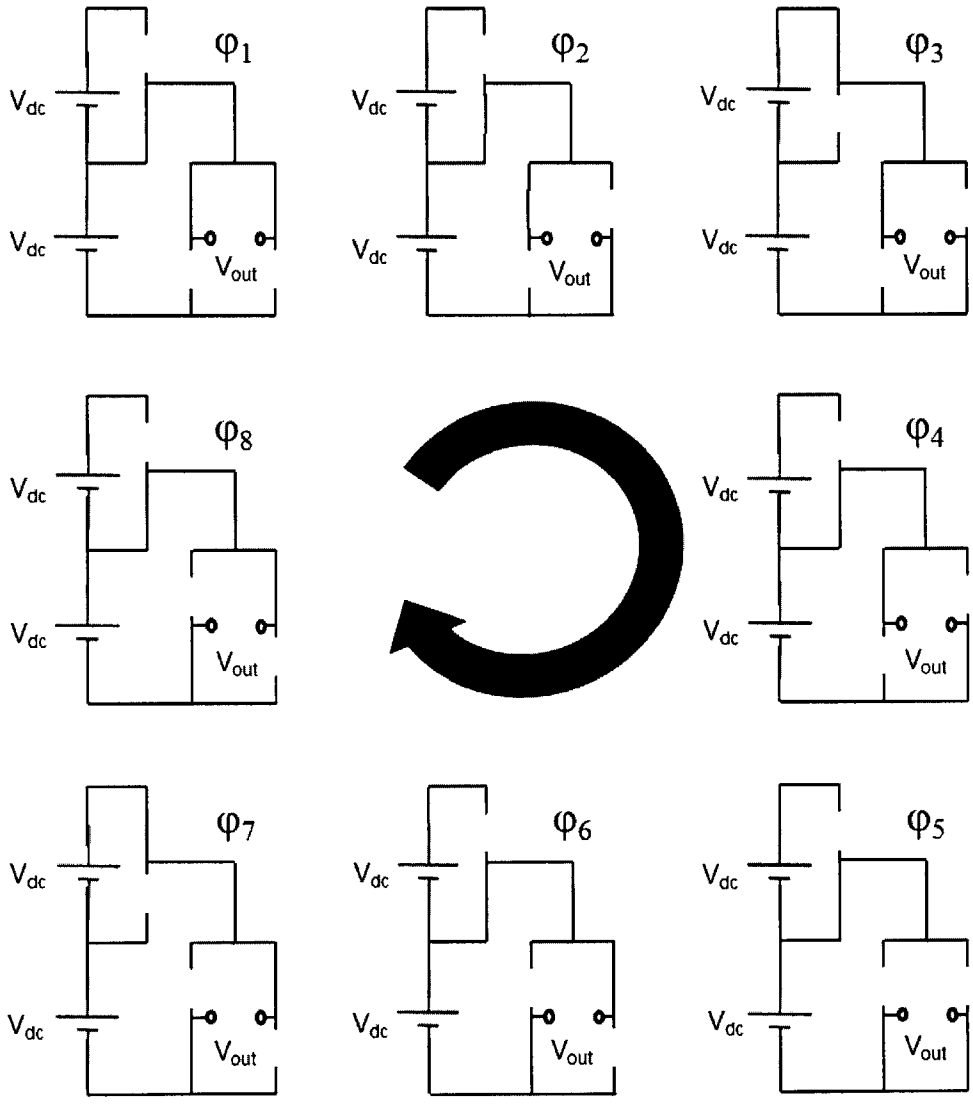

Figure 4.2: Diagram of switching scheme that generates the desired output voltage waveform.

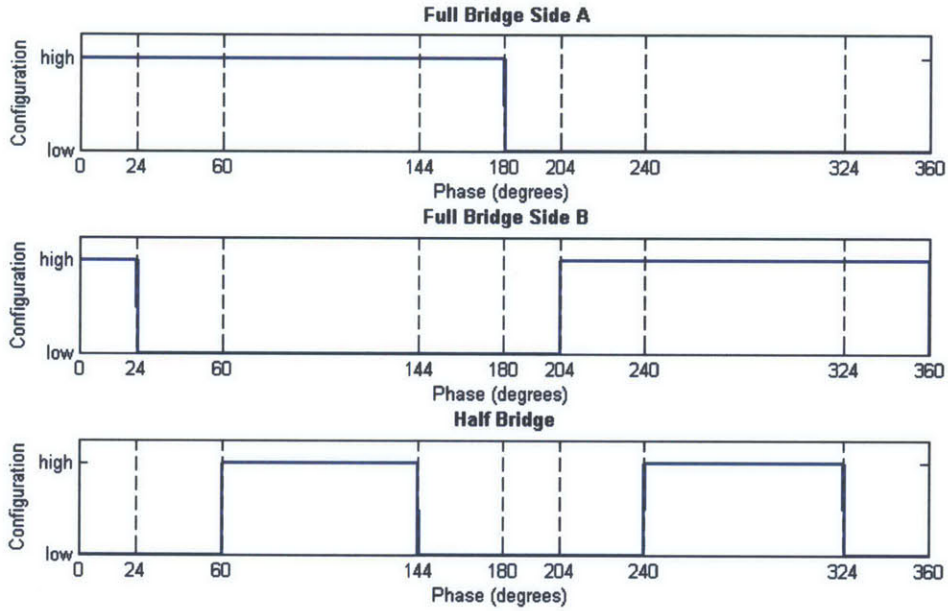

Figure 4.3: High/low bridge configurations that produce the desired output voltage waveform.

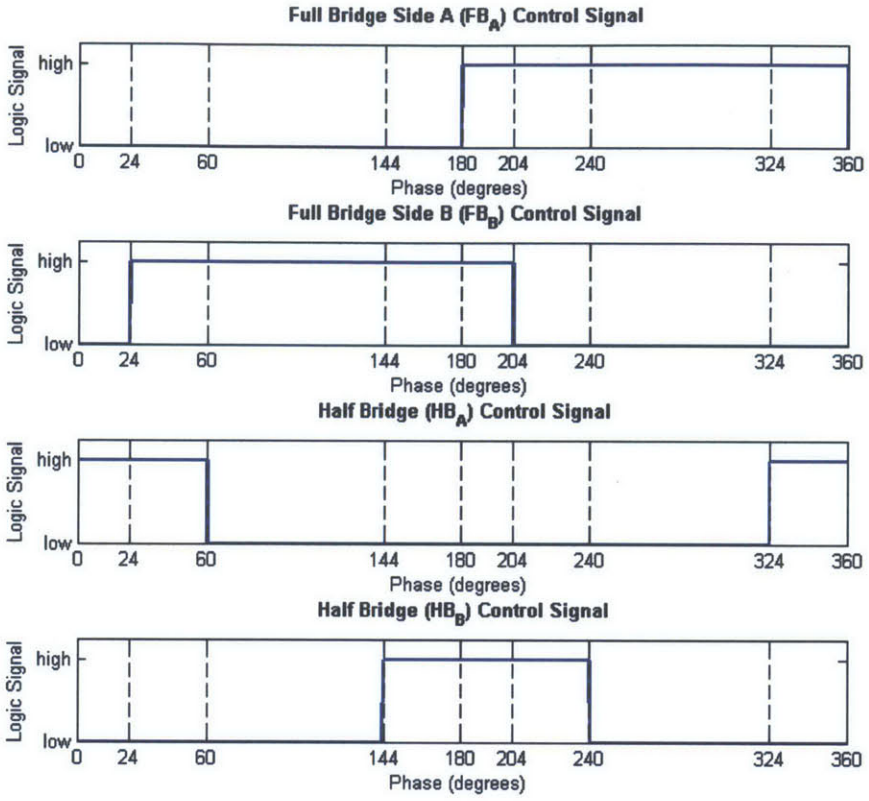

Figure 4.4: Required control signals.

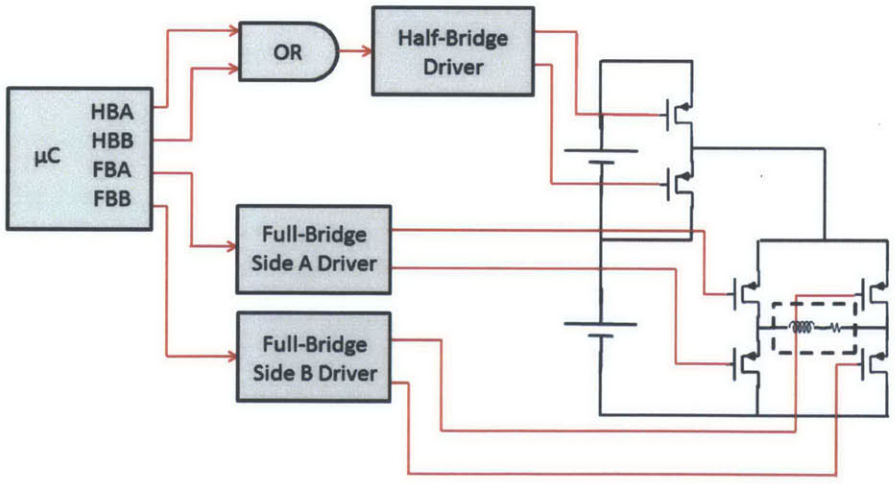

Figure 4.5: Diagram of controls signals (red) for CHBFB inverter. All control signals originate in the microcontroller (μC).

4.2 Control Implementation

4.2.1 Overview

The control program must meet two objectives: production of the required control signals across the frequency spectrum of this application (1–100 kHz), and user-controlled frequency. Meeting the first objective entails generating the control signals depicted in Fig. 4.4. Meeting the second objective entails implementing a user interface (such as a dial or number pad) which the control program reads in order to set the frequency of the control signals. The software and hardware must be designed to meet these objectives.

The software generates the steady-state control loop described by the block diagram of Fig. 4.6. The specifics of how it achieves this control loop (which is discussed in detail in the following section) depends on the hardware.

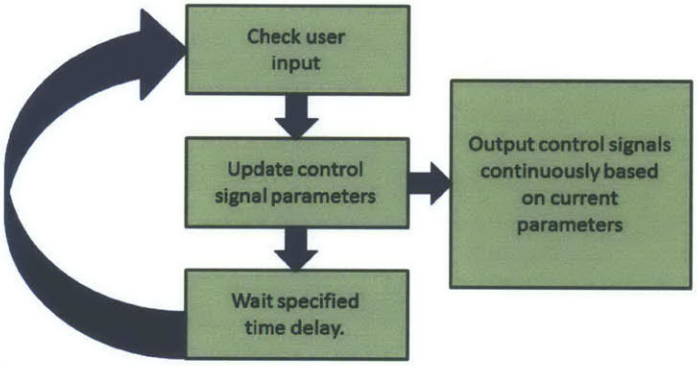

Figure 4.6: Block diagram of control program loop.

The device selected to provide the control signals must be able to achieve the control loop in Fig. 4.6 across a frequency range of 1–100 kHz. It must provide a mechanism for a user interface, generate four synchronized control signals with different duty cycles and relative phase, generate the control signals at up to 100 kHz, and continuously output the control signals while interacting with the user interface.

Based on these requirements, the Microchip dsPIC33FJ16GS502 microcontroller was selected to provide the control signals. This microcontroller has an analog-digital converter (ADC) feature that provides a user interface; the user can specify the control frequency by setting the ADC pin voltage to a given value. Additionally, its PWM module allows for the continuous generation of four synced control signals with independent duty cycle and phase. Finally, the microcontroller's clock system allows for a maximum internal oscillator frequency of 118 MHz, which is more than adequate to generate sharp control signals at the maximum operating frequency of 100 kHz.

4.2.2 Program Design

The microcontroller program implements the control loop in Fig. 4.6 across a frequency range of 1–100 kHz. This is accomplished by configuring the microcontroller's static features, features that may

not be changed during program operation, as well as its dynamic features, which may be changed during program operation.

The static features enable the generation of four synchronized control signals with independent duty cycles and phases. Control signal generation is achieved using the PWM module because it is capable of generating continuous control signals; the periodic checking of the user input does not interrupt the control signals. The generation of four signals is provided for by setting the Output Mode of the PWM module to Redundant, such that four (out of possible eight) PWM channels are activated. The channels are synchronized by selecting the Master Time Base Mode, such that the frequency of all four PWM channels is set by one control register. The Master Time Base Mode also specifies that the phase of each PWM channel is determined by independent control registers. Independent duty cycles are achieved by selecting the Independent Duty Cycle Registers Mode.

In addition, the static features must enable control signal generation at up to 100 kHz. Since the microcontroller speed is determined by its internal clock frequency, this objective is accomplished by selecting an internal clock with a high frequency. Specifically, the internal clock frequency must be high enough such that "jittering", in which microcontroller timing errors cause the control signals to change slightly each cycle, is not significant. In order to obtain control signals at 100 kHz that had no significant jitter, it is necessary to use an internal clock frequency of 117.92 MHz. This is achieved by using the internal RC oscillator with a frequency of 7.37 MHz and activating the phase-locked loop (PLL), which multiplies the clock frequency by 16. However, for an internal clock frequency of 117.92 MHz, certain PWM function registers are too small to store the values required to generate signals at below 20 kHz. Therefore, in the 1–10 kHz range, the PLL is not activated, and the internal clock frequency is 7.37 MHz. Because the control signal frequency is lower, the jittering is not significant despite the lower internal oscillator frequency. However, because the PLL state cannot be changed during program operation, two separate programs are required to cover the entire range of 1--100 kHz: one for 1–20 kHz, and another for 20–100 kHz. The system must be shut down and the microcontroller manually reprogrammed in order to transfer between the two frequency ranges.

The dynamic features must be set during program operation in order to update the required control signals according to the user input. First, the program must check the user input. This is achieved by the periodic reading of the ADC pin, which is connected to a potentiometer configured as an adjustable voltage divider. The potentiometer is connected to a constant voltage source, and the sliding contact is connected to the ADC pin. As the user turns the knob of the potentiometer, the sliding contact moves such that the proportion of resistances changes, causing a changing voltage at the ADC pin. The microcontroller is programmed to read the value of the ADC pin every 6.5 μs (an arbitrary value). The value on the ADC pin is stored in the microcontroller in the function register ADCBUF0 as a value between 0–1023.

After checking the user input, the microcontroller must update the control signal parameters based on the voltage of the ADC pin. This is achieved through a lookup table; a series of "if...else if..." statements implements a portion of code determined by the ADC value. Each portion of code configures the microcontroller to generate the required control signals at a different frequency in the range of 1–100 kHz, proportional to the voltage on the ADC pin.

The frequency (f) of all four PWM channels and the duty cycle (D_x) and phase (θ_x) of each

PWM channel must be defined for a different frequency in each portion of the lookup table. These parameters are determined by the values of certain key function registers:

$$f = \frac{8 \times \text{ACLK}}{(\text{PCLKDIV} \times \text{PTPER})} \tag{4.1}$$

$$D_x = \frac{\text{PDCx}}{\text{PTPER}} \tag{4.2}$$

$$\theta_x = 360° \times \frac{\text{PHASEx}}{\text{PTPER}} \tag{4.3}$$

Thus, producing the desired control signals is achieved by programming the control registers ACLK, PCLKDIV, PTPER, PDCx, and PHASEx (where x indicates the PWM channel, e.g., PDC1, PDC2, etc.). In order to achieve the control signals in Fig. 4.4, the duty cycle and phase shifts of the four PWM signals must be defined as in Table 4.1.[1] It should be noted that the function registers can only store integers, and not decimals, which introduces rounding errors to the duty cycles and phases. However, since the values stored in the PWM function registers are in the hundreds and thousands, this rounding error is insignificant.

PWM Channel	Control Signal	Duty Cycle (%)	Phase (degrees)
1	FBA	$\frac{180}{360} = 0.500$	0
3	FBB	$\frac{180}{360} = 0.500$	156
2	HBA	$\frac{84}{360} = 0.233$	216
4	HBB	$\frac{84}{360} = 0.233$	36

Table 4.1: PWM Control Signal Duty Cycle and Phase Settings

Because the function registers that configure the internal clock –ACLK (the internal oscillator frequency, 7.37 MHz or 117.92 MHz) and PCLKDIV (the internal oscillator frequency divider, set to the default value of 1) –cannot be changed during operation, the control signals must be updated by changing the PTPER, PDCx, and PHASEx function registers. Therefore, each portion of the lookup table defines the PTPER, PDCx, and PHASEx function registers such that the PWM module generates the control signals at a different frequency. For example, if the potentiometer knob is turned all the way clockwise, then the ADC pin reads its maximum value, corresponding to the highest frequency, 100 kHz (considering the program for the 20–100 kHz range). For this case, the PTPER, PDC4, and PHASE4 control registers (corresponding to the HBB control signal) are defined based on (4.4)–(4.4):

$$\text{PTPER} = \frac{8 \times \text{ACLK}}{(\text{PCLKDIV} \times f_{PWM})} = \frac{8 \times 117.92 \times 10^6\,\text{Hz}}{(1 \times 1 \times 10^5\,\text{Hz})} = 9433.6 \approx 9434 \tag{4.4}$$

$$\text{PDC4} = D \times \text{PTPER} = 0.233 \times 9433.6 = 2198 \tag{4.5}$$

[1]Note that the PWM channel/control signal pairings were determined by circuit board layout considerations (i.e. PWM channel 1 corresponds to FBA because it was most convenient based on how the microcontroller was layed out with respect to the driver for the full bridge side A driver).

$$\text{PHASE4} = \frac{36^\circ}{360^\circ} \times 9433.6 = 943.4 \approx 943 \tag{4.6}$$

In this manner, PDCx and PHASEx are defined for all four PWM channels for each frequency. The control register values are placed in the lookup table so that different control signals are generated based on the rotation of the potentiometer knob.

In addition to calculating the function register values as shown in the previous paragraph, manual fine tuning was required in order to achieve the desired load voltage waveform. Although the control signals follow the timing scheme determined by the control registers, the timing syncronization of the signals is not perfectly preserved during propagation to the load terminals. As a result, achieving the calculated control signals in Fig. 4.4 does not perfectly produce the ideal load waveform in Fig. 4.1. The discrepancy in the durations of the different output voltage states (φ_n) was up to 19% of the calculated durations. This was a significant problem because inaccuracy in the durations of the output voltage states impairs the required cancellation of low-frequency voltage harmonics. In order to solve this problem, the PDCx and PHASEx values of each control signal were fine-tuned at each frequency until the measured duration of each state was reduced to within 5% of the calculated duration.

By configuring the function registers as described, the microcontroller may be programmed to achieve the tasks illustrated in Fig. 4.6. By following this control loop, the microcontroller generates the control signals required to switch the CHBFB inverter to create the desired waveform. The final program used to generate the control signals from the dsPIC is included in Appendix B.

Chapter 5

Circuit Performance

We have constructed a test bench to characterize induction heating loads in order to study nonlinearities in the load with respect to the strength of the current and frequency. To this end, the test bench has been designed to meet certain objectives:

- **Low harmonic distortion.** We need to be able to take data in which we excite the inductive load with current at a single frequency. Due to filtering by the inductive load, this can be achieved by generating near-sinusoidal voltage waveforms with low harmonic distortion. To this end, the test bench has been designed to cancel all voltage harmonics that are multiples of 2, 3, and 5 (all 2n, 3n, and 5n harmonics). For the purpose of load characterization experiments, we require that the magnitude of these harmonics be below 3% of the magnitude of the fundamental.

- **Range of currents and frequencies.** We want to measure how the load impedance changes with respect to the current and the frequency. This test bench has a frequency range of 1–100 kHz. At each frequency, a range of currents can be achieved by varying the voltage from 1–160 V_{pp}.

- **High power levels.** We must be able to drive inductive loads at the current levels at which they approach saturation, such that their nonlinear behavior can be easily observed. This requires operating at high power levels. To meet this objective, the test bench can drive the load with up to 2 kW of power.

- **Inductive Loads.** The test bench must be capable of driving inductive loads, for which the impedance angle is between 0° and 90°. The test bench meets all of the above objectives while driving the kind of load used with induction cooktops: induction coils coupled to cookware (i.e. pots and pans) made out of a magnetic material, such as cast iron.

The following chapter presents an analysis of the test bench with respect to these objectives. A comprehensive assesment of the test bench can be made by operating the test bench in two regimes: at low current levels, in which the load is approximately linear, and at high current levels, in which the load approaches saturation. Accordingly, two experiments are presented. In the first, we use the test bench to drive an inductive load at a low voltage across the frequency range of the test bench,

corresponding to a low to medium current range. This test allows for measurement of the harmonic distortion produced by the test bench while driving inductive loads and an analysis of how well the voltage waveform produced by the test bench matches the ideal waveform. In the second test, we use the test bench to drive an inductive load at the maximum voltage across the frequency range of the test bench, corresponding to a high power and field strength range. This test demonstrates the ability of the test bench to drive an inductive load such that it approaches saturation and its nonlinearity can be easily observed.

The resulting circuit board and test bench setup are shown in Fig. 5.1 and Fig. 5.2. The full schematic and a list of the components are included in Appendix C.

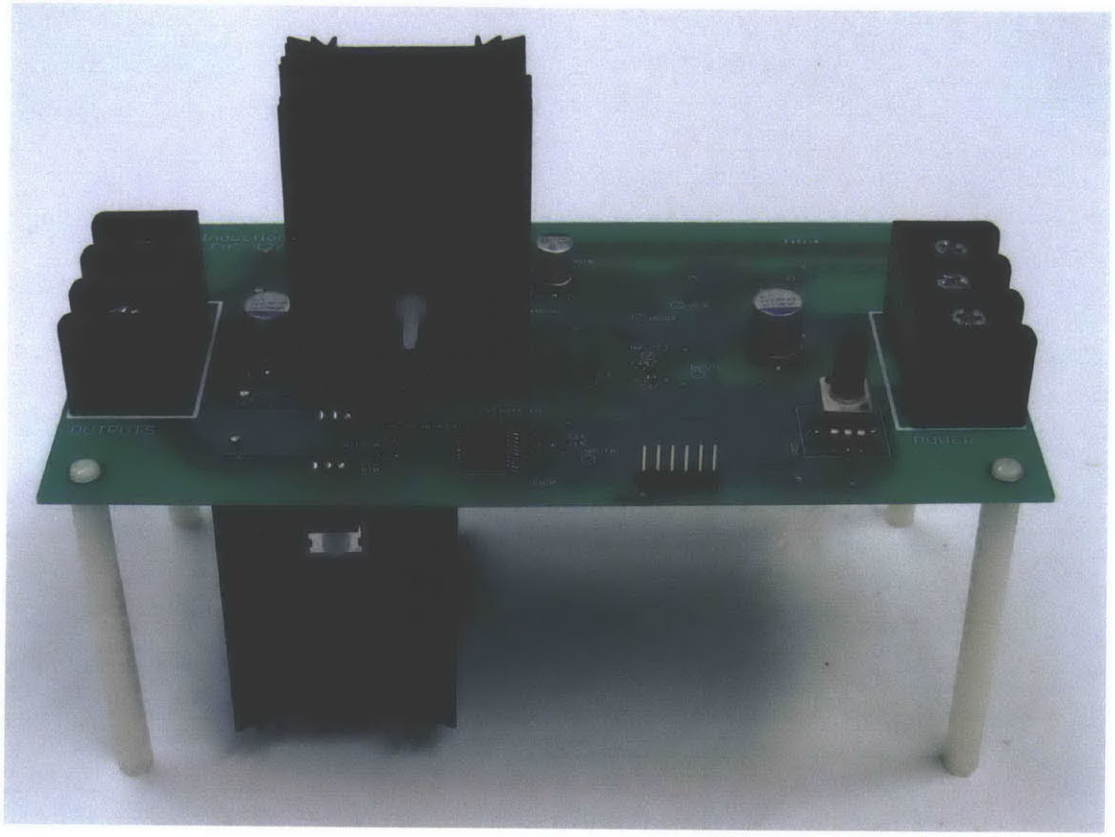

Figure 5.1: Completed circuit board.

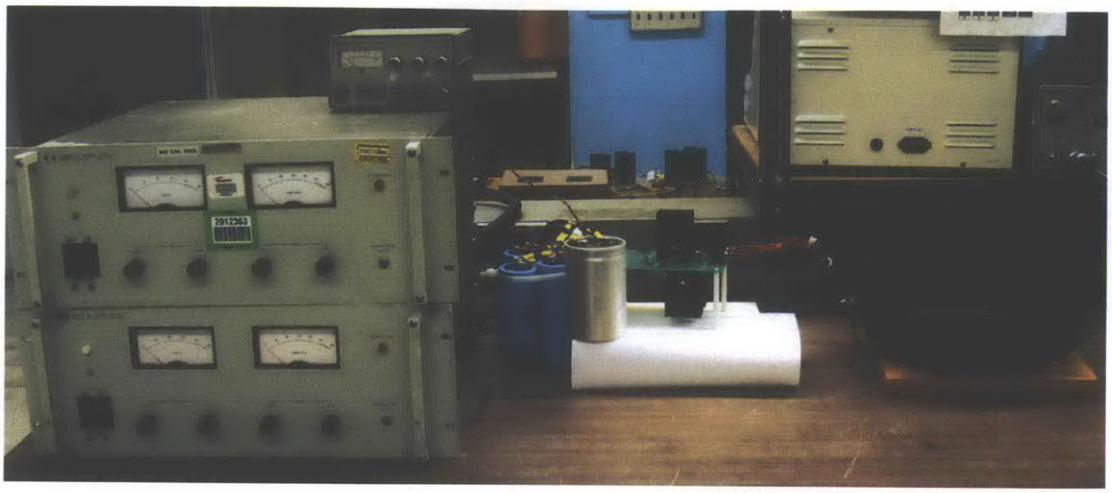

Figure 5.2: Setup of test bench, including the power supplies, external capacitors across power supplies, circuit board, and pan-loaded coil.

5.1 Low to Medium Current Test

The purpose of the first test is to assess the test bench's ability to produce near-sinusoidal voltage and current excitations across an inductive load. The extent to which the voltage waveforms are sinusoidal depends on how well the circuit reduces all higher harmonics. To this end, the circuit topology and switching scheme were designed to cancel all 2n, 3n, and 5n harmonics, such that the lowest present harmonics are the 7th and 11th. Thus, the quality of the voltage waveforms produced by the test bench depends on how well the test bench achieves the ideal waveform described in Chapter 2, Background.

The quality of the current waveforms, on the other hand, is dependent on the load. By driving the load at low currents, where the magnetic material does not approach saturation, we expect the current harmonic content to be free from nonlinear effects. Therefore, because the load is inductive, we expect the harmonic distortion present in the current waveforms to be lower than that of the voltage waveforms.

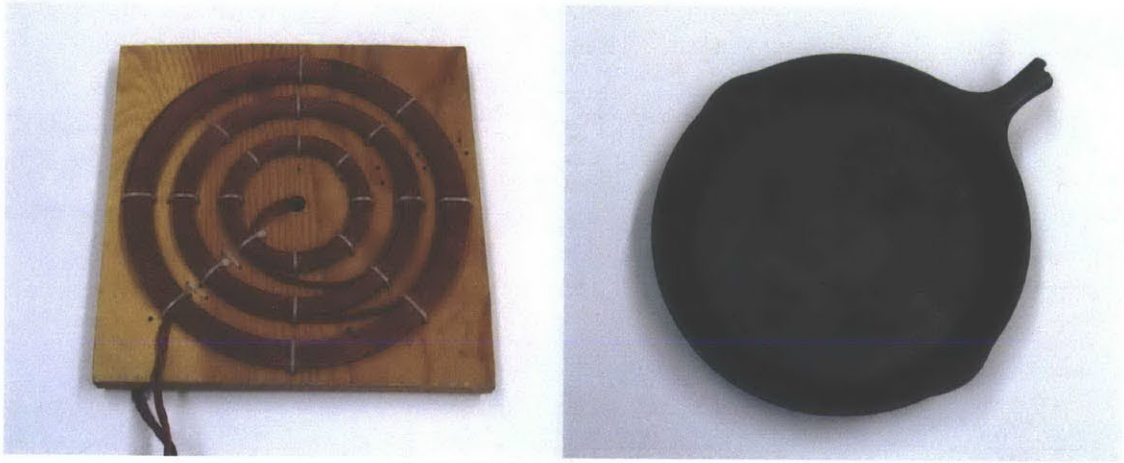

(a) Induction coil (b) Pan

Figure 5.3: Picture of the custom induction coil and 11.5 in diameter cast iron pan

5.1.1 Experiment Design

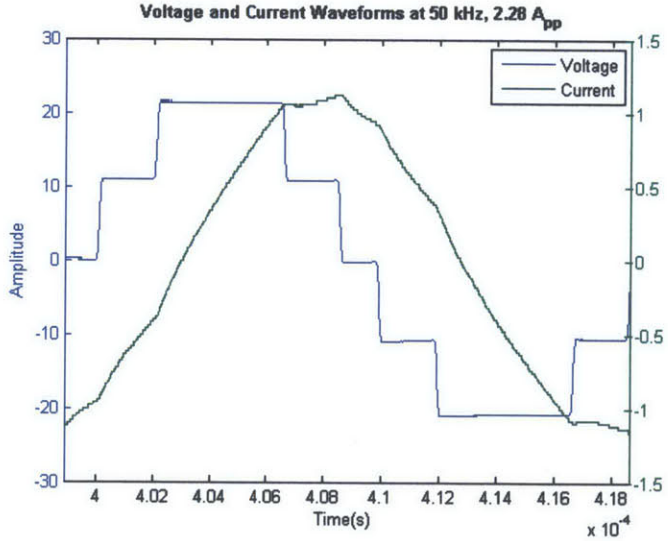

Figure 5.4: Voltage and current of the pan-loaded coil at 40 Vpp and 50 kHz, 2.28 A_{pp}. Note that the discrete appearance of the current waveform is due to the resolution of the oscilloscope.

The experimental process consisted of measuring the voltage and current waveforms across the inductive load for 1-100 kHz in 10 kHz increments at 40 V_{pp}. The load for this experiment was the custom induction coil coupled to the cast iron frying pan shown in Fig. 5.3. A Tektronix TPS2024B oscilloscope was used to collect at least one period of voltage and current data, averaged over 128 cycles to reduce experimental noise. Due to the inductive nature of the load, 40 V_{pp} generates different current levels at different frequencies. Across the test bench's frequency range of 1-100 kHz,

at 40 V_{pp} the corresponding current and power ranges were 42-1.4 A_{pp} and 14-0.5 W. An example of the voltage and current waveforms produced by the test bench across the pan-loaded coil are shown in Fig. 5.4.

The collected data was processed in order to extract information about the harmonic content of the voltage and current signals: the time-domain signal was transformed into the frequency domain using a Fourier transform. This was achieved by selecting exactly one period of data and calculating the discrete fast Fourier transform (FFT) using Matlab. The code used to do this is in Appendix D. This calculation decomposes each waveform into its harmonics, allowing us to easily compare the relative magnitudes of the undesirable higher harmonics with the magnitude of the fundamental. To this end, all the amplitudes were normalized such that the amplitude of the fundamental was 1; an amplitude of 1×10^{-2}, for example, corresponds to 1% of the fundamental.

5.1.2 Main Results

The sinusoidal quality of the voltage and current waveforms is illustrated by the FFTs of the waveforms. A perfectly sinusoidal signal has a fundamental amplitude of 1, and all the higher harmonics have amplitudes of 0. Although this test bench is not designed to achieve a perfectly sinusoidal signal, the goal of this test bench is to cancel the 2n, 3n, and 5n harmonics to below 3% of the fundamental, such that the waveforms are approximately sinusoidal.

The sample of FFTs shown in Fig. 5.5 shows that across the tested range of field strengths, the degree of harmonic cancellation of the 2n, 3n, and 5n harmonics consistently meets the 3% benchmark (marked by a dashed line). This suggests that the test bench produces excitations that are sinusoidal enough for the desired load characterization experiments. The large value of the 7th and 11th harmonics is expected; the switching scheme was not designed to cancel the 7th and 11th harmonics, so their magnitudes are approximately what they would be for a square wave: $\frac{1}{n}$ of the fundamental. Also as expected, the harmonics in the current waveform are generally lower than the harmonics in the voltage waveform due to load filtering. Differences in the harmonic content across the current range is the result of differences in the voltage waveforms, as will be discussed in greater detail in Sec. 5.1.3.

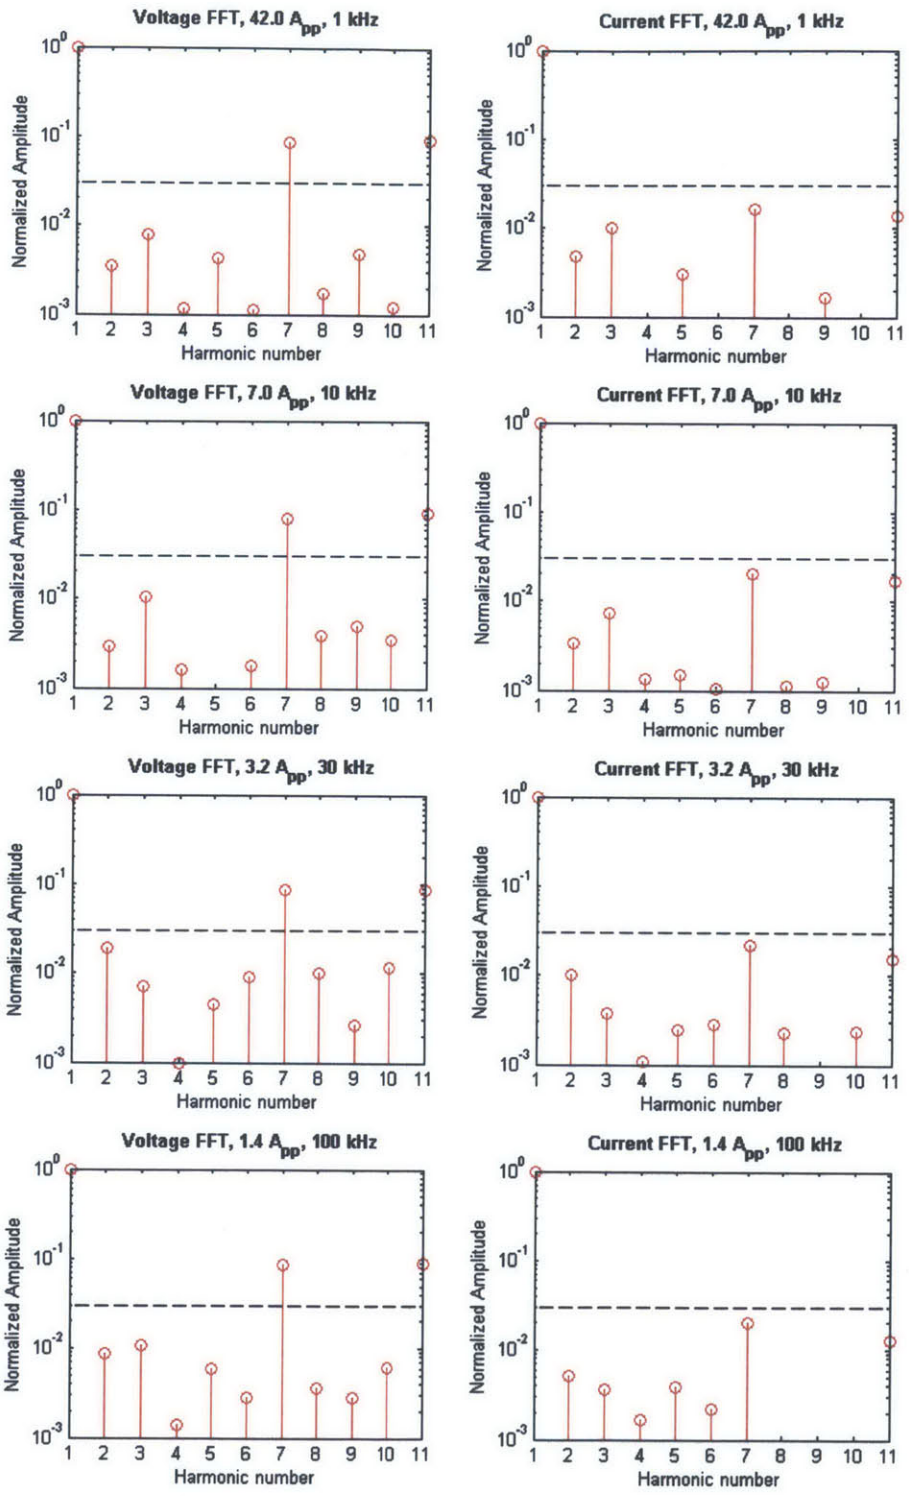

Figure 5.5: Samples of FFTs taken at 40 V_{pp} across 42-1.4 A_{pp}, corresponding to 1-100 kHz and 14-0.5 W. If no value is shown for a given harmonic, then the magnitude of that harmonic is less than 10^{-3} of the fundamental. The dashed line marks 0.03 of the fundamental.

5.1.3 Waveform Analysis

Unexpected harmonics are present in the FFTs in Fig. 5.5 because the voltage waveforms produced by the test bench do not perfectly match the theoretical waveform (described in Chapter 2, Background), in which the 2n, 3n, and 5n harmonics are cancelled.

By comparing the achieved and ideal waveforms, we can identify the non-ideal characteristics present in the achieved waveforms that are the source of higher harmonics. Fig. 5.6 shows the ideal load voltage waveform as well as the achieved waveforms at the limits of the frequency range of the test bench, 1 kHz and 100 kHz.

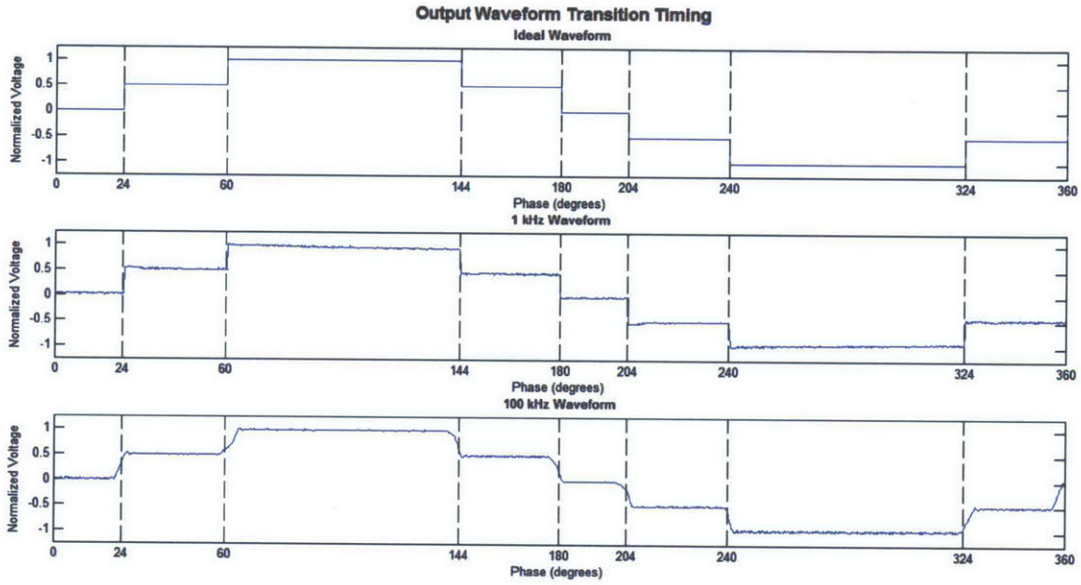

Figure 5.6: Ideal and measured load voltage waveforms. Measuremements were taken at 160 V_{pp}.

To highlight the flaws in the switching cycle, we subtracted the ideal waveform from the achieved waveform. Therefore, points in the difference waveform with nonzero amplitude correspond to points where the achieved waveform does not match the ideal waveform. The ideal and achieved voltage waveforms, the difference between them, and the FFT of the difference are shown at 1 kHz in Fig. 5.7 and across the frequency range in Fig. 5.8. In these figures, the FFTs are normalized to the same scale used in Fig. 5.5: the magnitude of the fundamental of the voltage waveform is 1.

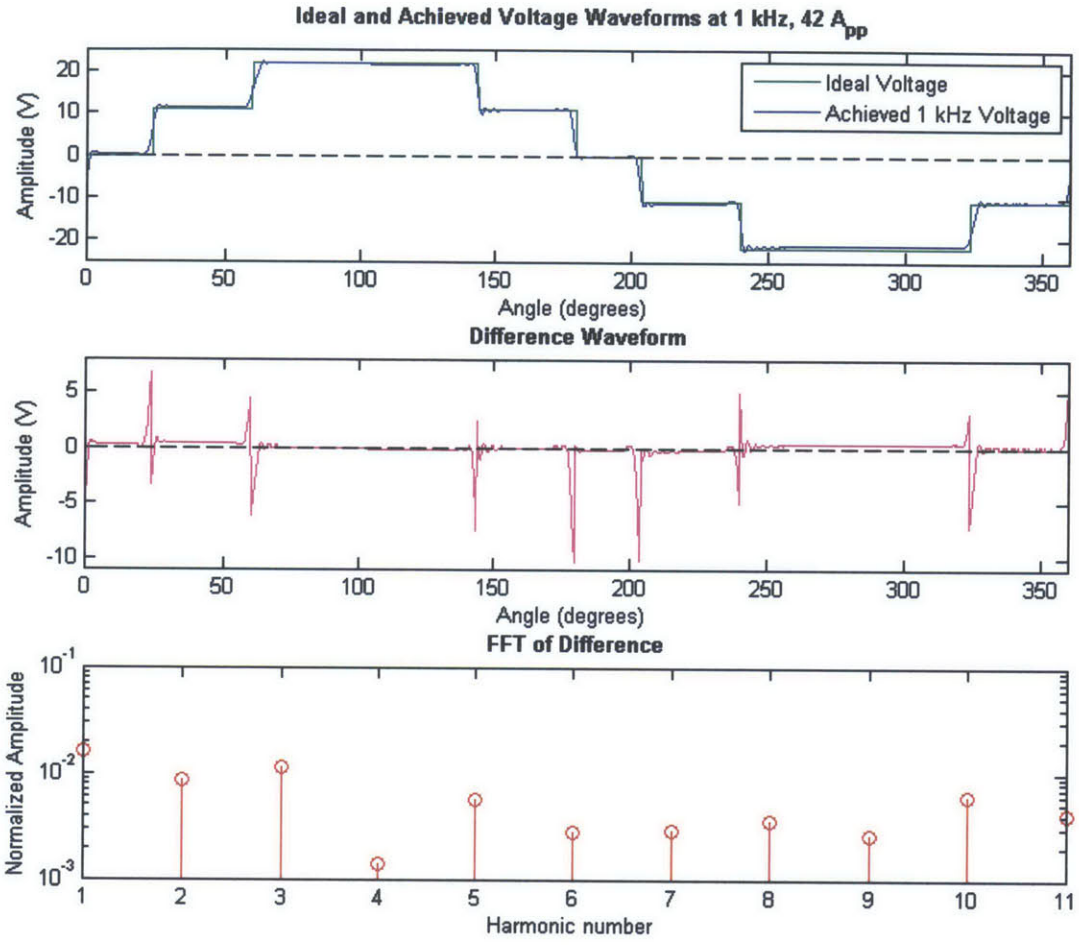

Figure 5.7: Top: Ideal and achieved load voltage waveforms at 1 Hz, 40 V_{pp}; Middle: Difference between ideal and achieved waveforms; Bottom: FFT of difference waveform normalized such that the magnitude of the fundamental of the voltage waveform is 1.

The "Difference Waveform" plots in Fig. 5.7 and Fig. 5.8 show three non-ideal features:

- **Fundamental Frequency Oscillations.** This can most easily be observed in the 42.0 A_{pp} difference waveform in Fig. 5.7. The cause of this behavior is the power supply "droop" problem described in Chapter 3, Circuit Design & Implementation.
- **Voltage spikes.** Examples of voltage spikes are at 360° in the 7.0 A_{pp} waveform or at 60° in the 1.4 A_{pp} waveform in Fig. 5.8. This problem is caused by imperfect timing: errors due to both inaccurate switching timing and the finite transition time of the MOSFET switches.
- **High-frequency Oscillations.** This can most easily be observed immediately after the 240° voltage spike in the 1.4 A_{pp} difference waveform in Fig. 5.8. This corresponds to ringing across the load voltage.

The non-ideal characteristics observed above suggest that the source of unexpected harmonics in the waveforms are caused by voltage droop, imperfect switching, and ringing. Furthermore, Fig. 5.8 highlights the fact that the achieved voltage waveforms differ across the tested currents, which explains why the harmonic content also varies.

The FFTs of the difference waveforms are evidence that the difference between the ideal and achieved waveform is the source of the undesired harmonics. Fig. 5.9 compares the FFTs of the original voltage waveforms with the FFTs of the difference waveforms. As expected, the FFTs are comparable for for 2n, 3n, and 5n harmonics; but the fundamental and the 7th and 11th harmonicsare not present in the difference FFT.

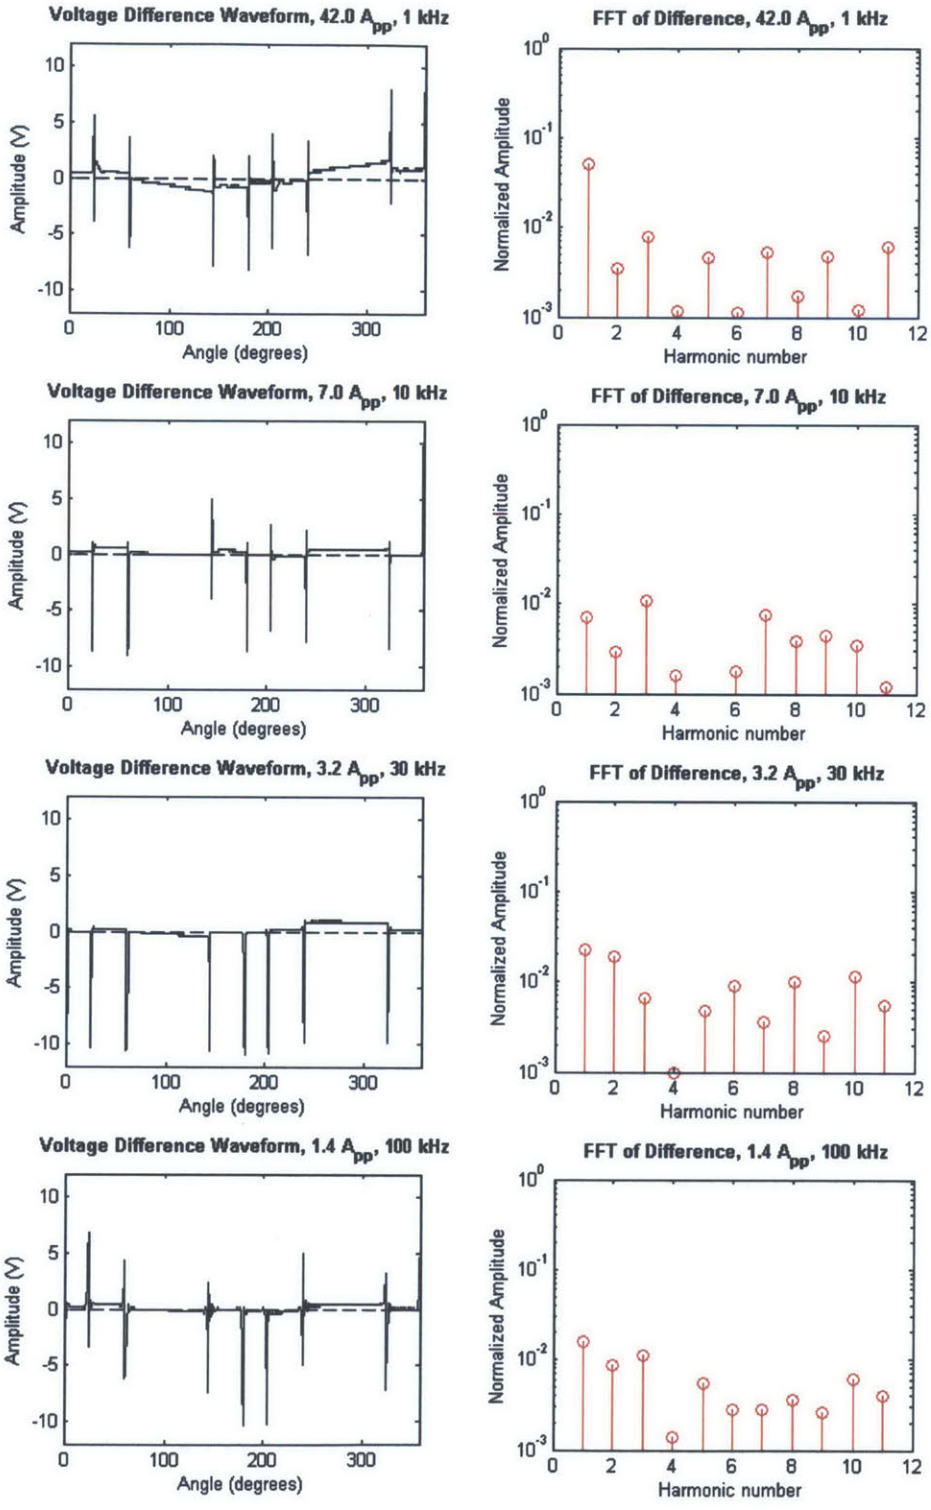

Figure 5.8: (a) the difference between the ideal and achieved voltage waveforms across the low-mid current range; (b) corresponding FFTs of the difference waveforms.

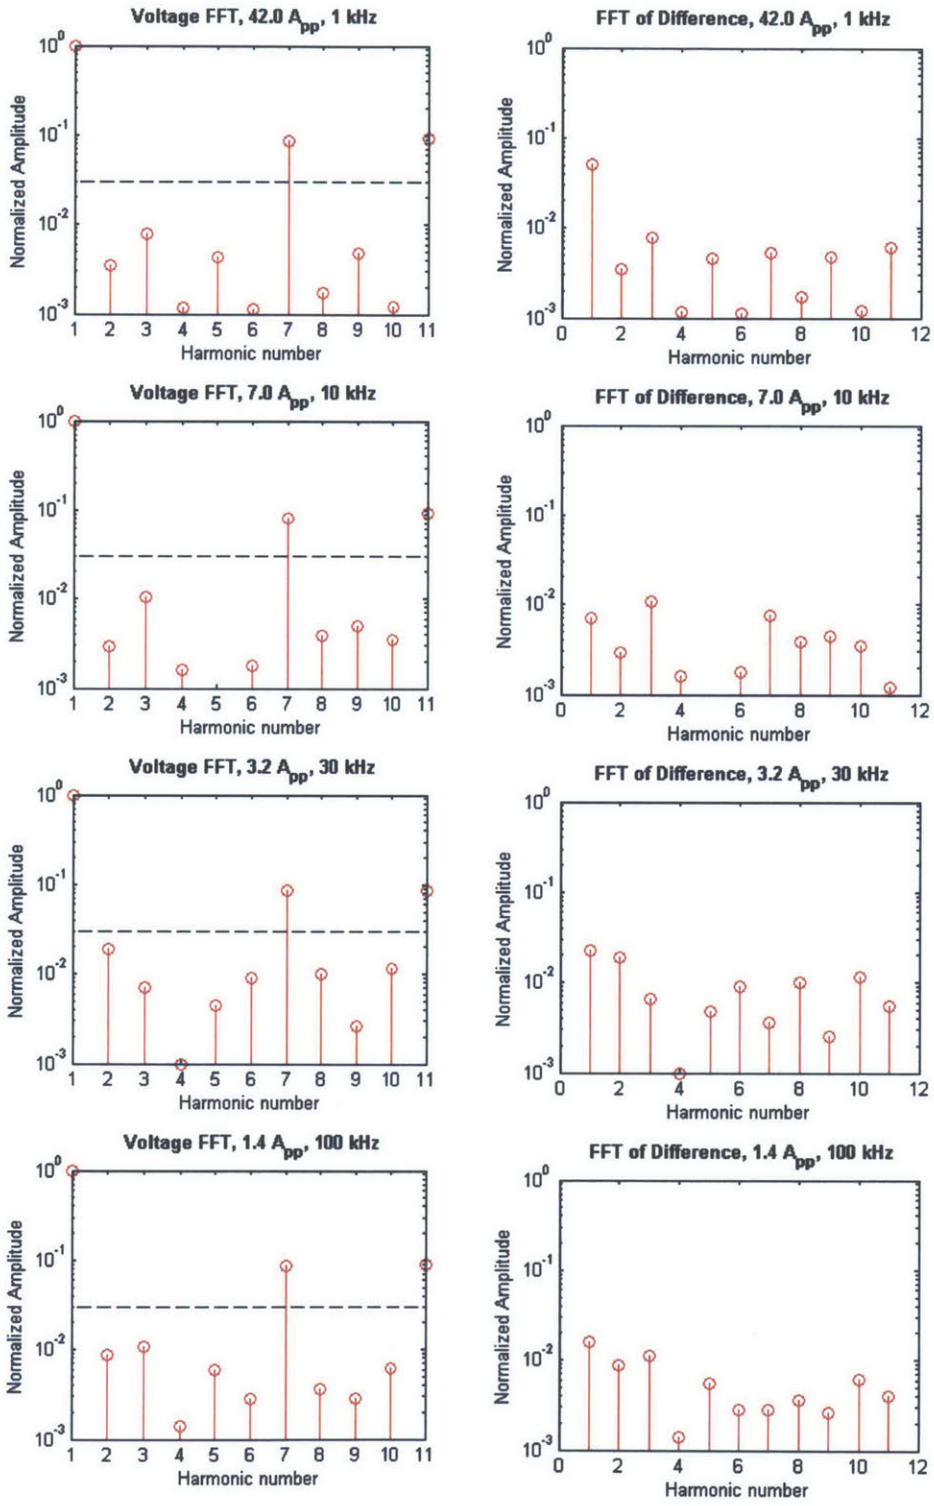

Figure 5.9: Comparison of the FFTs of the voltage waveforms (a) and the difference waveforms (b).

5.2 High Current Test

The purpose of the second test is to demonstrate the ability of the test bench to operate at current levels where the load's nonlinearity can be easily observed. The test bench was designed to operate at up to 2 kW, where it is expected that typical induction heating loads approach saturation. An indicator of load saturation is an increase in the odd current harmonics, so it is not necessary to measure the load impedance to detect load saturation [25, 26].

5.2.1 Experiment Design

The experimental process consisted of measuring the voltage and current waveforms across the inductive load for 1-100 kHz in 10 kHz increments at 160 V_{pp}, such that the maximum power level of the test bench was achieved. The experimental setup was the same as for the first experiment. Across the test bench's frequency range of 1-100 kHz, at 160 V_{pp} the corresponding current and power ranges were 132.8-4.5 A_{pp} and 1807.8-169.7 W. The voltage and current waveforms produced by the test bench at the highest power level are shown in Fig. 5.10. As in the first experiment, the sinusoidal quality of the waveforms is observed by taking the FFT of the waveforms using Matlab, normalized such that the amplitude of the voltage fundamental is 1.

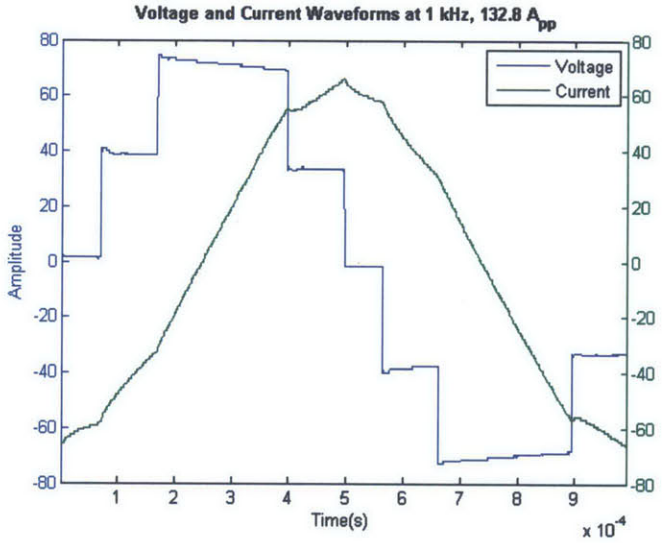

Figure 5.10: Voltage and current of the pan-loaded coil at 160 Vpp and 1 kHz, 132.8 A_{pp}.

5.2.2 Main Results

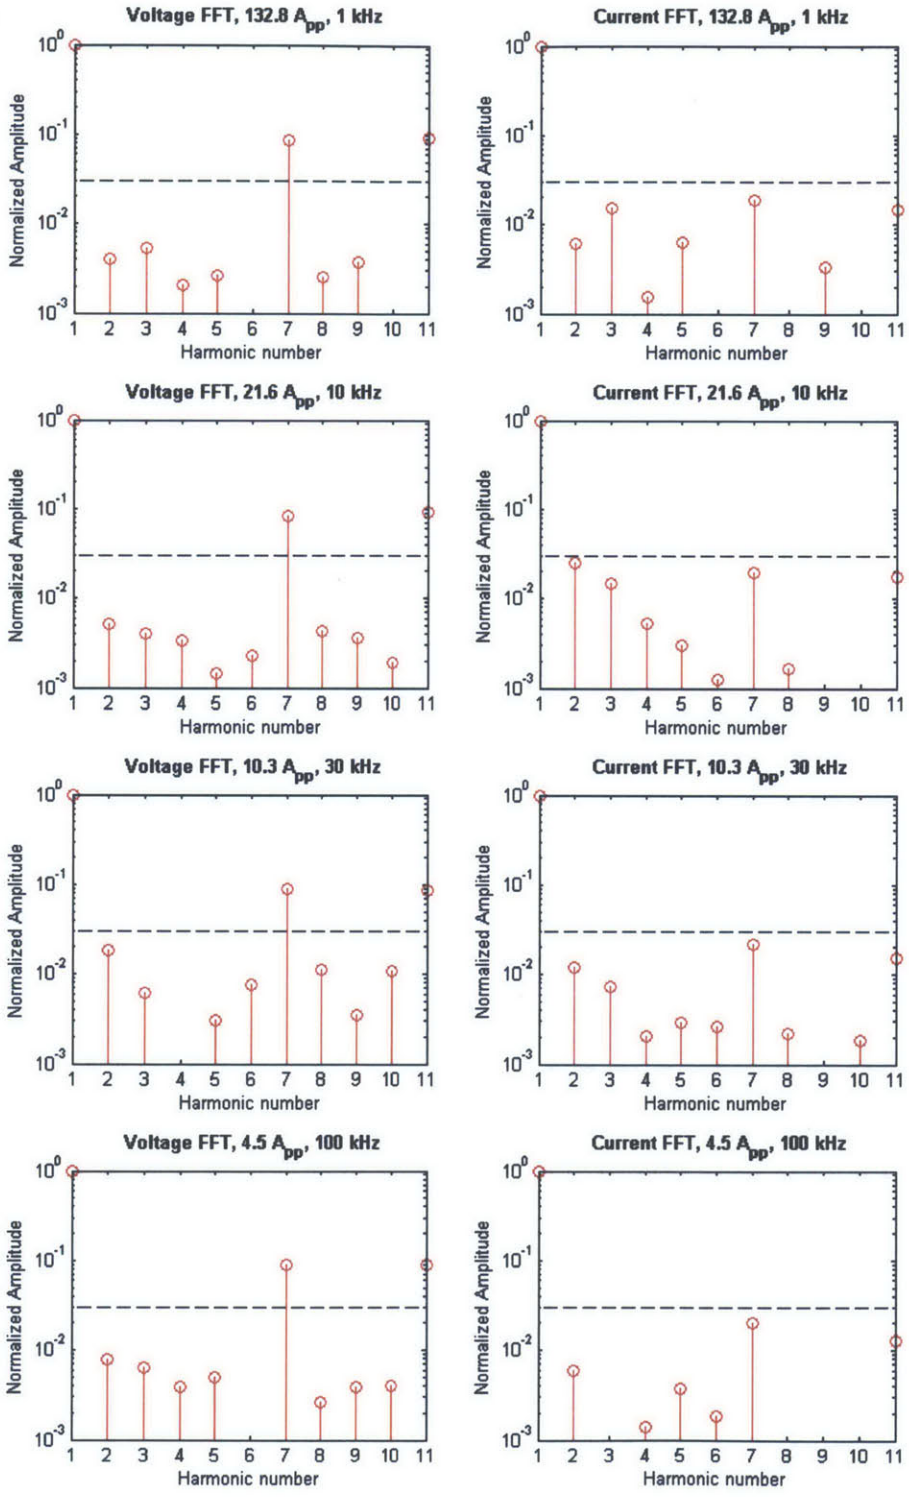

Figure 5.11: Samples of FFTs taken at 160 V_{pp} across 132.8-4.5 A_{pp}, corresponding to 1-100 kHz and 1.8 kW-69.7 W.

The sinusoidal quality of the voltage and current waveforms is illustrated by the FFTs of the waveforms. Having ascertained in Sec. 5.1 that the test bench generates voltage and current waveforms in which all 2n, 3n, and 5n harmonics are cancelled to below 3% of the fundamental, we expect to see similar harmonic content in this test. However, unlike the low current test, we expect the odd current harmonics to increase with the current level due to saturation of the magnetic material. A sample of the resulting FFTs taken at currents across the tested current range is shown in Fig. 5.11.

The conclusions made in Sec. 5.1.2 extend to the case of maximum power: the test bench produces excitations that are near-sinusoidal for the purposes of induction heating load characterization, despite minor harmonic distortion due to waveform non-idealities. Overall, the FFTs produced during this test are similar to the FFTs produced at lower power levels: the magnitudes of most of the 2n, 3n, and 5n harmonics are below 3% of the magnitude of the fundamental.

Furthermore, Fig. 5.11 shows the predicted rise in odd current harmonics with increasing current. The increase is most noticable in the third current harmonic; for 1 kHz at 132.8 A_{pp} and 1.8 kW, the magnitude of the third harmonic is 1.6% of the magnitude of the fundamental, whereas for the same frequency at 42 A_{pp} and 14 W (as measured in Sec. 5.1), the magnitude is only 1.0%. Although it is less obvious under visual inspection, the magnitudes of most of the odd harmonics are largest in the high power case. This increase in the odd harmonics can be observed in the current waveforms as a more defined peak, as shown in Fig. 5.12.

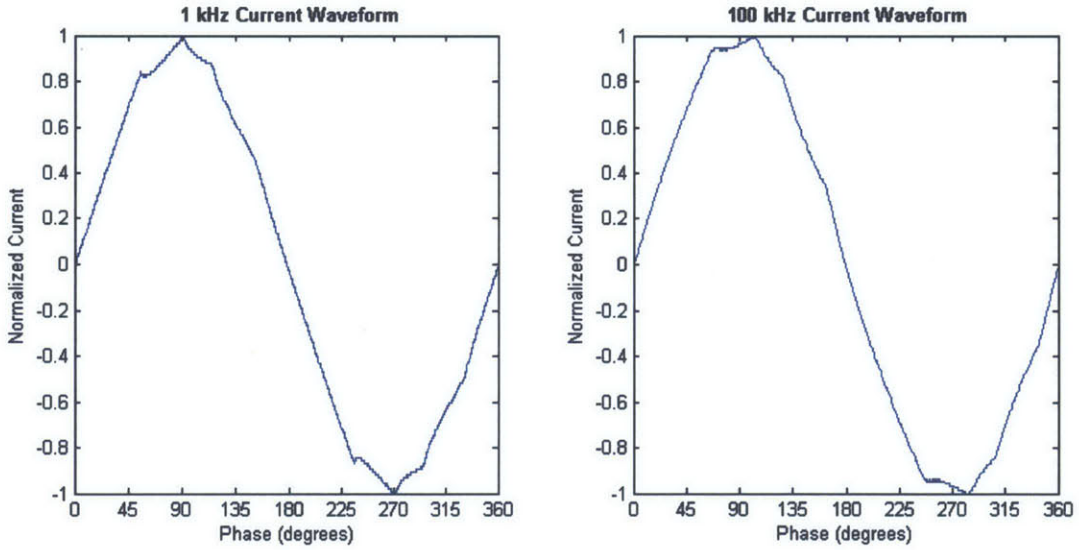

Figure 5.12: Current waveforms at 160 Vpp taken at 132.8 A_{pp} (1 kHz, 1.8 kW), showing a defined peak, and at 4.5 A_{pp} (100 kHz, 70 W), showing a less defined peak.

We can test the hypothesis that the increase in odd current harmonics is caused by load saturation. First, if the load is driven towards saturation, then we expect the odd current harmonics to increase, but the corresponding voltage harmonics to be unaffected. Second, if the experiment is repeated with the magnetic material removed, then we expect the odd current harmonics to be unaffected by the increasing current level. To validate the hypothesis that saturation of the load is

the source of the increase in the odd current harmonics, these tests are presented in the following two sections.

5.2.3 Comparison of Voltage & Current Harmonics

First, we tested the prediction that when the load approaches saturation, the odd current harmonics increase but the corresponding voltage harmonics are unaffected. This was done by comparing the odd current harmonics with the odd voltage harmonics across a range of currents, which corresponds to a range of frequencies. We expect the magnitude of the current harmonics to increase with decreasing frequency, but the magnitude of the voltage harmonics to be unaffected. The difference in the behavior of the odd harmonics of the voltage and current waveforms collected in the above section (Section 5.2.2) can be observed in Fig. 5.13.

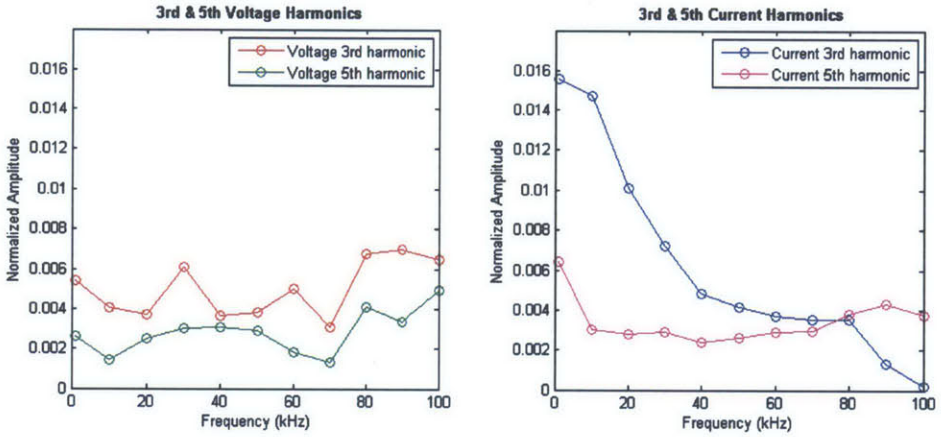

Figure 5.13: The largest odd harmonics of the voltage and current waveforms for the pan-loaded coil, taken at 160 V_{pp} across 132.8-4.5 A_{pp}, corresponding to 1-100 kHz and 1.8 kW-69.7 W.

As expected, the odd voltage harmonics are largely flat across the range of frequencies tested, while the odd current harmonics show a rise with decreasing frequency (increasing current). This agrees with the hypothesis that the pan approaches saturation at large currents, causing the odd current harmonics to increase in magnitude. While the rise in the third current harmonic is more significant than in the fifth current harmonic, the fifth harmonic does have a spike at 1 kHz, suggesting that it is beginning to saturate. Because no saturation behavior is observed in the voltage harmonics, the variation in voltage harmonic amplitudes is caused solely by the waveform defficiencies described in the low power test in Sec. 5.1.3.

5.2.4 Comparison of Harmonics with Air-loaded Coil & Pan-loaded Coil

Second, we validated the prediction that when the magnetic material is removed, the odd current harmonics will be unaffected by the current level. This was done by comparing the odd current harmonics of the pan-loaded coil with the odd current harmonics of an air-loaded coil, a coil with nothing resting on top of it. In order to make this comparison, we duplicated the experiment described in Sec. 5.2.1 using an air-loaded coil. For this load, across the test bench's frequency

range of 1-100 kHz, at 160 V_{pp} the corresponding current and power ranges were 126.4-2.2 A_{pp} and 1225-33 W. These ranges are comparable to those achieved in the original experiment. The difference between the odd current harmonics for the air-loaded coil and the pan-loaded coil can be observed in Fig. 5.14.

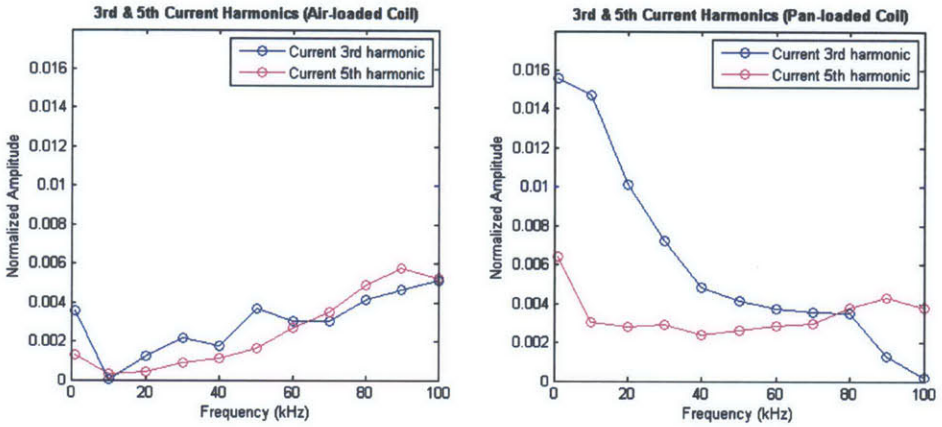

Figure 5.14: The largest odd current harmonics for the air-loaded coil and the pan-loaded coil taken at 160 V_{pp} across 126.4- 2.2 A_{pp}, corresponding to 1-100 kHz and 1.2 kW-33 W.

As expected, when the nonlinear load material is removed the odd current harmonics do not show a rise with decreasing frequency (increasing current). In fact, the odd harmonics in the air-loaded case even tend to decrease with decreasing frequency. The lack of increase in the odd current harmonics in the air-loaded test agrees with the hypothesis that the observed increase in odd current harmonics is due to nonlinear behavior of the pan.

Chapter 6

Load Characterization Case Study

In this thesis, we have designed and constructed a test bench for the purpose of induction heating load characterization. In the previous chapter, we demonstrated that the test bench is capable of meeting the established requirements: production of near-sinusoidal current excitations across inductive loads from 1–100 kHz and up to 2 kW power.

This chapter demonstrates how this test bench can be used for load characterization experiments: a case study is provided of the pan-loaded coil load used in Chapter 5, Circuit Performance. Magnetization curves are produced using data collected with the test bench.

6.1 Method

Data was collected by driving the load across a range of frequencies and currents using this test bench. At each frequency-current operating point, the voltage and current waveforms were measured. Matlab was then used to process the data to produce the magnetization curves. The code is shown in Appendix E.

Due to the resolution of our measurement apparatus, the current probe and oscilloscope, the voltage and current curves exhibited undesired quantization. This can be observed in the "Original Voltage and Current" waveforms in Fig. 6.1 as small steps in the curves that do not reflect actual voltage and current behavior. In order to correct for this, the data was dithered, filtered using a 3rd order finite impulse response (FIR) filter, and resampled at a slower sample rate.[1] The resulting waveforms, which more accurately express the voltage and current, are shown in the "Filtered Voltage and Current Data" plot in Fig. 6.1.

[1]The reduction factor between the original and new sampling rate was determined based on how much it was oversampled, which is proportional to the number of data points in the period. For example, the 1 kHz data was ~2500 data points, and a factor of $\frac{1}{9}$ of the original sampling rate was used; the 30 kHz data was ~1600 data points, and a factor of $\frac{1}{6}$ was used.

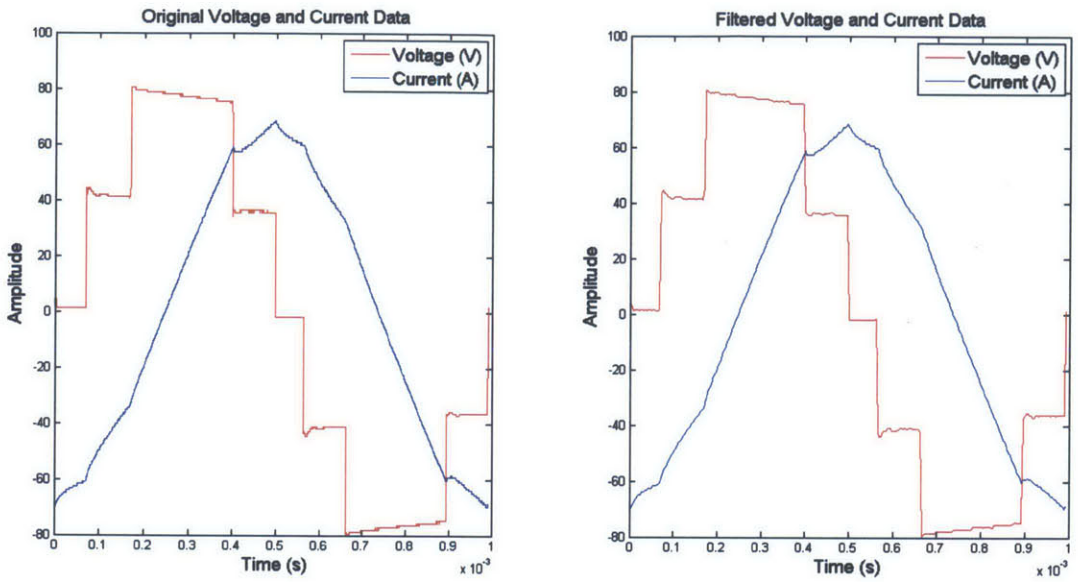

Figure 6.1: Raw and filtered voltage and current waveforms.

6.2 Magnetization Curves

The magnetic properties of the pan-loaded coil can be observed by plotting the flux linkage $\lambda = \int V \cdot dt$ as a function of the current I. Fig. 6.2 shows sample magnetization curves for this load for two cases: Fig. 6.2a shows the magnetization curves for various current levels at a constant frequency of 1 kHz, and Fig. 6.2b shows the magnetization curves for various frequencies at a constant current level of 4.36 ± 0.06 A_{pp}. The data points shown in the figure were selected in order to span the range of frequencies and currents tested.

The overall shapes of the λ-I curves shown in Fig. 6.2 do not match the typical hysteresis curve shape. Specifically, one unexpected phenomenon exhibited by all of the λ-I curves is an atypical hook-shaped behavior that occurs prior to the positive and negative peak current. It can be observed in the figure that this behavior is most significant at large current levels and low frequencies. The hook corresponds to peaking in the current waveform when the voltage waveform changes levels, which can be observed in Fig. 6.1. The physical cause of this behavior is unclear, and further investigation is required to better understand this phenomenon.

Other than the hook behavior, Fig. 6.2 exhibits various characteristics of interest. First, the λ-I curves do not exhibit saturation; for each of the frequency-current operating points shown, as I is increased, λ does not exhibit any apparent decrease. However, saturation remanence and hysteresis are observed in all cases. The saturation remanence and the amount of hysteresis —the area inside the λ- curve— increase with increasing peak current and decreasing frequency. Additionally, the shape of the λ-I curve exhibits little change in shape across different current levels but does change with frequency. In Fig. 6.2a, the curves at different currents appear to be scaled versions of each other. However, in Fig. 6.2b, as the current increases the B-H curve rotates clockwise in the λ-I

plane and becomes thinner.

One source of uncertainty in interpreting the magnetization curves —particularly the unexpected hook-shaped behavior, but also the saturation, saturation remanence, and hysteresis— is the effects of the coil impedance in the λ-I curves. The inductance coil has both resistance and inductance that influence the shape of the λ-I curves. The measurements used to calculate these curves are the voltage and current across the pan-loaded coil; the impedance of the induction coil affects the magnetization curves. Characterization of the induction coil λ-I curves would allow for a more knowledgable interpretation of the hook-shaped behavior observed in Fig. 6.2.

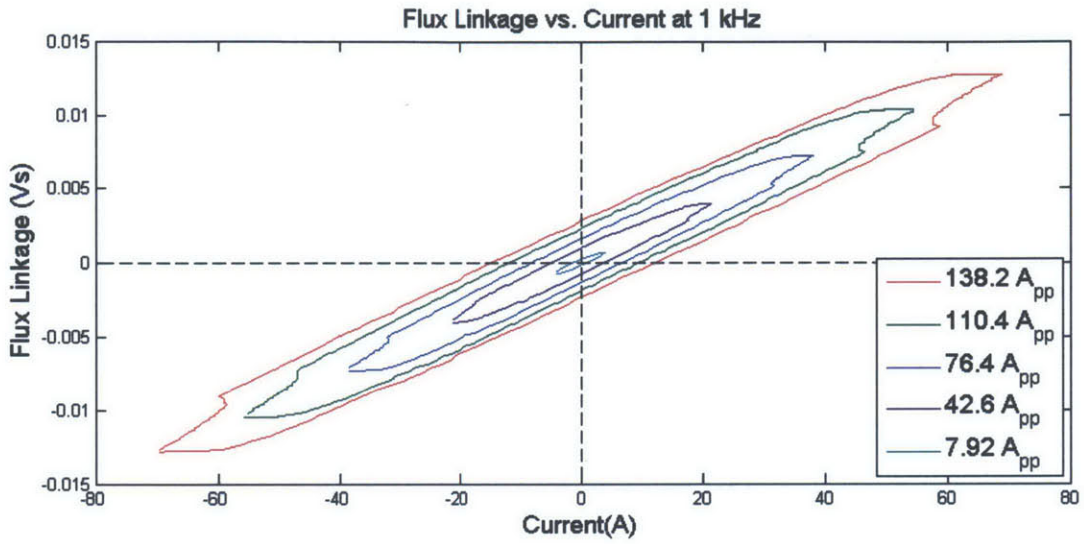

(a)

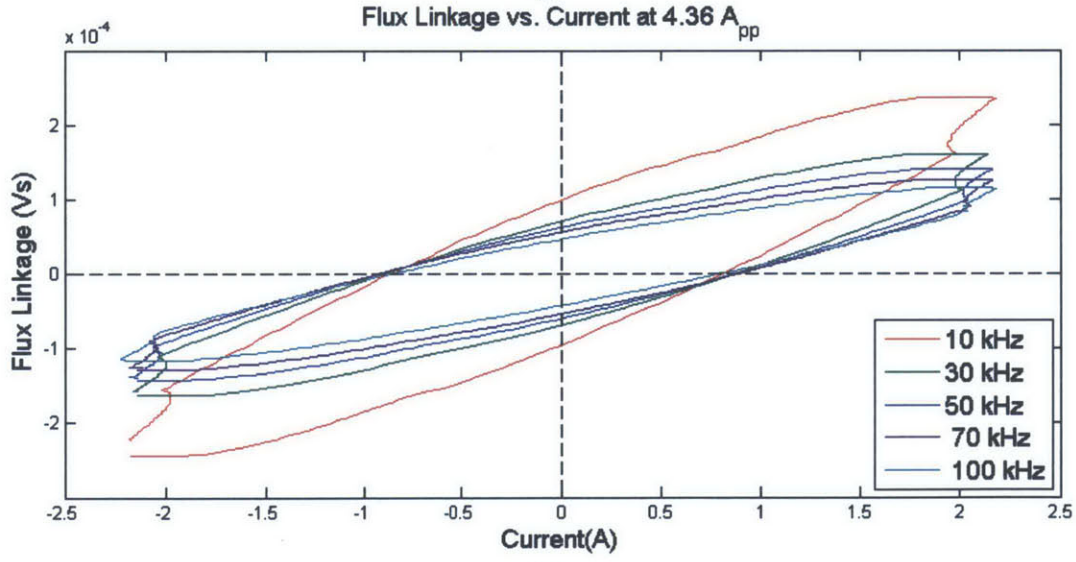

(b)

Figure 6.2: Plots of the flux linkage λ vs. current I for (a) various current levels at a constant frequency of 1 kHz, and (b) various frequencies at a constant current level of $4.36 \pm 0.06\, A_{\mathbf{pp}}$.

Chapter 7

Conclusion & Future Work

In this thesis, we have presented a novel test bench for inductive load characterization across a wide range of currents and frequencies. The test bench is capable of driving inductive loads with near-sinusoidal voltage and current excitations across 1–100 kHz at up to 2 kW power. The use of the novel CHBFB multilevel inverter topology and a switching scheme that provides harmonic cancellation enables our test bench to provide low-distortion waveforms across the operating range.

We have demonstrated that our test bench is capable of characterizing induction heating loads by conducting a case study of a custom induction coil loaded with a cast iron frying pan. We generated magnetizing curves for the load using voltage and current waveforms acquired across the current and frequency range of the test bench.

The case study suggests further investigation into the effects of the coil impedance on the magnetization curves. Specifically, we recommend characterization of the magnetization curves of the pan alone. This can be achieved by characterizing the curves of the unloaded induction coil. Producing the induction coil magnetization curves will allow for the isolation of the magnetization curves of the pan. Absent the linear behavior of the induction coil, these curves will more clearly depict the nonlinear characteristics of the pan's magnetic material.

Finally, we also recommend investigation into the effects of magnetic nonlinearities on power dissipation in domestic induction cooktop loads, which was the primary motivation behind this thesis. While the impedance of a pan-loaded coil is known to change with frequency, the load impedance is typically assumed constant across the current levels used in domestic induction heating applications [2, 3, 7]. The measurement of the load impedance at both low and high current levels across the test bench's frequency range would allow for the validation or invalidation of this assumption. Currently, at high frequency levels the test bench is not able to produce the high current needed for this experiment. However, using more powerful dc power supplies or adding a a series capacitor to the induction coil would allow for higher current at high frequencies.

Appendix A

Transition Waveforms

Finally, the MOSFET transitions exhibit acceptable speed and overshoot. Samples of gate-source and drain-source transition waveforms taken with a drain-source voltage of 40 V at the minimum and maximum frequencies are shown in Fig. A.1. The transition waveforms show several noteworthy characteristics:

- The MOSFET turn-off transition is less clean than the turn-on transition at both 1 kHz and 100 kHz, and particularly at 1 kHz. At 1 kHz, the drain-source voltage exhibits peak overshoot of $\approx$2.4, corresponding to a peak voltage of 97.6 V, and significant subsequent ringing that lasts approximately 200 ns. Because the MOSFETs are rated to 100 V, the amount of overshoot does not cause MOSFET breakdown, and the test bench is able to continue functioning. Furthermore, the duration of the ringing is short compared to the period of the signal, which is 1 ms. Therefore, although undesired, the amount of overshoot and ringing is acceptable.
- The 1 kHz transition waveforms are uniformly less clean than the 100 kHz waveforms. This is because the load draws significantly more current at 1 kHz (42 A_{pp}) than at 100 kHz (1.4 A_{pp}).
- At both 1 kHz and 100 kHz, the drain-source voltage rise times during MOSFET turn-off were approximately 50 ns, and the drain-source voltage fall times during MOSFET turn-on were less than 65 ns.
- As expected, the fall times of the gate-source voltage are shorter than the rise times due to the anti-parallel gate diode.
- The delay between rise/fall of the gate-source voltage and fall/rise of the drain-source voltage is approximately 100 ns.

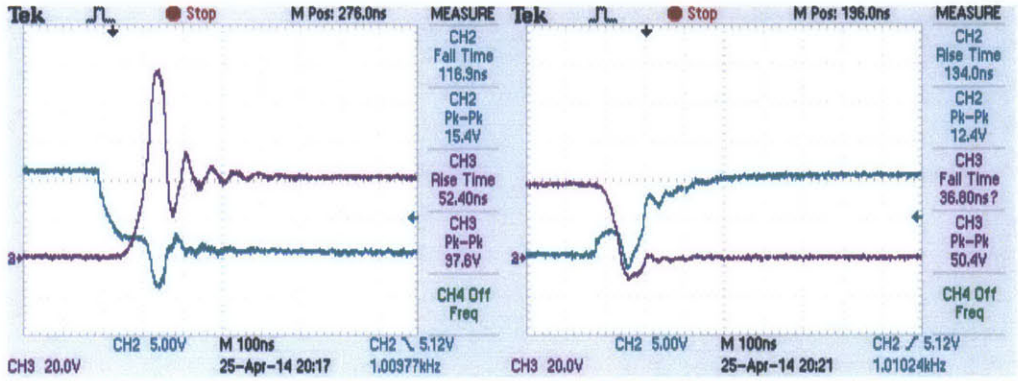

FBB low-side turn-off 40 Vds

(a) Turn-off at 1 kHz

FBB low-side turn-on

(b) Turn-on at 1 kHz

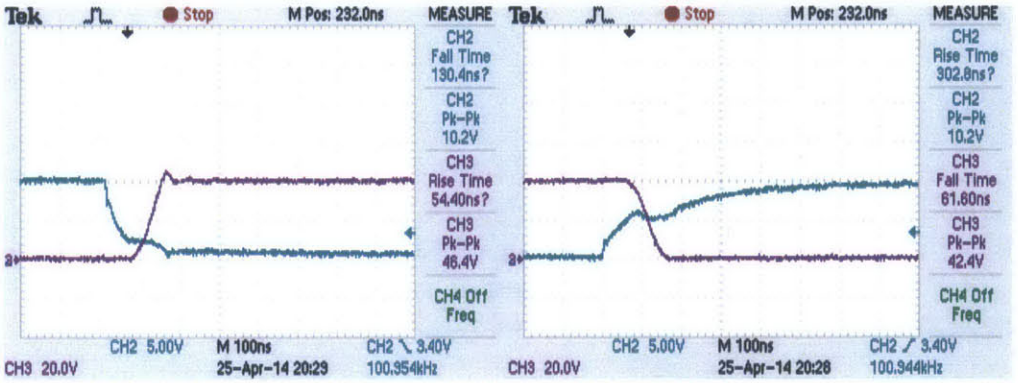

FBB low-side turn-off

(c) Turn-off at 100 kHz

FBB low-side turn-on

(d) Turn-on at 100 kHz

Figure A.1: Samples of gate-source (blue, channel 2) and drain source (purple, channel 3) waveforms collected at a drain-source voltage of 40 V at 1 kHz (corresponding to maximum power conditions) and 100 kHz. Waveforms collected from the low-side MOSFET of the side B of the full bridge (FB_B). For this measurement, low-inductance probe tips were used.

Appendix B

Microcontroller Code

```
////////////////////////////////////////////////////////////////////////////
// Description:
// This code is written in C. It programs the Microchip dsPIC33FJ16GS502 microcontroller.
// This program produces independent PWM outputs on pins PWM1L, PWM2L, PWM3L, and PWM4L
// (which is on pin RP7). When timer1 overflows, an ADC measurement is taken of the
// potentiometer. The frequency, duty cycles, and phase shifts are calculated based on
// the value of the  pot. The frequency range is from 20kHz - 100kHz in 10kHz steps.
////////////////////////////////////////////////////////////////////////////

//Include header files.
#include "p33FJ16GS502.h"
#include "FBcontrol.h"

//Set Configuration Bits.
_FOSCSEL(FNOSC_FRC) //System Clock Selection: Fast RC internal oscillator: Fin=7.37MHz
_FOSC(FCKSM_CSECMD & OSCIOFNC_ON %IOL1WAY_OFF)
_FWDT(FWDTEN_OFF)
_FPOR(FPWRT_PWR128)
_FICD(ICS_PGD1 & JTAGEN_OFF)    //Communication Channel 1 is selected.

//Main code.
int main() {
// Set up I/O
    TRISAbits.TRISA0 = 1;       // AN0/RA0 (pot) input
    TRISB=0;                    //Port B is PWM output.

    __builtin_write_OSCCONH(0x01);      // New Oscillator FRC w/ PLL
    __builtin_write_OSCCONL(0x01);

    //Set RP bits as PWM4 outputs.
    RPOR3=0b0010110100101100;   //RP7 is PWM4L, RP6 is PWM4H.
                                //This line must occur after write to OSCCONL.

//Set up internal oscillator.
    while(OSCCONbits.COSC != 0b001);        // Wait for new Oscillator to become FRC w/ PLL
     while(OSCCONbits.LOCK != 1);           // Wait for PLL to Lock

//Set up the ADC and PWM clock.
    ACLKCONbits.FRCSEL = 1;         // FRC provides input for Auxiliary PLL
    ACLKCONbits.SELACLK = 1;        // Auxiliary Oscillator provides clock source for PWM & ADC
    ACLKCONbits.APSTSCLR = 0b111;   // N = 1
    ACLKCONbits.ENAPLL = 1;         // M1=16
```

```
                                    // ACLK = (FRC * M1) / N = (7.37MHz * 16) / 1 = 117.92MHz
                                    // In the 1-10 kHz program, ACLK is set to 7.37 MHz.

    while(ACLKCONbits.APLLCK != 1);         // Wait for Auxiliary PLL to Lock
    ADC_Init();

//Set up timer (that triggers the interrupt in which the ADC is read).
    Timer1_Init();

//Set up the PWM module.
    init_PWM();

//Set up interrupts.
    EnableInterrupts();

//MAIN PROGRAM: Infinite loop.
    while(1);   {
    }
}


//Set up the PWM module.
void init_PWM()
{
    //Set up PWM Master Time Base Register.
    //PTPER = (ACLK * 8)/(PCLKDIV * fpwm) = (117.92MHz * 8)/(1 * 100.02kHz) = 9432
    PTPER = 9434;

    //Configure PWM1.
    IOCON1bits.PENH      = 0;               //PWM1H is not a PWM output pin.
    IOCON1bits.PENL      = 1;               //PWM1L is a PWM output pin.
    IOCON1bits.PMOD      = 1;               //PWM Redundant Output Mode.
    PWMCON1bits.MDCS=0;                     //PDC1 sets duty cycle.
    PDC1=PTPER/2;
    DTR1     = 0;                           //No deadtime.
    ALTDTR1 = 0;                            //No deadtime.
    PHASE1 = 0;                             //No phase shift on PWM1L

    //Configure PWM2.
    IOCON2bits.PENH      = 0;               //PWM2H is not a PWM output pin.
    IOCON2bits.PENL      = 1;               //PWM2L is a PWM output pin.
    IOCON2bits.PMOD      = 1;               //PWM Redundant Output Mode.
    PWMCON2bits.MDCS=0;                     //PDC2 sets duty cycle.
    PDC2=PTPER/2;
    DTR2     = 0;                           //No deadtime.
    ALTDTR2 = 0;                            //No deadtime.
    PHASE2 = 0;                             //No initial phase shift.

    //Configure PWM3.
    IOCON3bits.PENH      = 0;               //PWM3H is not a PWM output pin.
    IOCON3bits.PENL      = 1;               //PWM3L is a PWM output pin.
    IOCON3bits.PMOD      = 1;               //PWM Redundant Output Mode.
    PWMCON3bits.MDCS=0;                     //PDC3 sets duty cycle.
    PDC3=PTPER/2;
    DTR3     = 0;                           //No deadtime.
    ALTDTR3 = 0;                            //No deadtime.
    PHASE3 = 0;                             //No initial phase shift.

    //Configure PWM4.
    IOCON4bits.PENH      = 0;               //PWM4H is not a PWM output pin.
    IOCON4bits.PENL      = 1;               //PWM4L is a PWM output pin.
    IOCON4bits.PMOD      = 1;               //PWM Redundant Output Mode.
    PWMCON4bits.MDCS=0;                     //MDC sets duty cycle.
```

```
    PDC4=PTPER/2;
    DTR4    = 0;                            //No deadtime.
    ALTDTR4 = 0;                            //No deadtime.
    PHASE4 = 0;                             //No initial phase shift.

    //Enable PWM Module.
    PTCONbits.PTEN       = 1;
}


//Set up timer.
void Timer1_Init(void)
{
    T1CONbits.TCKPS = 0b11;         // Prescale=256: Fosc=39.6 MHz/256=154.7 kHz (this is arbitrary)
    TMR1 = 0;                       // Initialize Timer 1.
    T1CONbits.TON = 1;              // Start timer.
}


//Set up ADC.
void ADC_Init(void)
{
    ADPCFGbits.PCFG0=0;              //PCFG0=AN0 is enabled for ADC.
    ADCONbits.SLOWCLK = 1;           //ADC is clocked by auxiliary PLL.
    ADSTAT = 0;                      // Clear the ADSTAT register.
                                     //This register contains status of data-ready in buffer.
    ADCPC0bits.TRGSRC0=0b00001;      // Individual software triggers of ADC of channels AN1 and AN0.
    ADCONbits.ADON=1;                // Turn on ADC.
}


//Set up interrupts.
void EnableInterrupts(void)
{
     IPC0bits.T1IP = 0b111;          // Timer1 interrupt highest priority.
    IFS0bits.T1IF = 0;               // Clear Timer1 interrupt flag.
    IEC0bits.T1IE = 1;               // Timer1 interrupting enabled.
}


//Interrupt Service
#pragma interrupt InterruptService
void __attribute__((interrupt, no_auto_psv)) _T1Interrupt(void) //Timer 1 ISR
{
    // Check for AccumulatorA Trap Flag Interrupt (if timer overflowed).
    if  (IFS0bits.T1IF)
    {
        //Clear Timer1 overflow flag.
        IFS0bits.T1IF = 0;

        //Take ADC measurement.
        ADCPC0bits.SWTRG0=1;                //Start conversion of AN0.
        while(ADCPC0bits.PEND0 != 0);       //Wait for ADC to be completed.

        //Calculate fpwm based on ADC value.
        //Harmonic Reduction Scheme. These values roughly follow the equations in Ch. 4,
        // but they have been fine-tuned, so they do not match perfectly.
        if (ADCBUF0<200)
            {PTPER=47168;   //min freq: 20 kHz
            PDC1=24128;
            PDC2=12578;
            PDC3=23206;
            PDC4=12578;
```

```
    PHASE1 = 0;
    PHASE2 = 4626;
    PHASE3=19989;
    PHASE4=27755;}

else if (ADCBUF0<300)
    {PTPER= 31438;  //30 kHz
    PDC1=16074;
    PDC2=8383;
    PDC3=15453;
    PDC4=8383;
    PHASE1 = 0;
    PHASE2 = 3089;
    PHASE3=13353;
    PHASE4=18534;}

else if (ADCBUF0<400)
    {PTPER=23584;       //40 kHz
    PDC1=12078;
    PDC2=6289;
    PDC3=11607;
    PDC4=6289;
    PHASE1 = 0;
    PHASE2 = 2317;
    PHASE3= 10010;
    PHASE4=13899;}

else if (ADCBUF0<500)
    {PTPER=18860;   //50 kHz
    PDC1=9661;
    PDC2=5029;
    PDC3=9297;
    PDC4=5029;
    PHASE1 = 0;
    PHASE2 = 1853;
    PHASE3=8016;
    PHASE4=11121;}

else if (ADCBUF0<600)
    {PTPER=15720;       //60 kHz
    PDC1=8066;
    PDC2=4192;
    PDC3=7741;
    PDC4=4192;
    PHASE1 = 0;
    PHASE2 = 1544;
    PHASE3= 6673;
    PHASE4=9262;}

else if (ADCBUF0<700)
    {PTPER=13475;   //70 kHz
    PDC1=6957;
    PDC2=3593;
    PDC3=6638;
    PDC4=3593;
    PHASE1 = 0;
    PHASE2 = 1319;
    PHASE3=5703;
    PHASE4=7912;}

else if (ADCBUF0<800)
    {PTPER=11792;       //80 kHz
    PDC1=6038;
```

```
            PDC2=3145;
            PDC3=5801;
            PDC4=3145;
            PHASE1 = 0;
            PHASE2 = 1162;
            PHASE3= 5018;
            PHASE4=6970;}

        else if (ADCBUF0<900)
            {PTPER=10480;    //90 kHz
            PDC1=5379;
            PDC2=2795;
            PDC3=5169;
            PDC4=2795;
            PHASE1 = 0;
            PHASE2 = 1030;
            PHASE3=4453;
            PHASE4=6182;}

        else
            {PTPER=9432;     //max freq: 100 kHz
            PDC1=4849;
            PDC2=2515;
            PDC3=4639;
            PDC4=2515;
            PHASE1 = 0;
            PHASE2 =928;
            PHASE3=4004;
            PHASE4=5568;}
    }
}   //Return from interrupt.
```

Appendix C

Circuit Details

C.1 Schematic

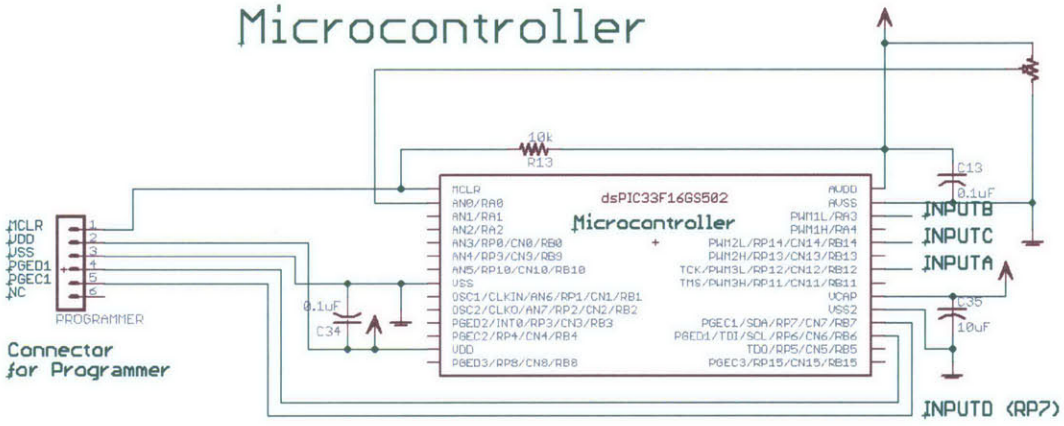

Figure C.1: Microcontroller schematic.

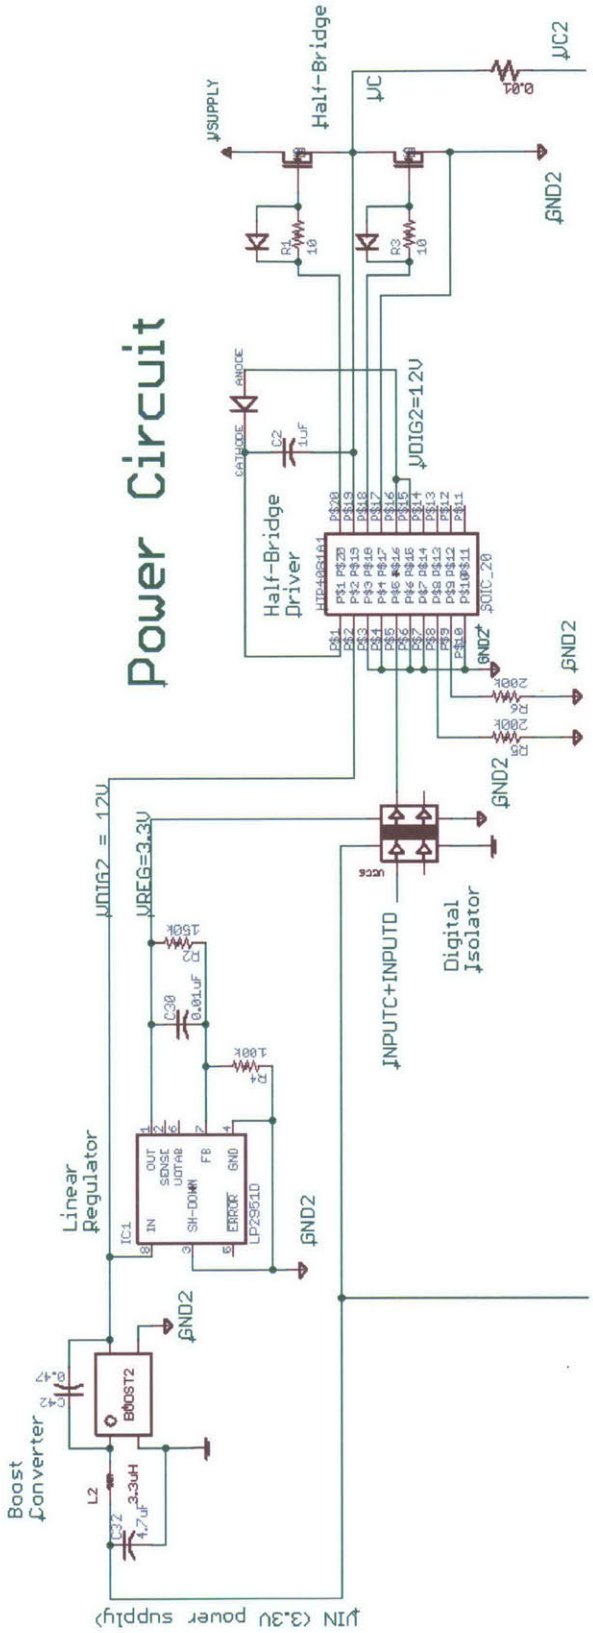

Figure C.2: Schematic of the inverter (continued in Fig. C.3).

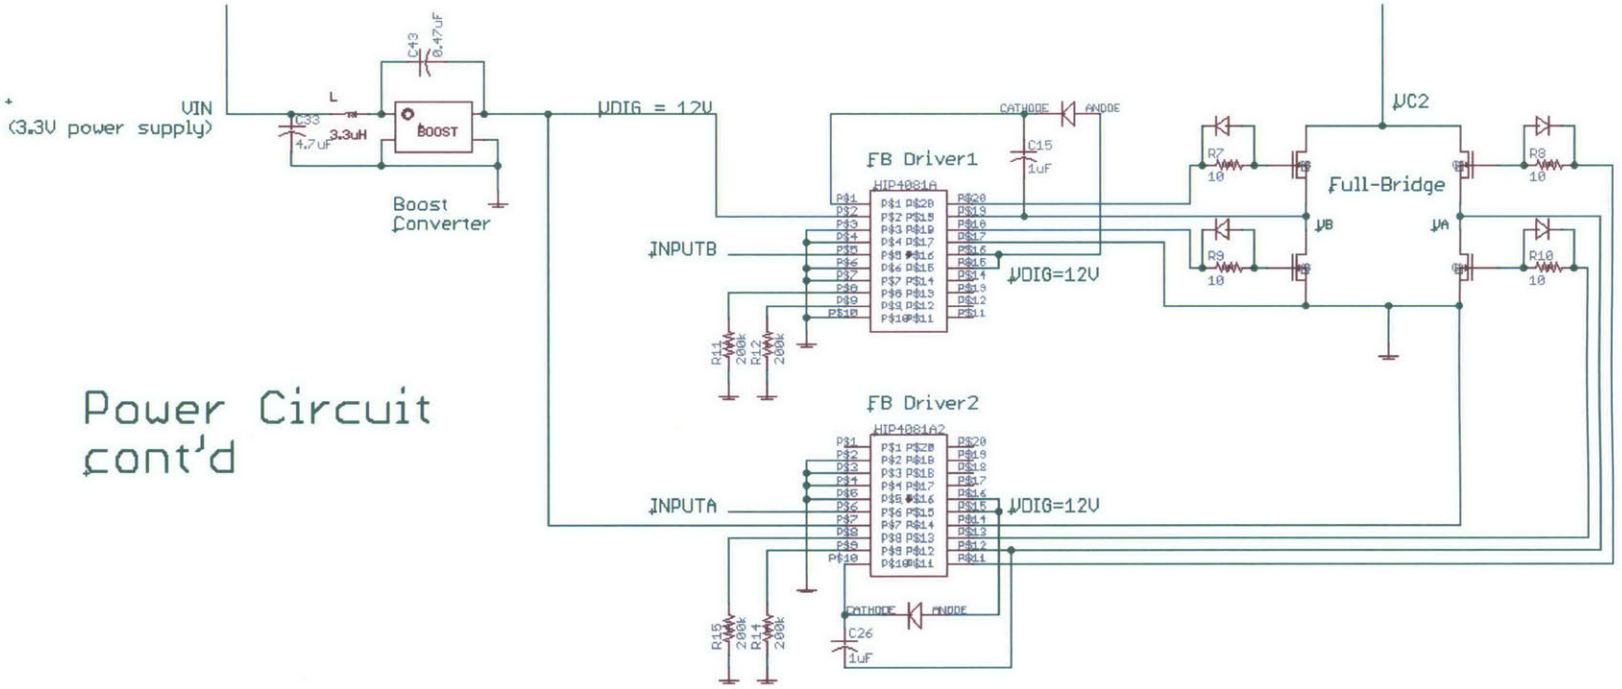

Figure C.3: Continuation of the inverter schematic.

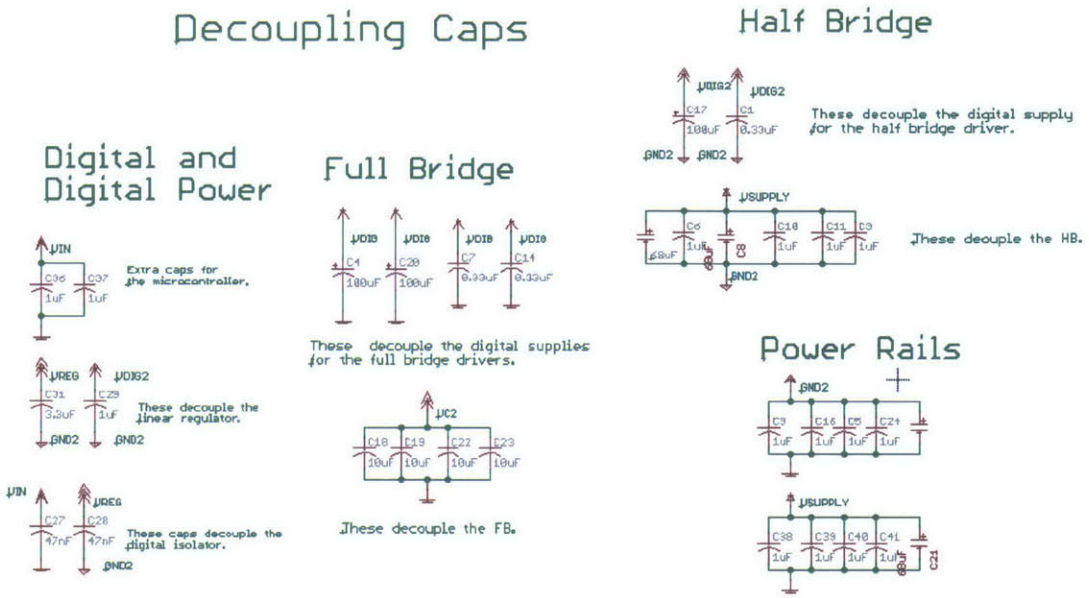

Figure C.4: Schematic of decoupling capacitors.

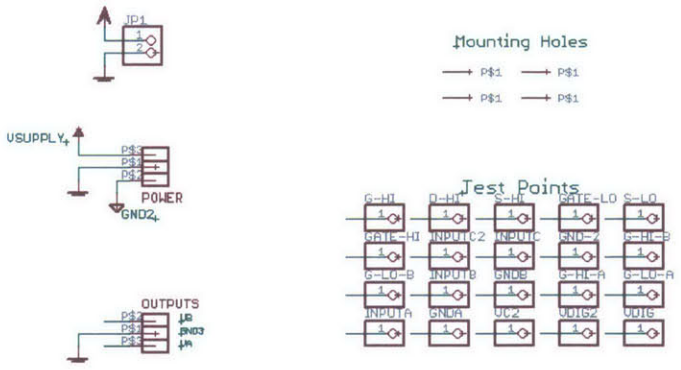

Figure C.5: Schematic of Miscellaneous components.

C.2 Selected Parts List

Component	Manufacturer Part Number	Notes
microcontroller	DSPIC33FJ16GS502-50I/SO	
potentiometer	PTV09A-4225F-B203	
decoupling capacitor	50SVPF68M	OS-CON capacitors recommended for decoupling capacitors because they have capacitance and high ripple current rating.
3.3 V–12 V isolated boost converter	R1S-3.312	
adjustable voltage regulator	LP2951DR	Set to regulate 12 V down to 3.3 V.
digital isolator	IL711T-1E	
driver	HIP4081AIBZ	
bootstrap capacitor	C2012X5R1H475K125AB	Two 4.7 μF capacitors used in parallel for bootstrap capacitor.
bootstrap diode	MSS1P6-M3/89A	
gate diode	DB2X20600L	
sense resistor	FCSL110R001JER	
MOSFET	TK100E10N1,S1X	
thermal pad	V7466Z	
heatsink	173-7-240A	

Table C.1: List of selected components used in circuit board. Standard components (resistors, capacitors, etc.) are not included, but they are shown on the schematic.

Appendix D

Fourier Transform Code

```
function [rmspower, irms, ipp] = fft_select

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% This function is written in Matlab. It performs several tasks:
%   1. Imports current and voltage data (one period must be selected manually);
%   2. Takes the FFT of the data using Matlab's DFT function;
%   3. Finds the frequency, value, and phase of the greatest component (i.e the
%       fundamental component);
%   4. Plots and returns these values.
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%Import data. In this code sample, only one data set is used. In the real code,
    %all the data is processed iteratively.
'f=1 kHz, Ipp=41.8 A';
voltage = importfile('Data\Pan-Coil Load Data 4-09-14.xlsx','Sheet1','J16:K2487');
    %One period must be selected manually by the user.
current = importfile('Data\Pan-Coil Load Data 4-09-14.xlsx','Sheet1','D16:E2487');
currentscale=importfile('Data\Pan-Coil Load Data 4-09-14.xlsx','Sheet1','B1');
    %The current is read by the oscilloscope as a voltage. The conversion to amps is
    %determined by the converstion ratio on the current probe amplifier. This conversion
    %ratio must be included in the data file, and it is read into Matlab as the
    %variable "currentscale".

%Calculate important values from the data.
current(:,2)=current(:,2)*currentscale/0.01;    %Calculate the current.
power=voltage(:,2).*current(:,2);
rmspower=rms(power);
irms=rms(current(:,2));
ipp=max(current(:,2))-min(current(:,2));
n=length(current);            %Data length

%Calculate FFT.
    %Voltage:
    fftv = fft(voltage(:,2));       %DFT
    fftv=2*fftv(2:n/2)/n;             %Normalized fft.
    magv=abs(fftv);   %magnitude of normalized FFT
    maxv = max(magv);   %Find max.

    %Current:
    fftc = fft(current(:,2));       %DFT
    fftc=2*fftc(2:n/2)/n;             %Normalized fft.
    magc=abs(fftc);   %magnitude of normalized FFT
    maxc = max(magc);   %Find max.
```

```
%Normalize FFT (for plotting purposes) s.t. the magnitude
    %of the fundamental harmonic is 1.
magv_norm=magv/maxv;
magc_norm=magc/maxc;

%Plot the data.
subplot (3,1,1)
[ax h1 h2] = plotyy(voltage(:,1),voltage(:,2),current(:,1),current(:,2));
set(ax(1),'xLim',[current(1,1) current(end,1)]) set(ax(2),'xLim',[current(1,1) current(end,1)])
xlabel('Time(s)')
ylabel('Amplitude')
title('{\bf Voltage and Current Waveforms, 1 kHz, 41.8 A_{pp}}')
legend('Voltage','Current','Location','Northeast')

%Plot the Voltage FFT.
subplot(3,1,2)
s1=stem(magv_norm, 'r')
set(gca,'yscal','log')
hbase1 = get(s1,'Baseline');
set(hbase1,'BaseValue',10^-4)
line([0 11],[0.03 0.03],'Color','k','LineStyle','--');
title('{\bf Voltage FFT, 1 kHz, 42.0 A_{pp}}')
xlim([1 11])
set(gca,'Xtick',[1 2 3 4 5 6 7 8 9 10 11])
ylim([1e-3 1])
xlabel('Harmonic number')
ylabel('Normalized Amplitude')

%Plot the Current FFT.
subplot(3,1,3)
s1=stem(magc_norm, 'r')
set(gca,'yscal','log')
hbase1 = get(s1,'Baseline');
set(hbase1,'BaseValue',10^-4)
line([0 11],[0.03 0.03],'Color','k','LineStyle','--');
title('{\bf Current FFT, 1 kHz, 42.0 A_{pp}}')
xlim([1 11])
set(gca,'Xtick',[1 2 3 4 5 6 7 8 9 10 11])
ylim([1e-3 1])
xlabel('Harmonic number')
ylabel('Normalized Amplitude')

end
```

Appendix E

Magnetization Curve Code

```
function [current1_up_avg, L1_up_avg]=L_1kHz
%This function is written in Matlab. It performs several tasks:
%  1. Imports current and voltage data (one period must be selected manually);
%  2. Filters and resamples the data;
%  3. Calculates the flux linkage;
%  4. Plots the magnetization curve: flux linkage vs. current.

%Define colors for plotting
red=rgb('Red');
green=rgb('Green');
blue=rgb('Blue');
turquoise=rgb('DarkTurquoise');
purple=rgb('Purple');

%Import data. In this code sample, only one data set is used. In the real code,
    %all the data is processed iteratively.
%Ipp=138.4 A
    current = importfile('Data\1 kHz Pan-Coil Load Data 5-13-14.xlsx','Sheet1','D19:E2492');
    currentscale=importfile('Data\1 kHz Pan-Coil Load Data 5-13-14.xlsx','Sheet1','B1');
    voltage = importfile('Data\1 kHz Pan-Coil Load Data 5-13-14.xlsx','Sheet1','J19:K2492');

%Make voltage ac.
voltage_dc=mean(voltage(:,2));
voltage(:,2)=voltage(:,2)-voltage_dc;

%Calculate current
current(:,2)=current(:,2)*currentscale/0.01;

%Dither.
n=length(current);
R=0.4*rand(n,1);  %0.4 is half the resolution
voltage_dith=voltage(:,2)+R;
R=0.*rand(n,1);    %0.4 is half the resolution
current_dith=current(:,2)+R;

%Filter by resampling at 1/9 the original sample rate and using a 3rd order FIR.
current_filt(:,1)=decimate(current(:,1),9,3,'fir');
current_filt(:,2)=decimate(current_dith,9,3,'fir');
voltage_filt(:,1)=decimate(voltage(:,1),9,3,'fir');
voltage_filt(:,2)=decimate(voltage_dith,9,3,'fir');

%Calculate flux density (lambda). Make ac.
lambda=cumtrapz(voltage_filt(:,1),voltage_filt(:,2));
```

```
lambda_dc=mean(lambda);
lambda=lambda-lambda_dc;

%Plot flux linkage vs. current.
plot(current_filt(:,2),lambda,'Color',red)
line([-80 80],[0 0],'Color','k','LineStyle','--');
line([0 0],[-0.015 0.015],'Color','k','LineStyle','--');
xlim([-80 80])
ylim([-0.015 0.015])
xlabel('Current(A)')
ylabel('Flux Linkage (Vs)')
title('Flux Linkage vs. Current at 1 kHz, 138.2 A_{pp}')

end
```

Bibliography

[1] J. G. Kassakian, M. F. Schlecht, and G. C. Verghese, *Principles of power electronics*, vol. 46. Addison-Wesley Reading, MA, 1991.

[2] J. Acero, R. Alonso, J. Burdio, L. Barragan, and D. Puyal, "Analytical equivalent impedance for a planar circular induction heating system," *IEEE Transactions on Magnetics*, vol. 42, pp. 84–86, Jan. 2006.

[3] F. Forest, E. Laboure, F. Costa, and G. J. Yves., "Principle of a multi-load/single converter system for low power induction heating," *IEEE Transactions on Power Electronics*, vol. 15, no. 2, pp. 223–230, 2000.

[4] O. Lucia, J. Acero, C. Carretero, and J. Burdio, "Induction Heating Appliances," *IEEE Industrial Electronics Magazine*, vol. 1, pp. 35–47, Sept. 2013.

[5] C. P. Steinmetz, *Theory and calculation of electric circuits*. McGraw-Hill book Company, Incorporated, 1917.

[6] R. Groessinger, N. Mehboob, M. Kriegisch, a. Bachmaier, and R. Pippan, "Frequency Dependence of the Coercivity of Soft Magnetic Materials," *IEEE Transactions on Magnetics*, vol. 48, pp. 1473–1476, Apr. 2012.

[7] J. Acero, J. Burdio, L. Barragan, D. Navarro, R. Alonso, J. Garcia, F. Monterde, P. Hernandez, S. Llorente, and I. Garde, "An overview of recent research," *IEEE Industry Applications Magazine*, vol. 16, no. 2, pp. 39–47, 2010.

[8] S. Llorente, F. Monterde, J. M. Burdío, and J. Acero, "A Comparative Study of Resonant Inverter Topologies Used in Induction Cookers," in *Applied Power Electronics Conference and Exposition*, vol. 2, (Dallas, TX.), pp. 1168–1174, IEEE, 2002.

[9] D. Puyal, C. Bernal, J. M. Burdío, J. Acero, and I. Millán, "Versatile High-Frequency Inverter Module for Large-Signal Inductive Loads Characterization Up to 1 . 5 MHz and 7 kW," *IEEE Transactions on Power Electronics*, vol. 23, no. 1, pp. 75–87, 2008.

[10] G.-J. Su, "Multilevel DC-Link Inverter," *IEEE Transactions on Industry Applications*, vol. 41, pp. 848–854, May 2005.

[11] M. D. Bellar, T.-S. Wu, A. Tchamdjou, J. Mahdavi, and M. Ehsani, "A review of soft-switched DC-AC converters," *IEEE Transactions on Industry Applications*, vol. 34, no. 4, pp. 847–860, 1998.

[12] J.-s. Lai and F. Z. Peng, "Multilevel Converters-A New Breed of Power Converters," *IEEE Transactions on Industry Applications*, vol. 32, no. 3, pp. 509–517, 1996.

[13] J. Rodríguez and Lai, Jih-sheng, "Multilevel Inverters : A Survey of Topologies, Controls, and Applications," *IEEE Transactions on Industrial Electronics*, vol. 49, no. 4, pp. 724–738, 2002.

[14] L. Franquelo, J. Rodríguez, J. I. Leon, S. Kouro, R. Portillo, and M. A. M. Prats, "The Age of Multilevel Converters Arrives," *IEEE Industrial Electronics Magazine*, pp. 28–39, June 2008.

[15] C. Newton and M. Sumner, "Novel technique for maintaining balanced internal," *IEEE Proceedings*, vol. 146, no. 3, 1999.

[16] M. Malinowski, K. Gopakumar, J. Rodriguez, and M. Perez, "A survey on cascaded multilevel inverters," *IEEE Transactions on Industrial Electronics*, vol. 57, no. 7, pp. 2197–2206, 2010.

[17] A. K. Panda and Y. Suresh, "Research on cascade multilevel inverter with single DC source by using three-phase transformers," *International Journal of Electrical Power & Energy Systems*, vol. 40, pp. 9–20, Sept. 2012.

[18] M. Escalante and J. Arellano, "Harmonics and Reactive Power Compensation Using a Cascaded H-Bridge Multilevel Inverter," in *2006 IEEE International Symposium on Industrial Electronics*, (Montreal, Quebec, Canada), pp. 1966–1971, Ieee, July 2006.

[19] N. Mohan, T. M. Undeland, and W. P. Robbins, *Power Electronics: Converters, Applications, and Design.* New York: Wiley, 2nd ed., 1995.

[20] C. Blake and C. Bull, "IGBT or MOSFET : Choose Wisely," tech. rep., International Rectifier.

[21] "Application Note AN-9052," tech. rep., Fairchild Semiconductor Corporation, 2008.

[22] V. Barkhordarian, "Power MOSFET Basics," tech. rep., International Rectifier, El Segundo, CA.

[23] P. Viriya, S. Sittichok, and K. Matsuse, "Analysis of high-frequency induction cooker with variable frequency power control," in *Power Conversion Conference*, vol. 3, (Osaka), pp. 1502–1507, IEEE, 2002.

[24] P. W. Hammond, "A New Approach to Enhance Power Quality for Medium Voltage AC Drives," *IEEE Transactions on Industry Applications*, vol. 33, no. 1, pp. 202–208, 1997.

[25] G. Chang, C. Hatziadoniu, W. Xu, P. Ribeiro, R. Burch, W. M. Grady, M. Halpin, Y. Liu, S. Ranade, D. Ruthman, N. Watson, T. Ortmeyer, J. Wikston, A. Medina, A. Testa, R. Gardinier, V. Dinavahi, F. Acram, and P. Lehn, "Modeling devices with nonlinear voltage-current characteristics for harmonic studies," *IEEE Transactions on Power Delivery*, vol. 19, no. 4, pp. 1802–1811, 2004.

[26] I. Daut, S. Hasan, S. Taib, R. Chan, and M. Irwanto, "Harmonic content as the indicator of transformer core saturation," in *International Power Engineering and Optimization Conference*, no. June, (Selangor, Malaysia), pp. 23–24, IEEE, 2010.