Contents

1 Introduction: The World Needs Green Electronics **7**
- 1.1 The Energy Crisis . . . 7
- 1.2 Green Electronics: Part of the Solution . . . 8
  - 1.2.1 Example 1: Photovoltaic Generation . . . 8
  - 1.2.2 Example 2: An Electric Car . . . 10
- 1.3 The World Needs More Green Engineers . . . 11
  - 1.3.1 Organization of this Book . . . 12
- 1.4 Bibliographic Notes . . . 12

2 The Buck Converter **15**
- 2.1 Converter Efficiency . . . 15
- 2.2 Buck Topology . . . 17

3 Analysis of the Buck Converter **21**
- 3.1 Single-Cycle Analysis . . . 21
- 3.2 Low-Frequency Analysis . . . 23
- 3.3 Exercises . . . 24

4 Simulation of the Buck Converter **25**
- 4.1 High-Frequency Model . . . 25
- 4.2 Low-Frequency Model . . . 29
- 4.3 Exercises . . . 29

5 Controlling the Buck Converter **31**
- 5.1 Controller Derivation . . . 31
- 5.2 Controller Simulation . . . 34
- 5.3 Exercises . . . 36

6 Boost and Buck-Boost Converters **39**
- 6.1 Boost Converter . . . 39
- 6.2 Buck-Boost Converter . . . 41
- 6.3 Exercises . . . 43

7 Analysis and Control of the Boost Converter 45
7.1 Analysis of the Boost Converter . . . 45
7.2 Layered Control of the Boost . . . 47
7.3 Simulation . . . 49
7.4 Exercises . . . 51

8 Small-Signal Analysis of the Boost Converter 53
8.1 Impulse Response . . . 53
8.2 Small-Signal Model . . . 55
8.3 Controller Design . . . 56
8.4 Exercises . . . 58

9 Isolated Converters 61
9.1 The Ideal Transformer . . . 61
9.2 A Real Transformer . . . 65
9.2.1 Magnetizing Inductance . . . 66
9.2.2 Leakage Inductance . . . 67
9.3 A Full-Bridge Converter . . . 67
9.4 Effect of Leakage Inductance . . . 70
9.5 Final Comments . . . 72
9.6 Exercises . . . 73

10 The Flyback Converter 75
10.1 Flyback Topology and Operation . . . 75
10.2 Effect of Leakage Inductance . . . 78
10.3 Diagonal Flyback . . . 80
10.4 Primary-Side Sensing . . . 82
10.5 Exercises . . . 82

11 The Forward Converter 83
11.1 Forward Topology and Operation . . . 83
11.2 Comparison of Converters . . . 86
11.3 Exercises . . . 87

12 Discontinuous Conduction Mode 89
12.1 The Buck Converter Operating in DCM . . . 90
12.2 Simulation . . . 91
12.3 Control . . . 91
12.4 Exercises . . . 95

13 Capacitive Converters 97
13.1 A Capacitive Voltage Doubler . . . 97
13.2 Analysis of the Capacitive Voltage Doubler . . . 98
13.3 Exercises . . . 100

14 The Grid **101**
14.1 Structure of the Grid . . . 101
14.2 Generation . . . 102
14.3 Transmission and Distribution . . . 102
14.4 Control of the Grid . . . 103
14.5 Future of the Grid . . . 104
14.6 Exercises . . . 104

15 Inverters **105**
15.1 Basic Inverter . . . 105
15.2 Inverter Simulation . . . 106
15.3 Exercises . . . 110

16 Power Factor Correction **111**
16.1 Power Factor . . . 111
16.2 Power Factor Correction . . . 114
16.3 Simulation of a PFC Input Stage . . . 117
16.4 Grid-Tied Inverters . . . 120
16.5 Three-Phase Systems . . . 120
16.6 Exercises . . . 120

17 Switches **123**
17.1 DC Characteristics of Real Switches . . . 123
17.2 AC Characteristics of FETs and IGBTs . . . 125

18 Metal-Oxide Field Effect Transistors (MOSFETs) **129**
18.1 Simple Model . . . 129
18.2 Structure . . . 130
18.3 Dynamics . . . 133
18.4 Parasitic Elements . . . 133
18.5 Avalanche Energy . . . 136
18.6 Typical MOSFETs . . . 136
18.7 Exercises . . . 138

19 Diodes **139**
19.1 Simple Model . . . 140
19.2 Key Parameters . . . 142
19.3 Reverse Recovery . . . 144
19.4 Forward Recovery . . . 149
19.5 Current Hogging . . . 151
19.6 SPICE Models . . . 155
19.7 Measurements . . . 156
19.8 Exercises . . . 156

20 Switching Losses **159**
20.1 The Half Bridge . . . 159
20.2 Turn-On Transient . . . 161
20.3 Turn-Off Transient . . . 164
20.4 Parasitic Losses . . . 167

21 Anatomy of a Switching Cycle **169**
21.1 An IGBT Half-Bridge . . . 169
21.2 A Switching Cycle . . . 172
21.3 High-Side Turn-On . . . 172
21.4 High-Side Turn Off . . . 175
21.5 Exercises . . . 175

22 Gate Drive **177**
22.1 Gate Drive Circuit . . . 177
22.2 Dead Time . . . 179
22.3 High-Side Supplies . . . 183
22.4 Transformer Gate Drive . . . 186
22.5 Transient Immunity . . . 186
22.6 Source Reference . . . 187
22.7 Protection . . . 188
22.8 Exercises . . . 188

23 Magnetic Components **189**
23.1 Magnetic Circuits . . . 189
23.2 Equivalent Magnetic Circuits . . . 192
23.3 Inductance . . . 195
23.4 Transformers . . . 196
23.5 Exercises . . . 199

24 Magnetic Materials **201**
24.1 Magnetic Materials . . . 201
24.1.1 Saturation . . . 201
24.1.2 Hysteresis . . . 202
24.2 Exercises . . . 206

25 Wire for Magnetics **207**
25.1 Skin Effect . . . 207
25.2 Proximity Effect . . . 208
25.3 Litz Wire . . . 208
25.4 Winding Geometry . . . 208
25.5 Exercises . . . 209

26 Designing Magnetic Components **211**
26.1 Inductor Design . . . 211
26.2 Inductor Example . . . 213
26.3 Core Selection . . . 217
26.4 Details . . . 217
26.5 Exercises . . . 218

27 Transformer Design **219**
27.1 Design Procedure . . . 219
27.2 Transformer Design Example . . . 220
27.3 Flyback Transformers . . . 221
27.4 Exercises . . . 221

28 Photovoltaic Cells and Modules **223**
28.1 Photovoltaic Cells . . . 223
28.2 Photovoltaic Modules . . . 223
28.3 Single-Phase PV Systems . . . 223
28.4 Three-Phase PV Systems . . . 223
28.5 Module Imbalance . . . 223
28.6 PV Economics . . . 223
28.7 Exercises . . . 223

29 Electric Motors **229**
29.1 Lorentz Force and Faraday Induction . . . 229
29.2 The Brushed Permanent Magnet Motor . . . 231
29.3 Motor Model . . . 232
29.4 The Torque Curve . . . 233
29.5 Step Response . . . 235
29.6 Motor Drive . . . 237
29.7 Exercises . . . 239

30 Motor Control **241**
30.1 A Naive Controller . . . 241
30.2 PI Motor Control . . . 243
30.3 Current-Limited Control . . . 243
30.4 Exercises . . . 246

31 Brushless Permanent Magnet Motors **247**
31.1 Motor Geometry . . . 247
31.1.1 Back EMF . . . 249
31.1.2 Torque . . . 249
31.2 Equivalent Circuit . . . 251
31.3 Brushless Motor Drive . . . 251
31.4 Hard Electronic Commutation . . . 253
31.5 Excercises . . . 253

32 AC Induction Motors **255**
32.1 Motor Geometry 255
32.2 Induction Motor Operation 255
32.3 Equivalent Circuit 259
32.4 Torque Curve 260
32.5 Induction Motor Control 261
32.6 Exercises 268

33 Soft Switching **273**
33.1 The Quasi-Square-Wave Converter 273
33.2 Control of the Quasi-Square-Wave Converter 279
33.3 Exercises 282

Chapter 1

Introduction: The World Needs Green Electronics

1.1 The Energy Crisis

An increasing, industrialized world population is putting undue pressures on our planet's delicate ecosystem and on our natural resources. In the United States, 85% of energy generated is from fossil fuels (DOE Annual Energy Outlook). While there are large reserves, that are increasing as people get better at exploration and recovery, the supply is finite. We will ultimately run out of fossil fuels.

Even if our supply of fossil fuels was infinite, burning fossil fuels generates greenhouse gasses like CO_2 that accumulate in the atmosphere leading to climate change. In recent years, atmospheric CO_2 has been increasing at a rate of about 20ppm/decade (5% per decade) (NOAA data). We need to burn less fossil fuels, or sequester the carbon generated by burning these fuels, or both — or suffer the consequences of climate change.

To deal both with the limited supplies of fossil fuels and their environmental impact, we need to develop technologies for economic renewable energy sources. Wind and solar power are two of the most promising candidates — both depend strongly on green electronics. Photovoltaic controllers, generator controllers, and inverters are important components in efficiently getting energy from solar cells and windmills to the grid — and ultimately to the end user.

Conservation plays a large role in reducing our consumption of fossil fuels and our generation of greenhouse gasses. Replacing inefficient incandescent lamps with highly-efficient LED lamps and using intelligent control to have lamps only illuminate areas where people are looking can dramatically reduce the energy consumption due to lighting. Data centers, which consume about 2% of the electric power in the US, use power supplies that are often only 70-80% efficient and processors that spend the bulk of their power on overhead. More efficient power supply and processor design would dramatically reduce

electricity consumption. There are many more examples where efficient power electronics and intelligent control can greatly reduce consumption.

1.2 Green Electronics: Part of the Solution

A key technology for both sustainable energy generation and conservation is the combination of intelligent control with power electronics — the brains and brawn of energy systems. The power electronics provides the brawn — doing the heavy work of converting energy from one form to another. This includes electrical conversion — as in a photovoltaic system where the DC output voltage from a string of solar panels is converted to AC power at a different voltage to drive the grid — and mechanical conversion where the system includes a motor or generator — as in a wind farm or an electric vehicle.

Intelligent computer control provides the brains of the system, optimizing operation to improve efficiency. A photovoltaic controller searches to find the optimum operating point for each solar cell in a string — optimizing efficiency. A controller electrically switches the windings of a brushless generator on a wind mill, achieving greater efficiency than with an alternator or a generator with a commutator. A lighting controller uses cameras to sense the location and gaze of people in a building to turn on lights only where they are needed.

1.2.1 Example 1: Photovoltaic Generation

As an example of a *green-electronics system*, consider the grid-connected photovoltaic system of Figure 1.1. A *photovoltaic (PV) array* of solar panels converts solar energy to electricity. With a typical configuration of 10 40V panels in each series *string*, the output of this array is 400V DC. The power depends on the number of strings and for a residential system is typically in the range of 4-10kW. So the DC current out of the PV array is in the range of 10-25A.

A *PV Controller* and *Inverter* connects the PV array to a bank of batteries (that operate at 48V DC) and to the electric grid (240V AC). The PV controller converts electrical energy from one form to another. For a 4kW system, it converts 10A at 400V DC to 16.7A at 240V AC to drive the grid, or to 83.3A at 48V DC to charge the batteries.

The controller also sets the operating point of the PV array to maximize the energy produced. As shown in Figure 1.2, a solar panel acts as a *current source* — producing a nearly constant current up until a maximum voltage where the current falls quickly. To get maximum power out of the panel, it must be operated at the *knee* of this curve — at the voltage where the current just begins to drop. The PV controller adjusts its output impedance to operate the array at this point.

Because of non-uniformity of illumination and manufacturing not all panels may have the same opimal operating point. In a simple system like that shown in Figure 1.1 the controller must set a single operating point that is a compromise — it gives the highest power over all but may not be the best for each panel.

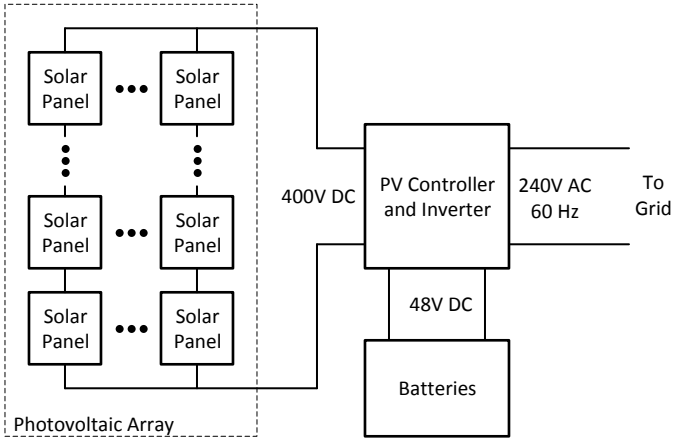

Figure 1.1: A photovoltaic system. The PV controller converts between 400V DC (string of panels), 48V DC (batteries), and 240V AC power and sets the operating point for the PV array to maximize efficiency.

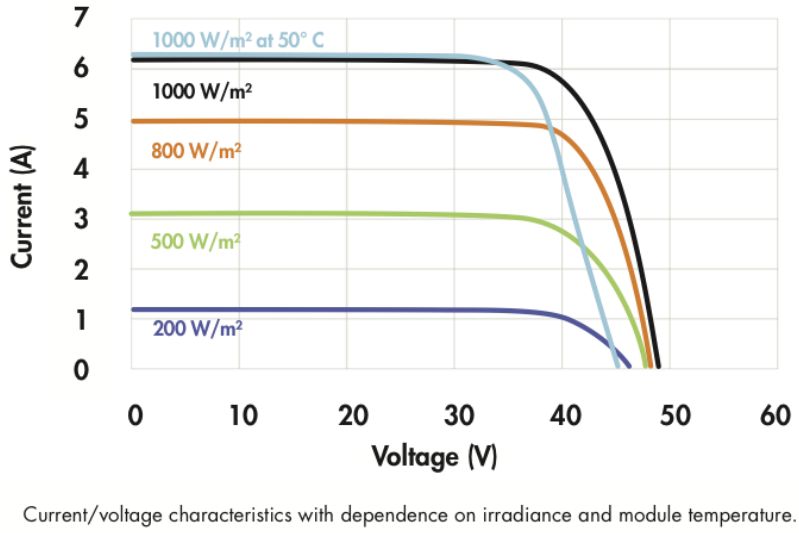

Figure 1.2: Current/voltage (IV) curve of a SunPower E19/240 solar panel. Optimum power is produced when the panel is operated at the *knee* of the curve.

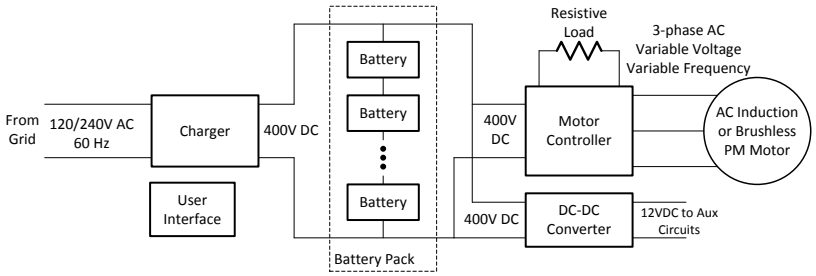

Figure 1.3: Block diagram of an electric car. Energy is stored in a central battery pack that can be charged from the grid when the car is parked. A controller can transfer energy in either direction between the battery pack and an electric motor.

More sophisticated PV controllers monitor the voltage and current across each panel independently and by bypassing current around some panels (using a switched-inductor shunt) operate every panel at its peak power point.

In a grid-connected system that includes battery storage, an intelligent controller decides when to charge the batteries to make best use of *time-of-day* power rates. It will charge the batteries during off-peak times when electricity is substantially cheaper than during peak times[1]. Not only does this optimize economics, it also optimizes efficiency because the *peaking plants* used by the electric utility to provide power at peak times are substantially less efficient than the plants used to provide base capacity.

The controller will also tailor the charging profile to maximize efficiency and battery life. A sophisticated controller will monitor the voltage of individual battery cells and balance the cells to optimize operation.

1.2.2 Example 2: An Electric Car

A second example of a green electronics system is the electric car of Figure 1.3. The electronics here charge and manage the batteries and control the motor that provides vehicle propulsion.

When the car is parked, a *charger* converts 120V or 240V AC power to 400V DC power to charge the battery pack. While individual battery cells range from 2V (lead-acid) to 3.7V (Lithium ion), hundreds of these cells are connected in series to give an efficient operating voltage of 400V DC. The charger must carefully manage voltage, current, and temperature during the charging process to maximize efficiency and battery life.

The motor controller converts the 400V DC voltage to a three-phase AC

[1]Under the California E-7 Tariff, the peak rate is nearly four times the off-peak rate.

voltage to drive the motor. The frequency of the motor drive depends on the motor speed. The effective voltage of the motor drive is varied using *pulse-width modulation* to control the motor current which determines the torque produced by the motor.

The motor controller employs a *control law* to decide how much current, and hence torque, to apply at any given time. The decision depends on throttle position, vehicle speed, and battery state. The decision made is a compromise between drivability, energy efficiency, and battery longevity.

To recycle energy when slowing the car, the motor controller employs *regenerative braking* by operating the motor as a generator. In this case, the kinetic energy of the vehicle is converted to electrical energy that is used to charge the battery. During this process, the motor controller acts as a generator controller — sequencing the motor phases to optimize generator efficiency and controlling the battery charging process. If regenerative braking occurs when the batteries are fully charged the excess energy is disspated as thermal energy in a resistive load.

In all modes of operation — whether accelerating or braking — the battery pack must be carefully managed to balance cell voltages, and avoid excessive currents and temperatures. Lithium Ion cells are particularly sensitive to proper battery management.

In a hybrid vehicle where a heat engine can be turned on to provide propulsion and charge batteries the controller has an additional degree of freedom. The controller must decide when to turn on the heat engine, how to divide the torque load between the heat engine and the electric motor, and how much energy to divert to charging batteries when the heat engine is running. Making optimal power management decisions may involve predicting the future route of the car — for example, whether it is likely to arrive at a charging station soon, or whether it is likely to be doing up or down a large hill.

A user interface is provided to that the driver can understand how their control actions result in energy movement within the system. A good user interface can train the driver to drive more efficiently. Of course one must be careful with the user interface to avoid distracting the driver[2]

1.3 The World Needs More Green Engineers

To realize *green electronic* systems like efficient photovoltaic controllers, wind mill controllers, energy converters, electric- and hybrid-vehicle drives, etc... requires engineers who understand power electronics and computer control. There is a shortage of such talented individuals.

The goal of this course is to teach the basics of these two critical disciplines as an introductory engineering course. The course teaches *engineering thinking* using green electronic examples. In doing so, it conveys engineering fundamentals — learning the process of specification, design, and analysis and learning

[2]In the long run the computer will drive the car, resulting in more efficient operation.

techniques like modeling and simulation. At the same time, it gives a working knowledge of simple power electronics and microprocessor-based control.

1.3.1 Organization of this Book

This book is organized into four parts. Chapters 2 to 16 describes the toopology, analysis, and control of voltage converters. As evidenced by our photovoltaic and electric car examples above, such converters are central to all Green Electronics systems. These converters are realized from semiconductor switches (MOSFETs, IGBTs, and diodes) and magnetic components. The properties and design of these components and their uses are discussed in Chapters **??** to 26. The motors that are used to convert between electrical and mechanical energy are covered in Chapters 29 to **??**. The control systems that are used to control both voltage converters and motors are described in Chapters ?? to ??.

1.4 Bibliographic Notes

DOE Annual Energy Outlook NOAA data

Part I. Voltage Converters

Buck converter
Buck Analysis
Buck Simulation
Buck Control
Boost and Buck-Boost Converters
Control of the Boost Converter
Small-Signal Analysis of the Boost Converter
Isolated converters
Flyback
Discontinuous mode Capacitive converters
Inverters
Power Factor and PFC

Chapter 2

The Buck Converter

Efficient voltage conversion is at the core of most green-electronic systems. Converting power from a photovoltaic array to drive the AC line, controlling a motor, and providing an efficient power supply to a computer are at their core problems of efficient voltage conversion. In this chapter we introduce the *buck* converter which is one of the simplest and most widely used voltage converters. A simple buck regulator, for example, is used in most computer power supplies to convert an intermediate 12V DC voltage to a 1V DC supply voltage to power the processor.

2.1 Converter Efficiency

As shown in Figure 2.1, a voltage converter takes an input supply with voltage V_1 and supplies a load across its output with a different voltage, V_2. The converter shown in Figure 2.1 is a *non-isolated* converter because the source and load share a common terminal. In Chapter 9 we will consider the design of *isolated* converters in which the input and output voltages need not be referenced to a common point.

If the load in Figure 2.1 draws current I_2, then the power delivered to the load is $P_2 = I_2V_2$. If the input current is I_1, then the input power is $P_1 = I_1V_1$. With no energy storage or generation in our converter we must have $P_1 >= P_2$. We define the loss as

$$P_L = P_1 - P_2. \tag{2.1}$$

and the efficiency of the converter is given by

$$\nu = \frac{P_2}{P_1}. \tag{2.2}$$

A simple way to convert voltage levels is to drop the input voltage across a resistor as shown in Figure 2.2(a). Resistor R must be set so that $R = \frac{V_1 - V_2}{I_2}$

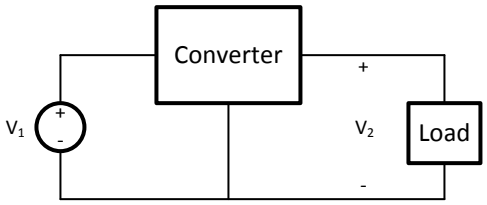

Figure 2.1: A voltage converter. A voltage converter accepts input voltage V_1 from a voltage source and converts the source power to output voltage V_2 that is supplied to a load.

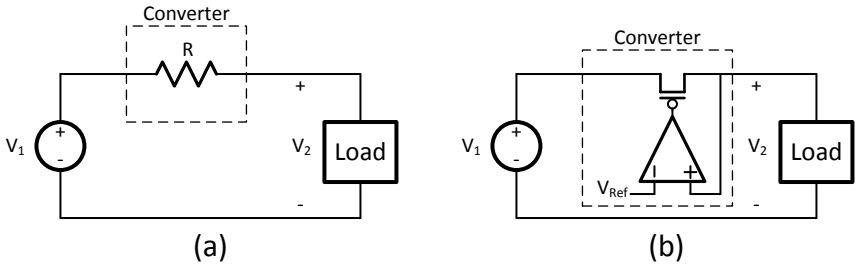

Figure 2.2: A dissipative voltage converter. (a) The input voltage is reduced by dropping the difference voltage across a resistor. (b) In practice the resistor is realized by a PFET and a feedback loop adjusts the resistance of the PFET to give the desired output voltage.

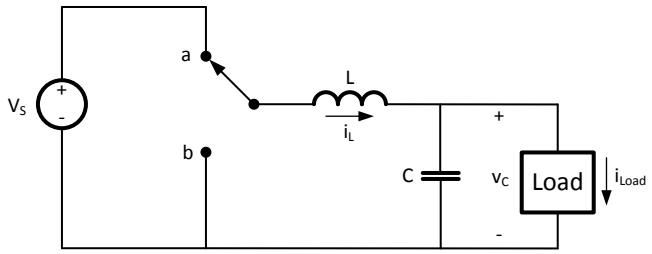

Figure 2.3: A buck converter. Input voltage V_S is converted to output voltage v_C without loss by storing energy in inductor L.

to give the correct output voltage V_2. Any time the load current changes the resistance R must be changed to maintain the proper voltage.

To simplify the continuous adjustment of resistance, the resistor is often implemented with an active device controlled by a feedback circuit to control its resistance. For example, Figure 2.2(b) shows a converter implemented using a PFET driven by an op-amp. Negative feedback adjusts the gate voltage of the PFET so that output voltage V_2 matches a reference voltage V_{Ref}.

The converters shown in Figure 2.2 are *dissipative* converters. They convert the input voltage to the output voltage by dissipating the voltage difference in a lossy element like a resistor. These converters are sometimes called *linear regulators* because they are built with linear elements. Such a converter, however is **not** linear. Its input-output relationship is not linear.

With such a *dissipative regulator* we have $I_1 = I_2$ which means that our efficiency, neglecting overhead losses, is

$$\nu_{\text{Dis}} = \frac{V_2}{V_1} \tag{2.3}$$

For example, if our computer power supply with a 12V input and 1V output used a dissipative regulator it would have an efficiency of no better than 1/12 or 8.3%. Clearly this is unacceptably inefficienct.

2.2 Buck Topology

To convert voltages more efficiently we use an energy storage device rather than a dissipative device in our converter. A *buck* converter is shown in Figure 2.3. The buck converter stores energy in an inductor to convert from a higher input voltage to a lower output voltage.

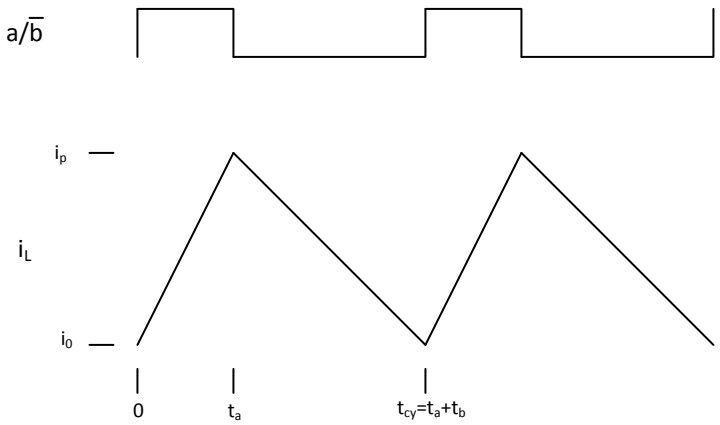

Figure 2.4: Waveforms showing operation of the buck converter in *periodic steady state*. The top trace shows the position of the switch. The bottom trace shows inductor current i_L.

Waveforms showing operation of the buck converter are shown in Figure 2.4. For this analysis we assume capacitor C is large large enough that the voltage across the capacitor is constant. Let the current in the inductor at the beginning of the cycle (at time $t = 0$) be i_0. When the switch is in position **a**, the voltage across the inductor is $V_L = V_S - v_C$ and the current increases with a slope of $\frac{di_L}{dt} = \frac{V_S - v_C}{L}$. When the switch changes position at time t_a the current has reached a peak value of

$$i_p = i_0 + \frac{t_a \left(V_S - v_C\right)}{L}. \tag{2.4}$$

After time t_a the switch is in position **b** and the voltage across the inductor is $V_L = -v_C$. During this period the current in the inductor decreases with slope $\frac{di_L}{dt} = \frac{-v_C}{L}$. At the end of the cycle (time $t = t_{cy}$), the current has returned to its initial value i_0.

When a circuit starts and ends a cycle with its state variables having the same values, as is shown in Figure 2.4, it is said to be in a *periodic steady state*. Its state variables change within a cycle, but from cycle-to-cycle it is in a steady state. At the same point in each cycle, its state variables have the same values.

For a buck converter to be in the periodic steady state, two conditions must hold. First, for the capacitor voltage to remain stable, the average inductor current must equal the load current:

$$|i_L| = \frac{i_0 + i_p}{2} = I_{\text{Load}}. \tag{2.5}$$

Second, for the inductor current to return to i_0 at the end of the cycle we must have the volt-seconds applied to the inductor sum to zero over the cycle:

$$t_a\left(V_S - v_C\right) - \left(t_{cy} - t_a\right)\left(v_C\right) = 0. \tag{2.6}$$

Which we can rewrite as

$$\begin{aligned} t_a V_S &= t_{cy} v_C & (2.7)\\ \frac{t_a}{t_{cy}} &= \frac{v_C}{V_S} \\ D &= \frac{v_C}{V_S} & (2.8) \end{aligned}$$

From (2.8) we see that in the periodic steady state the voltage ratio of the converter v_C/V_S is equal to the duty factor of the switch $D = t_a/t_{cy}$. Analysis of the converter when it is not in the periodic steady state will be covered in Chapter 3.

The principle of superposition is useful in analyzing switching power circuits. A simple way to directy write (2.7) is to consider each voltage source in turn. If we zero v_C and consider only V_S we see that a voltage of V_S is applied across the inductor for duration t_a. Then we zero V_S and consider only v_C and see that $-v_C$ is applied across the inductor during the entire cycle t_{cy}. Hence, the total *volt-seconds* seen by the inductor is $V_S t_a - v_C t_{cy}$.

The buck converter can also be understood by considering it to be a pulse width modulator followed by a low-pass filter. The switch produces a voltage V_x at the left terminal of the inductor with a duty factor $D = t_a/t_{cy}$ and an amplitude of V_S. Thus, the average voltage at this point is $|V_x| = DV_S$. The inductor and capacitor act as a low-pass filter to produce this average voltage at the output so that $v_C = |V_x| = DV_S$.

If the switch and the inductor in Figure 2.3 are ideal the converter is lossless: $P_2 = P_1$ and $\nu_{\text{Buck}} = 1$. Unfortunately real switches have finite switching times leading to *switching losses*. Also, real switches and inductors have non-zero resistance leading to *conduction losses*. We will consider ideal elements for now and revisit real losses in Chapter 17.

Chapter 3

Analysis of the Buck Converter

Suppose the buck converter is not in periodic steady state. How do its state variables i_L and v_C evolve over time as a function of duty factor D and load current i_{Load}? We can most easily address this question by considering two different time scales. First we consider how the state variables evolve over a single switching cycle. Then, using this relationship to determine the change in state variables over each cycle, we analyze the dynamics of the circuit at frequencies lower than the switching frequency.

3.1 Single-Cycle Analysis

Figure 3.1 shows the inductor current for a buck converter that is not in periodic steady state. The converter starts cycle 0 with current i_0 and ramps up to current i_{p0} as given by (2.4). However, at the end of the cycle, the current is left at $i_1 \neq i_0$, the starting current for cycle 1.

We can compute the difference current during cycle i Δi_i by considering the imbalance in volt-seconds applied to the inductor each cycle

$$\begin{aligned} \Delta i_i &= \frac{V_S t_a - v_i t_{cy}}{L} \\ &= \frac{t_{cy}\left(DV_S - v_i\right)}{L}. \end{aligned} \tag{3.1}$$

Here V_S is the source voltage applied to the input of the circuit and v_i is the voltage across the capacitor at the start of cycle i. We continue to assume it is approximately constant across the cycle to simplify our analysis.

The starting current for each cycle is simply this difference current added to the starting current from the previous cycle:

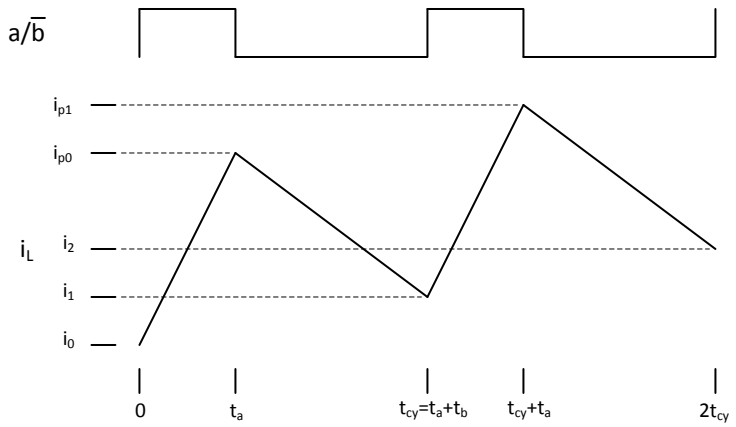

Figure 3.1: Waveforms showing operation of the buck converter when it is not in periodic steady state. The top trace shows the position of the switch. The bottom trace shows inductor current i_L.

$$i_{i+1} = i_i + \Delta i_i. \tag{3.2}$$

The peak current each cycle is given by adding the inital ramp to the starting current for that cycle:

$$i_{pi} = i_i + \frac{t_a\left(V_S - v_i\right)}{L}. \tag{3.3}$$

The average current over cycle i is given by

$$\hat{i}_i = \left(\frac{t_a}{2t_{cy}}\right) i_i + \left(\frac{1}{2}\right) i_{pi} + \left(\frac{t_{cy} - t_a}{2t_{cy}}\right) i_{i+1}. \tag{3.4}$$

$$\hat{i}_i = \frac{1}{2}\left(D i_i + i_{pi} + (1-D) i_{i+1}\right). \tag{3.5}$$

To simplify analysis we will make the approximation

$$\hat{i}_i \approx \frac{1}{2}\left(i_i + i_{pi}\right). \tag{3.6}$$

$$= i_i + \frac{t_a\left(V_S - v_i\right)}{2L}. \tag{3.7}$$

From this average current we can calculate the change in capacitor voltage each cycle

$$\Delta v_i = \frac{t_{cy}\left(\hat{i}_i - i_{\text{Load}}\right)}{C} \tag{3.8}$$

This analysis has assumed that the voltage across the capacitor is constant, even though it is changing each cycle. For large capacitors, and hence small voltage changes, this approximation is valid. The full analysis, considering the effects of capacitor volage changes within each cycle, is considerably more complex. The current waveforms are not piecewise linear but rather segments of sine waves with frequency $\omega = (LC)^{-1/2}$.

The increasing current in a buck converter not in steady state shown in Figure 3.1 highlights the importance of a controller that regulates the operation of the converter. Without a controller the current can quickly ramp to destructive levels.

3.2 Low-Frequency Analysis

At frequencies lower than the switching frequency we are not concerned with the detailed waveforms within each cycle, but rather only with the evolution of state variables from cycle to cycle. Equations (3.1) and (3.8) describe the evolution of i_L and v_C as a discrete-time system. We approximate these difference equations with differential equations to give a continuous time system to simplify analysis.

$$\frac{di}{dt} = \frac{DV_S - v_i}{L} \tag{3.9}$$

$$\frac{dv}{dt} = \frac{\hat{i}_i - i_{\text{Load}}}{C}. \tag{3.10}$$

Taking the Laplace transform of (3.9) and (3.10) gives:

$$I(s) = \frac{D(s)V_S - V(s)}{Ls} \tag{3.11}$$

$$V(s) = \frac{I(s) - I_{\text{Load}}(s)}{Cs} \tag{3.12}$$

Substituting (**??**) into (**??**) gives:

$$\begin{aligned} I(s) &= \frac{D(s)V_S}{Ls} - \frac{I(s) - I_{\text{Load}}(s)}{LCs^2} \\ LCs^2 I(s) &= CV_S D(s)s - I(s) + I_{\text{Load}}(s) \\ I(s)(1 + LCs^2) &= CV_s D(s)s + I_{\text{Load}}(s) \\ I(s) &= \frac{D(s)CV_s s + I_{\text{Load}}(s)}{1 + LCs^2} \end{aligned} \tag{3.13}$$

And substituting (**??**) into (**??**) gives:

$$\begin{aligned} V(s) &= \frac{D(s)V_S - V(s)}{LCs^2} - \frac{I_L(s)}{Cs} \\ LCs^2V(s) &= D(s)V_s - V(s) - LsI_L(s) \\ (1 + LCs^2)V(s) &= D(s)V_s - LsI_L(s) \\ V(s) &= \frac{D(s)V_s - LsI_L(s)}{1 + LCs^2} \end{aligned} \tag{3.14}$$

If we are concerned with the transfer function from $D(s)$ to $V(s)$ we can (for now) ignore I_{Load} and write:

$$\frac{V(s)}{D(s)} = \frac{V_s}{1 + LCs^2} \tag{3.15}$$

Equations (3.13) through (3.15) show that the response of the buck converter to a transient in D, V_S, or i_{Load} is an undamped oscillation with frequency $\omega = (LC)^{-1/2}$. The denominator of $I(s)$ has a damping factor of $\zeta = 0$. A dissipative element, such as a resistor in series with the source, is needed to damp these oscillations, or a controller is needed to prevent them.

Even if the oscillations are damped, an overshoot of more than a few percent can lead to destruction of the load. Consider our computer power supply. When the converter turns on, a transient in D from 0 to 0.083, the output voltage will oscillate between 0 and 2V before eventually settling to 1V. The first excursion above 1.2V is likely to do permanent damage to the load (the computer). Any transient in load current will excite similar oscillations. A controller is required to prevent destructive overshoot.

In a real buck converter there are some additional components that affect the dynamics of the system. In particular, the parasitic series resistance (ESR) of the output capacitor introduces a zero in the transfer function of (3.15). This zero must be taken into account when designing controllers for the buck. Also, most real loads are not pure current sources but have a resistive component. This resistance adds a damping factor. We examine the effects of these additional components in the exercises.

3.3 Exercises

relax constant capacitor voltage approximation
add series resistor to voltage source
add resistive load
add ESR to capacitor
response to change in input voltage

Chapter 4

Simulation of the Buck Converter

We can understand the dynamics of our buck converter by building a model that simulates these dynamics by evolving the state variables over small time steps. Using the simulation model we can see the *open-loop* response of the system to an input stimulus. We can also implement a *control law* and examine the closed-loop response of the system with a controller.

4.1 High-Frequency Model

To simulate our system we divide time into small steps of magnitude `dt` and compute the change to our state variables over each time step. We need to choose the time step small enough to make our simulation stable and to avoid inaccuracies due to discretization of time. In practice dividing each switching cycle into 100 time steps provides sufficient accuracy and gives reasonable simulation run times.

Figure 4.1 shows Matlab code that computes the evolution of the state variables i_L and v_C for one time step. The `if` statement computes the change in inductor current `di_l` depending on the position of the switch `a`. This is just applying the constituent equation of the inductor

$$i_L = \int \frac{v_L}{L} dt.$$

Which we can rewrite as

$$di_L = \frac{v_L}{L} dt.$$

The next line computes the change in voltage over the time step applying the equation

```
% simulate one time step
if(a==1) % if switch is in position a
  di_l = (v_s-v_c)*dt/l ;
else % switch is in position b
  di_l = -v_c*dt/l ;
end
dv_c = (i_l - i_load)*dt/c ;
i_l = i_l + di_l ; % integrate inductor current
v_c = v_c + dv_c ; % integrate capacitor voltage
```

Figure 4.1: Matlab simulation model to advance the state of the buck converter by one timestep, dt.

```
asteps = steps*(t_a/t_c) ;
for i2=1:steps
  a = (i2 <= asteps) ;
  step ;
end
```

Figure 4.2: Matlab code to simulate one switching cycle t_c of the buck converter. Each call to `step` invokes the code shown in Figure 4.1.

$$dv_C = \frac{i_C}{C} dt.$$

The remaining two lines of code perform the integration by adding the change in the state variables to their state at the last time step to compute their state at the new time step. It is important that we compute all of the changes to the state variables before updating any of the state variables so that all of the changes are computed using the state variables from the same (previous) time step.

Figure 4.2 shows Matlab code to simulate the buck converter for one switching cycle t_c. The variable `steps` is the number of steps per switching cycle, in this case 100. The code first computes the number of steps for which the switch should be in position `a`, `asteps`. It then uses this number to set variable `a` before calling `step`, the code from Figure 4.1.

Figure 4.3 shows the results of simulating three switching cycles of buck converter operation using the code of Figures 4.2 and 4.1 with a load current of $i_{\text{Load}} = 10\text{A}$ and initial conditions $v_{C0} = .9995\text{V}$ and $i_{L0} = 9.6\text{A}$. The top trace shows the inductor current ramping up with slope $(V_S - v_C)/L$ during the first part of each cycle and ramping down with slope $-v_C/L$ during the second part of each cycle. The current curve is actually a portion of a sine wave - since the voltage is changing during the cycle - but is very close to a straight line.

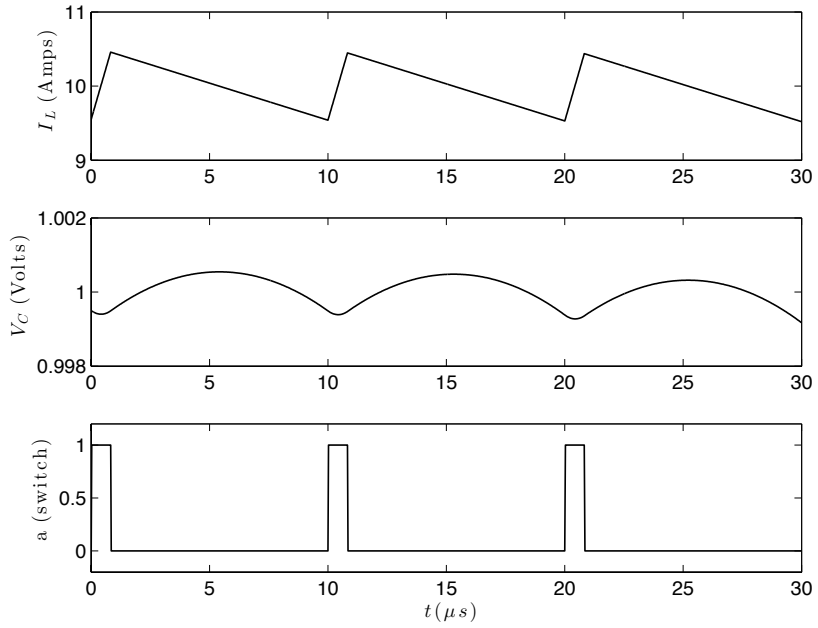

Figure 4.3: Waveforms showing simulated operation of the buck converter over three cycles while in the periodic steady state.

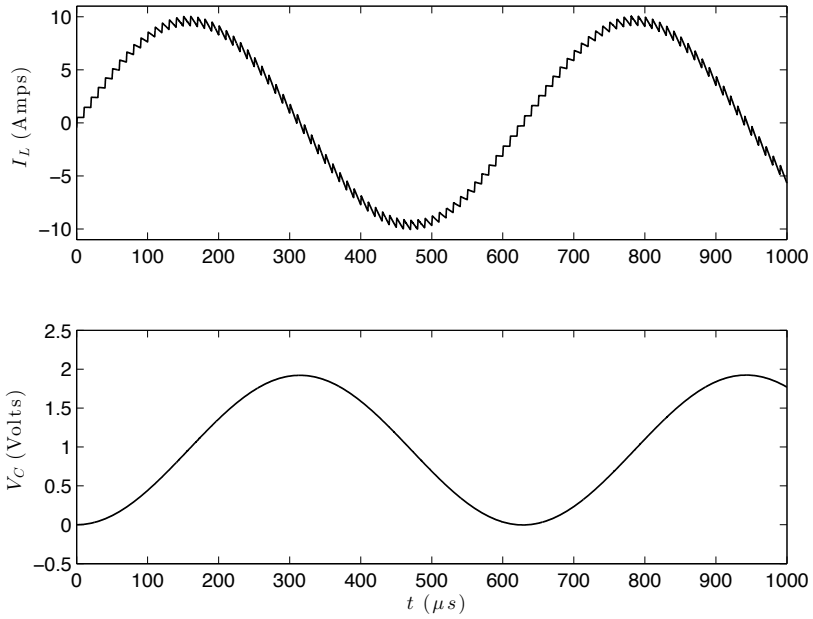

Figure 4.4: Waveforms showing simulated operation of the buck converter over 1ms during a *turn-on* transient with zero load current. The output voltage overshoots by a factor of two and oscillates with a period of 628 μs.

The second trace shows a very small ripple on v_C. Which varies by ± 0.5mV about the target voltage of $v_{\text{ref}} = 1$V. Again the actual curve is part of a sine wave - since its generated by a tank circuit. However, it is closely approximated by a parabola because it is the integral of the current which is closely approximated by a line.

The bottom trace of the figure shows the position of the switch. The duty factor is set to 1/12 to give the steady state condition for the 12:1 step down ratio of our example converter.

Figure 4.4 shows a simulation of the buck converter for 1ms of simulated time. The simulation shows a *startup transient* with initial conditions of $i_{l0} = v_{c0} = 0$. This simulation uses our high-frequency model to simulate the low-frequency behavior of the converter.

As expected from Equation (**??**), without a controller the converter oscillates in response to the transient with frequency $\omega = (LC)^{-1/2}$ and hence period $T = 2\pi\sqrt{LC}$. In this case $T = 628\mu$s.

If an actual converter oscillated in this manner it would most likely destroy the downstream circuit. A circuit designed for a 1V supply typically cannot withstand a 100% overshoot to 2V.

```
% simulate one time step - low frequency model
di_l = (d*v_s-v_c)*dt/l ;
dv_c = (i_l - i_load)*dt/c ;
i_l = i_l + di_l ;
v_c = v_c + dv_c ;
```

Figure 4.5: Matlab simulation model to advance the state of the buck converter by one timestep, *dt*, using a low-frequency approximation.

4.2 Low-Frequency Model

Figure 4.4 shows that we can use our high-frequency model to simulate the low-frequency behavior of the converter. However, it is inefficient to do so. We can reduce our simulation time by an order of magnitude or more by using a model that directly reflects the low-frequency behvavior of the converter as captured in Equations (3.9) and (3.10) directly.

Figure 4.5 shows a low-frequency simulation model for the buck converter. The first line computes the change in inductor current averaged over the cycle. We are not concerned with the position of the switch, but rather abstract the switching into the duty factor variable `d`. Otherwise the model is identical to that of Figure 4.1. With this small change, however, we can get by with a much larger step size `dt` resulting in fewer steps per switching cycle and a much faster simulation.

Figure 4.6 shows waveforms generated using the Matlab model of Figure 4.5 with initial conditions of $v_C = i_L = 0$ and 20 steps per switching cycle (`dt`= 0.5μs. The step size here must be kept small enough to avoid excessive errors in our *rectangular* integration formula. Using a better integration formula would enable a smaller step size.

Except for the missing ripple, which is not modeled by the low-frequency model, the response is identical to that produced by our high-frequency model.

4.3 Exercises

simulate buck with different initial conditions

better integration formula

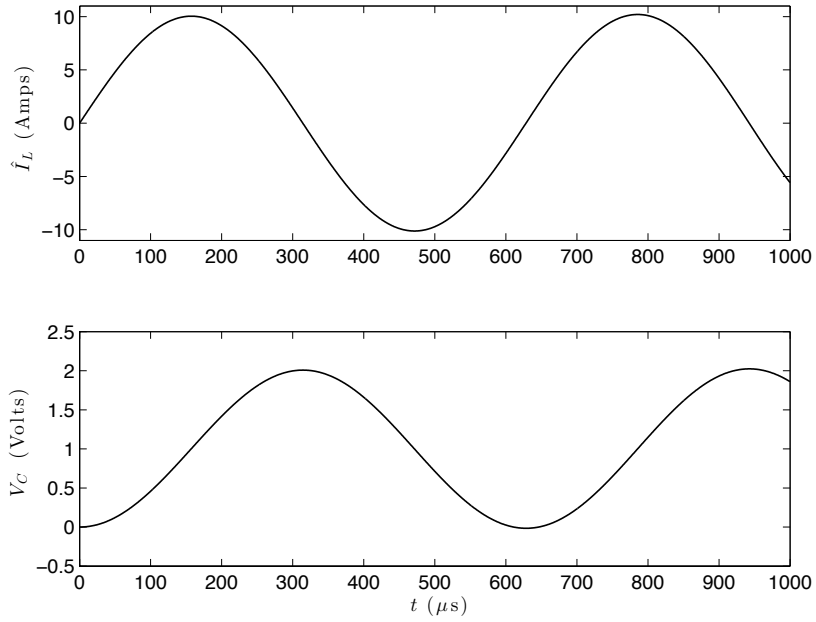

Figure 4.6: Waveforms showing simulated operation of the buck converter using a low-frequency model.

Chapter 5

Controlling the Buck Converter

In Chapters 3 and 4 we saw through both analysis and simulation that the buck converter, by itself, is unstable. A transient in either duty factor or load leads to an undamped oscillation in both output voltage v_C and inductor current i_L. In many cases such oscillations would destroy the load.

A good controller modifies the dynamics of the system to make it stable, avoid overshoot, and give fast response to transients. There are many strategies for controlling a buck converter. In this chapter we illustrate the principles of control using a simple PD (proportional-derivative) feedback controller.

5.1 Controller Derivation

Figure 5.1 shows the block diagram of a PD feedback controller for the buck regulator. The output voltage $v(s)$ is subtracted from a target reference voltage $V_{\text{Ref}}(s)$ to generate an error signal $e(s)$. A PD controller applies a proportional

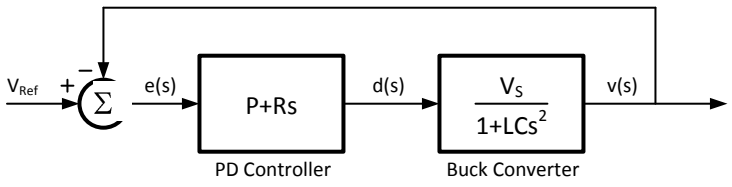

Figure 5.1: Block diagram of a feedback controller.

gain P and a derivative gain Rs to the error signal to generate a control signal $d(s) = (P + Rs)e(s)$. Here the duty factor is the control input to the buck converter.

The *open-loop* transfer function of this system is

$$\begin{aligned} H(s) &= \frac{v(s)}{e(s)} \\ &= \frac{V_S(P + Rs)}{1 + LCs^2} \\ &= \frac{(KP + KRs)}{s^2 + \frac{1}{LC}} \end{aligned} \tag{5.1}$$

Where $K = V_S/LC$. Closing the feedback loop gives a *closed-loop* transfer function of

$$\begin{aligned} G(s) &= \frac{v(s)}{V_{\mathrm{Ref}}(s)} \\ &= \frac{H(s)}{1 + H(s)} \\ &= \frac{\frac{(KP+KRs)}{s^2+\frac{1}{LC}}}{1 + \frac{(KP+KRs)}{s^2+\frac{1}{LC}}} \\ &= \frac{KP + KRs}{s^2 + KRs + KP + \frac{1}{LC}} \end{aligned} \tag{5.2}$$

The denominator of equation (5.2) determines the frequency ω and damping factor ζ of our closed-loop system. Its natural frequency is the square root of the constant term

$$\omega = \sqrt{KP + \frac{1}{LC}} \tag{5.3}$$

Its damping factor is given by the linear term

$$\zeta = \frac{KR}{2\omega} \tag{5.4}$$

We can rewrite (5.3) to give an expression for P

$$P = \frac{\omega^2 - \frac{1}{LC}}{K} = \frac{LC\omega^2 - 1}{KLC} \tag{5.5}$$

and we can rearrange (5.4) to give an expression for R:

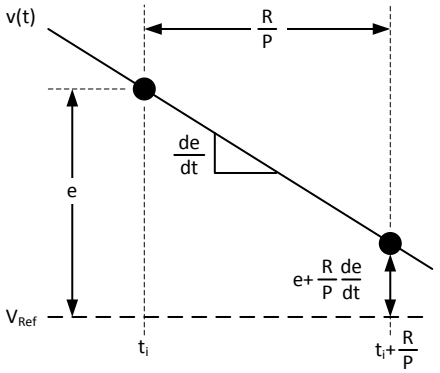

Figure 5.2: PD control is equivalent to looking ahead and providing proportional feedback at the value projected $\frac{R}{P}$ in the future.

$$R = \frac{2\zeta\omega}{K} \tag{5.6}$$

If we choose $\omega = 2.2 \times 10^4$ and $\zeta = 1$ (critically damped) we get $P = 1/3$ and $R = 3.7 \times 10^{-5}$. We need to choose ω to keep $f = \omega/2\pi$ well below (at least a factor of 10 below) the switching frequency to preserve our continuous time approximation. It should also be chosen to minimize non-linear behavior by avoiding saturating the control variable. Too high a gain will lead to duty factors being commanded outside the legal range.

The intuition behind proportional feedback, the P term, is simple. If we have an error in a certain direction, we apply an input proportional to the error. If our output voltage is below the reference voltage we apply a positive input to drive it higher. The larger the error, the larger the input.

Unfortunately, proportional feedback is not sufficient to control a second-order system like the buck converter. The inductor current gives the system *momentum*. Once the output voltage is rising, the inductor current will cause it to continue to rise - even if we set the duty factor to zero. To control such a system we need to compute a control input based not only on the error but also on the derivative of the error.

Derivative feedback has the effect of compensating for the *momentum* in the system by looking ahead and providing feedback based on the projected future value of the system output. This is illustrated in Figure 5.2. The figure shows a system where $e(t)$ has a positive value e and a negative slope de/dt at time t_i. A PD controller effectively *looks ahead* to time $t_i + R/P$ assuming a constant

```
% control law - runs once per cycle
er = v_ref - v_c ;
der = (er - old_er)/t_c ;
old_er = er ;
d = er*P + der*R + d0;
d = max(0,min(d_max,d)) ;
```

Figure 5.3: Matlab code for PD controller for buck converter.

slope and applies a control input based on that predicted value. By setting the derivative gain appropriately the look ahead interval can be set to critically damp the response - so that the error smoothly goes to zero without overshoot.

5.2 Controller Simulation

We can verify our controller, and the validity of the assumptions on which it is based, through simulation. Figure 5.3 shows Matlab code implementing our PD controller. This code runs once each switching cycle. It computes the error `er`, approximates the derivative of the error using forward differencing to calculate `der`. It then applies the control law to compute the duty factor `d`.

We add the duty factor needed to maintain the correct steady state output voltage `d0` = 1/12 to the control law to give the correct output voltage when the error is zero. This is an example of *feed-forward control*. The input to the plant includes a term directly proportional to the target voltage in addition to the output of the PD controller.

Finally, we limit the pulse width to its legal range. This limit is important because many PWM generators will *wrap* out of limit values. This can result in a small negative value of `d` resulting in a large pulse width — leading to catastrophic effects. The limit causes our controller to *saturate* rather than *wrap* when the control variable goes out of range.

In an optimized controller we would avoid the divide by folding $1/t_c$ into the constant multiplied by `der`. With this optimization the control law is four adds, two multiplies (by constants), and a limit operation. Even a slow microcontroller can perform the 7 operations needed to implement this loop fast enough to keep up with a 100kHz converter.

Figure 5.4 shows the response of the buck converter with PD controller to the startup transient. The conditions are identical to those shown in Figure 4.4 except for the addition of the controller. We see that the addition of feedback control, with appropriately selected parameters P and R, eliminates the oscillatory behavior. The output voltage v_C (middle panel) quicly converges to the target value with no overshoot.

The bottom panel shows the pulse width t_a and illustrates the non-linear nature of this controller. After four cycles, the inductor current is pumped up to

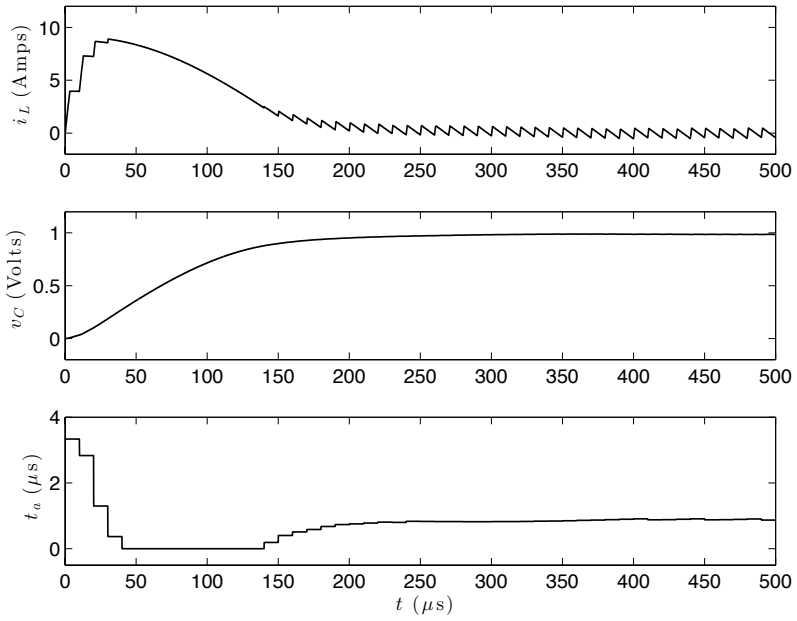

Figure 5.4: Response of the buck converter with a PD controller to the start up transient.

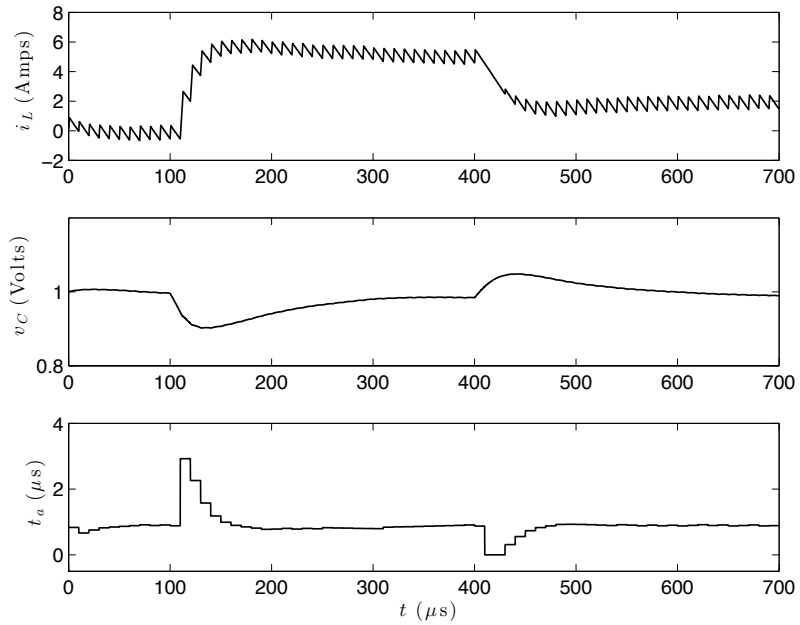

Figure 5.5: Response of the buck converter with a PD controller to a 5A current step at 100μs and a -3A current step at 400μs.

about 9A and the slope of v_c becomes so steep that the PD control law computes a negative duty factor. The last line of Figure 5.3 clamps this value to zero. The pulse width remains at zero for 11 cycles. Fortunately this non-linear behavior does not result in any overshoot. This is not always the case. Linear control theory does not completely predict the behavior of control systems where the control variables saturate.

Figure 5.5 shows the response of the controller to load current transients. Load currently steps from 0 to 5A at 100μs and then down to 2A at 400 μs. In each case the transient causes an initial disturbance in v_c. The controller then acts to drive the system back to $v_c = 1$ with no overshoot. The controller goes sligtly non-linear on the negative transient when a negative duty factor is computed and clamped to zero.

5.3 Exercises

Compute P and R for various values of L, C, V_S

What if V_S changes?

What is upper limit on P

Add integral control term - deal with windup

Explore different choices for ω for what values is saturation during startup transient avoided.

Add ESR to output capacitance.

Chapter 6

Boost and Buck-Boost Converters

Now that we understand the operation and control of the buck converter, we will explore its two close relatives: the *boost* converter, and the *buck-boost* converter. All three of these converters, the buck, boost, and buck-boost, are called *non-isolated* converters because their input and output share one terminal. We will visit *isolated* converters, where the input and output share no terminals and can *float* with respect to one another in Chapter 9.

These three types of converters are used to provide output voltages in different ranges. A buck converter is used to convert an input voltage V_S to a lower voltage $v_C \leq V_S$ of the same polarity. A boost converter, which is just a buck converter connected backwards, does the opposite. It *boosts* an input voltage V_S to a higher output voltage $v_C \geq V_S$ of the same polarity. A buck-boost converter is used to convert an input voltage V_S to an output voltage v_C of opposite polarity.

All three of the basic non-isolated converters (buck, boost, and buck-boost) use the same basic *cell* consisting of an inductor with one end connected to a two position switch. All three converters also have three terminals: input, output, and ground. The type of converter is determined by how the cell is connected to the terminals, specifically by which terminal is connected to the free end of the inductor. There are three possibilities: When the free end of the inductor is connected to the output we have a buck converter. When the free end of the inductor is connected to the input we have a boost converter (a buck connected backwards). Finally, when the free end of the inductor is connected to ground we have a buck-boost converter.

6.1 Boost Converter

Figure 6.1 shows a *boost* converter which is used to boost input voltage V_S to a higher output voltage $v_C \geq V_S$. A boost converter is just a buck converter with

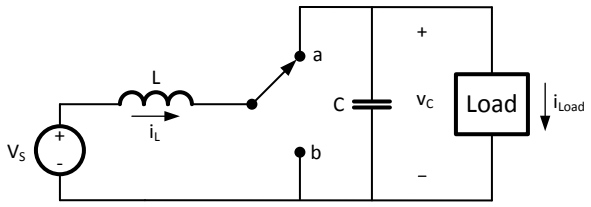

Figure 6.1: A boost converter. Input voltage V_S is *boosted* to output voltage $v_c \geq V_S$. A boost converter is simply a buck converter connected backwards — i.e., with input and output swapped.

its input and output swapped. Comparing Figure 6.1 with Figure 2.3 source V_S is swapped with the capacitor C (and the load) and the polarity of inductor current i_L is reversed. If a buck or boost is realized using switches that can carry current in either direction then power can flow in either direction and the two circuits are identical.

Boost converters are widely used any time an output voltage higher than the input voltage is required. They are used, for example, in power-factor correction input stages (Chapter 16) where they boost the rectified AC input voltage, which varies from 0V to 170V, to a higher intermediate DC voltage, often 200-400V, while keeping the input current proportional to the AC input voltage.

Figure 6.2 illustrates operation of the boost converter over two switching cycles in a periodic steady state. By convention we consider the cycle to start with the switch down, in position **b**. During the time t_b that the switch is down, the right end of the inductor is tied to ground, input voltage V_S is applied across the inductor, and the inductor current ramps with slope $\frac{di_L}{dt} = \frac{V_S}{L}$. Starting from an initial value of i_0, at time t_b the current has reached

$$i_p = i_0 + \frac{t_b V_S}{L}. \tag{6.1}$$

During the second half of the cycle the switch is up, in position **a**, and the right side of the inductor is connected to v_C. The voltage across the inductor is $V_S - v_C$, which has a negative value in the periodic steady state since $v_C > V_S$. During this part of the cycle, the inductor current decreases with slope $\frac{di_L}{dt} = \frac{V_S - v_C}{L}$.

For the converter to be in the periodic steady state we must have the volt-seconds across the inductor sum to zero over the cycle:

$$t_b V_S + t_a \left(V_S - v_C\right) = 0. \tag{6.2}$$

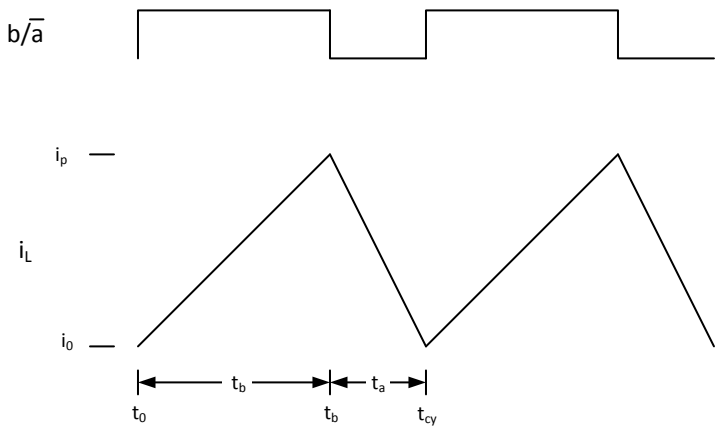

Figure 6.2: Waveforms showing operation of the boost converter.

If we denote the duty factor of the switch in position **a** as $D = \frac{t_a}{t_a+t_b} = \frac{t_a}{t_{cy}}$, we can rewrite this as:

$$\begin{aligned}(1-D)V_S + D\,(V_S - v_C) &= 0.\\ V_S - Dv_C &= 0\\ Dv_C &= V_S\\ v_C &= \frac{V_S}{D}.\end{aligned} \tag{6.3}$$

Which is just a rederivation of (2.8) with v_C and V_S reversed, $V_S = Dv_C$.

Many books describe the boost converter in terms of the duty factor of the switch in position **b**. If we define $D_b = \frac{t_b}{t_{cy}} = 1 - D$ we can rewrite (6.3) as

$$v_C = \frac{V_S}{1-D_b}. \tag{6.4}$$

Which is the more common form of (6.3) but one which makes it harder to see the relationship to the buck converter.

6.2 Buck-Boost Converter

Figure 6.3 shows a buck-boost converter which is used to produce a negative output voltage v_C from a positive input voltage V_S. When the switch is in

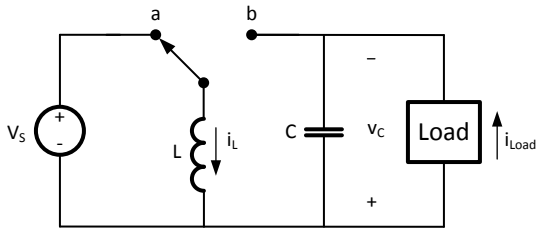

Figure 6.3: A buck-boost converter accepts input voltage V_S and produces output voltage v_C of the opposite polarity.

position **a** input voltage V_S is across the inductor and current increases with slope $\frac{di_L}{dt} = \frac{V_S}{L}$ reaching peak current:

$$i_P = i_0 + \frac{t_a V_S}{L}. \tag{6.5}$$

When the switch is in position **b** the capacitor is connected across the inductor and current ramps downward with slope $\frac{di_L}{dt} = \frac{-v_C}{L}$. We have labeled Figure 6.3 with the positive end of the capacitor downward so that v_C is a positive number but $-v_C$ appears on the output — and across the inductor when the switch is in position **b**.

To be in periodic steady state, the volt-seconds across the inductor must sum to zero giving:

$$\begin{aligned} t_a V_S &= -t_b v_C \\ D V_S &= -(1-D) v_C \\ v_C &= -\left(\frac{D}{1-D}\right) v_S. \end{aligned} \tag{6.6}$$

We see from (6.6) that the buck-boost converter negates the input voltage, generating an output voltage $V_C = -V_S$ when duty factor is $D = 0.5$. With a higher duty factor (in position **a**) a higher (negative) output voltage is generated. With a smaller duty factor a lower (negative) output voltage is generated. For example, when $D = 0.25$, $V_C = -\frac{1}{3}V_S$ and when $D = 0.75$, $V_C = -3V_S$.

Some people have trouble with the schematic of Figure 6.3 because the capacitor is *upside-down* with its positive terminal on the bottom. Figure 6.4 shows the buck-boost converter redrawn to eliminate this objection. The capacitor is *right-side-up* with the positive terminal on top, and the load current i_{Load} flows downward. This is achieved at the expense of crossing signal lines.

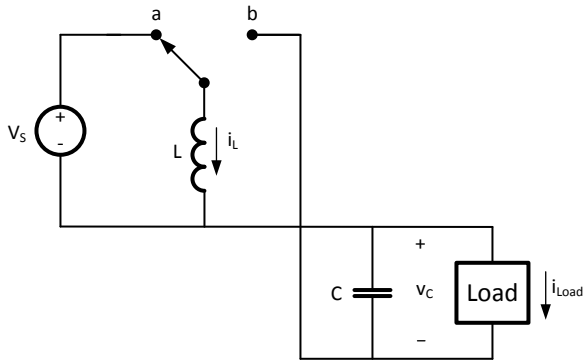

Figure 6.4: The buck-boost converter of Figure 6.3 redrawn so that the output capacitor has its positive terminal on top.

The circuit of Figure 6.4 is identical to that of Figure 6.3 however, some people find this version easier to understand.

6.3 Exercises

Cuk

Chapter 7

Analysis and Control of the Boost Converter

The boost converter, while it is just a buck converter connected backwards, has significantly different dynamics than the buck converter. A decrease in D_a of the boost gives a small immediate decrease and a long-term increase in the output voltage V_C. This short term decrease in V_C manifests itself as a right-half-plane zero in the transfer function of the boost which makes the boost significantly harder to control than the buck — which lacks this zero.

In this chapter we develop a dynamic model of the boost converter. We see that this model is non-linear and hence does not admit the easy control solution we applied to the buck. We use this non-linearity as an opportunity to introduce *layered control* and develop a layered controller for the boost converter. In Chapter 8 we solve the same problem in a more conventional manner by building a linearized small signal model of the boost converter and designing a conventional PD controller based on this model.

7.1 Analysis of the Boost Converter

When the boost converter is not in periodic steady state the change in inductor current over one cycle is given by dividing the volt-seconds imbalance by L:

$$\begin{aligned}\Delta i_i &= \frac{t_{cy}\left((1-D)V_S + D\left(V_S - v_i\right)\right)}{L} \\ \Delta i_i &= \frac{t_{cy}\left(V_S - Dv_i\right)}{L}\end{aligned} \tag{7.1}$$

The change in capacitor voltage over one cycle is given by the current imbalance:

$$\Delta v_i = \frac{t_{\text{cy}}\left(D\hat{i}_i - i_{\text{Load}}\right)}{C}. \tag{7.2}$$

The inductor only supplies current to the output capacitor when the switch is in position **a**. Thus $\hat{i}$ is muliplied by D in (7.2) to give the average inductor current seen by the output capacitor.

If we convert difference equations (7.1) and (7.2) to differential equations and take their Laplace transforms we get

$$i(s) = \frac{V_S - D(s)v(s)}{Ls} \tag{7.3}$$

$$v(s) = \frac{D(s)i(s) - i_{\text{Load}}(s)}{Cs} \tag{7.4}$$

Combining gives

$$\begin{aligned} i(s) &= \frac{V_S}{Ls} - \frac{D^2(s)i(s) - D(s)i_{\text{Load}}(s)}{LCs^2} \\ LCi(s)s^2 &= CV_s s - D^2(s)i(s) + D(s)i_{\text{Load}}(s) \\ LCi(s)s^2 + D^2(s)i(s) &= CV_s s + D(s)i_{\text{Load}}(s) \\ i(s) &= \frac{CV_s s + D(s)i_{\text{Load}}(s)}{LCs^2 + D^2(s)} \end{aligned} \tag{7.5}$$

Substituting into (7.4) gives

$$\begin{aligned} v(s) &= \frac{\frac{CV_s D(s)s + D^2(s)i_{\text{Load}}(s)}{LCs^2 + D^2(s)}}{Cs} \\ v(s) &= \frac{CV_s D(s)s^2 + D^2(s)i_{\text{Load}}(s)s}{LC^2s^2 + CD^2(s)} \end{aligned} \tag{7.6}$$

Ignoring the contribution from I_{Load} this simplifies to

$$v(s) = \frac{V_s D(s)s^2}{LCs^2 + D^2(s)} \tag{7.7}$$

From (7.7) we see that even though a boost is just a buck connected backwards it has significantly different dynamics. We cannot easily write a transfer function $v(s)/d(s)$ for the boost. The response from d to v is non-linear - there is a d^2 term in the (7.7). Also, the response from d to v acts in both directions as indicated by d appearing in both the numerator and denominator of (7.7). In the short term increasing d increases v because the increased duty factor provides a larger fraction of the inductor current to the capacitor. However, over time a larger d decreases i and hence decreases v.

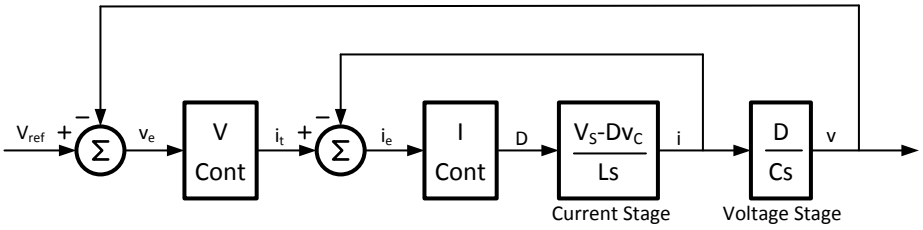

Figure 7.1: A current-mode controller for the boost converter. An inner control loop sets duty factor d to control inductor current i. The output control loop sets the target inductor current i_t to control voltage v.

7.2 Layered Control of the Boost

To simplify control of the boost we employ a *layered* or *nested* control system as illustrated in Figure 7.1. We use an inner control loop to control the inductor current i. An outer control loop controls capacitor voltage v using inductor current as the control variable. This approach splits a difficult second-order control problem into two simpler first-order control problems.

The downside of this approach is that it requires that we provide an inductor current input to our controller. Accurately measuring the inductor current with low noise will add cost to our converter.

We use proportional-integral (PI) controllers for both the current and voltage loops. The current loop computes a duty factor d to drive i_L to the target current i_t. The voltage loop computes a target current i_t to drive v_C to v_{ref}.

The PI controller for the current loop computes the current error $i_e = i_t - i_L$ and then computes the duty factor d as a weighted sum of the error and the integral of the error:

$$d = Pi_e + Q \int i_e dt. \tag{7.8}$$

Taking the laplace transform we have

$$\frac{d(s)}{i_e(s)} = P + \frac{Q}{s} \tag{7.9}$$

Multiplying this transfer function by (7.3) we get the open-loop transfer function of our current stage with controller:

$$H(s) = \frac{i(s)}{i_e(s)} = \frac{v_C\,(P + Q/s)}{Ls} \tag{7.10}$$
$$= \frac{v_C\,(Ps + Q)}{Ls^2}$$

Closing the loop gives:

$$G(s) = \frac{v_C\,(Ps + Q)}{Ls^2 + v_C Ps + v_C Q} \tag{7.11}$$
$$= \frac{\frac{v_C P}{L}s + \frac{v_C Q}{L}}{s^2 + \frac{v_C P}{L}s + \frac{v_C Q}{L}}$$

From the denominator of (7.12) we calculate:

$$\omega = \sqrt{\frac{v_C Q}{L}} \tag{7.12}$$
$$\zeta = \frac{v_C P}{2L\omega} \tag{7.13}$$
$$Q = \frac{\omega^2 L}{v_C} \tag{7.14}$$
$$P = \frac{2L\zeta\omega}{v_C} \tag{7.15}$$

For the current loop we choose $\omega = 6.28 \times 10^4$ (f=10kHz) and $\zeta = 1$. For our example converter with $v_C = 250$V and $L = 500\mu$H we derive $P = 0.251$ and $Q = 7.88 \times 10^3$.

The PI controller for the voltage loop computes the target current for the current mode controller as:

$$i_t = P' v_e + Q' \int v_e dt \tag{7.16}$$

For the voltage loop, combining the PI controller with (7.4), ignoring the load current, and considering D to be a constant we get the open-loop transfer function:

$$\frac{v(s)}{v_e(s)} = \frac{D\left(P' + \frac{Q'}{s}\right)}{Cs}$$
$$= \frac{D(P's + Q')}{Cs^2} \tag{7.17}$$

This gives a closed-loop transfer function of

```
% control law - execute once per t_c
% PI voltage controller
er = v_ref - v_c ;
ier = ier + er*t_c ;
i_t = er*P + ier*Q ;

% current controller
er2 = i_t - i_l ;
ier2 = ier2 + er2*t_c ;
d2 = er2*P2+ier2*Q2+D ;
t_a = d2*t_c ;
t_a = max(0,min(t_amax,t_a)) ;
```

Figure 7.2: Matlab code implementing the layered controller given by (7.8) and (7.16).

$$H(s) = \frac{\frac{DP'}{C}s + \frac{DQ'}{C}}{s^s + \frac{DP'}{C}s + \frac{DQ'}{C}} \tag{7.18}$$

So we have

$$\omega = \sqrt{DQ/C} \tag{7.19}$$

$$\zeta = \frac{DP}{2C\omega} \tag{7.20}$$

$$Q' = \frac{\omega^2 C}{D} \tag{7.21}$$

$$P' = \frac{2C\zeta\omega}{D} \tag{7.22}$$

For stability we must operate the voltage control loop with a lower natural frequency than the current control loop, so that transients to the current loop have settled out before the voltage loop advances. We choose $\omega = 5 \times 10^3$ ($f = 800$Hz). With $\zeta = 1$ and $C = 10\mu$F we have $P' = 0.25$ and $Q' = 7.88 \times 10^3$.

7.3 Simulation

To validate our layered controller, we simulate it using Matlab. Matlab code for our controller is shown in Figure 7.2. Figure 7.3 shows the response of the system to two steps on V_{ref}. V_{ref} starts at 250 V and is set to 240 V at 500μs and to 260V at 1ms.

The waveform for the target current in the bottom panel of Figure 7.3 shows the effect of controller non-linearity. At 1ms the voltage controller commands a

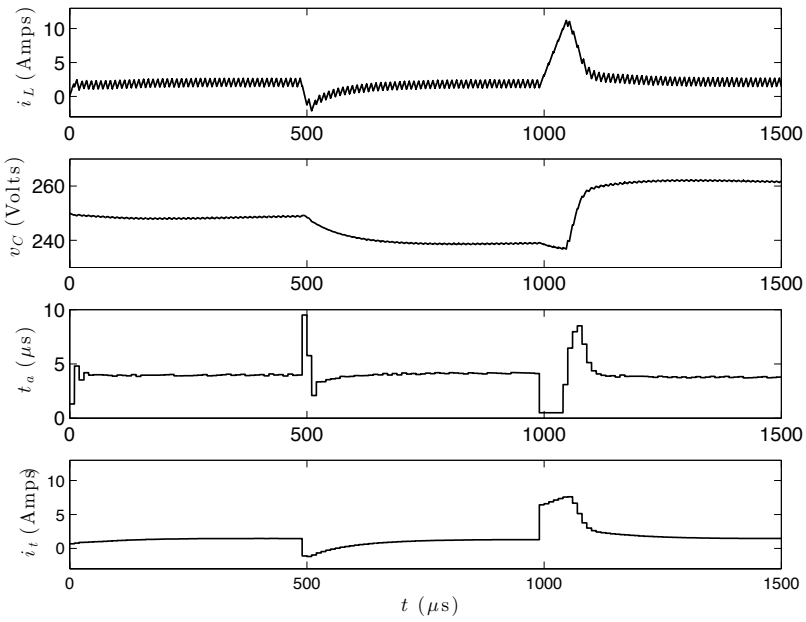

Figure 7.3: Simulated response of the boost controller with layered control.

steep increase in i_t. In an attempt to comply the current loop runs out of range on t_a (third panel of Figure 7.3). The waveform for i_L tracks that of i_t very closely except in cases like this where the control actuator saturates. In some cases control saturation can result in undesired overshoot as the response of the non-linear system deviates from the linear response. In such cases the control parameters should be adjusted to avoid saturation.

7.4 Exercises

controller with voltage loop at 1kHz - saturates variables and overshoots

apply current mode control to a buck

apply current mode control to a buck-boost

add ESR to capacitor

derive equation with load resistance

draw bode plots

Chapter 8

Small-Signal Analysis of the Boost Converter

In Chapter 7 we derived the dynamics of the boost converter. Because the converter was non-linear, we could not apply the direct approach to control that we used on the buck converter in Chapter 5. Instead we used a *layered* controller consisting of two PI control loops: an inner loop controlling inductor current and an outer loop controlling voltage.

In this chapter we take a different approach to controlling the boost regulator. We develop a linear transfer function for the boost converter by making a *small-signal* approximation. For small deviations about its steady-state operating point the boost converter will behave nearly linearly. We derive its linear transfer function for small perturbations about the operating point and in turn derive a control law based on this transfer function. We need to understand, however, that the resulting control law will only be valid in a small region around the original operating point.

8.1 Impulse Response

We can visualize the linearized transfer function of the boost converter by simulating its impulse response. Suppose the converter is operating in the steady state with current $i_L = I$, voltage $v_C = V$, and duty factor $d = D$. From this state we purturb it by applying a duty factor of $d = D + \hat{d}$ for one cycle. The resulting response is the scaled *impulse response* of the converter.

Figure 8.1 shows the simulated impulse response of our example boost converter. The converter is operating in periodic steady state until t=100μs when the duty factor is decreased by 0.1 for one cycle (i.e., $\hat{d} = -0.1$). The result is a waveform on v_C that first dips slightly and then oscillates with a 1.1ms period and an amplitude of 3.6V.

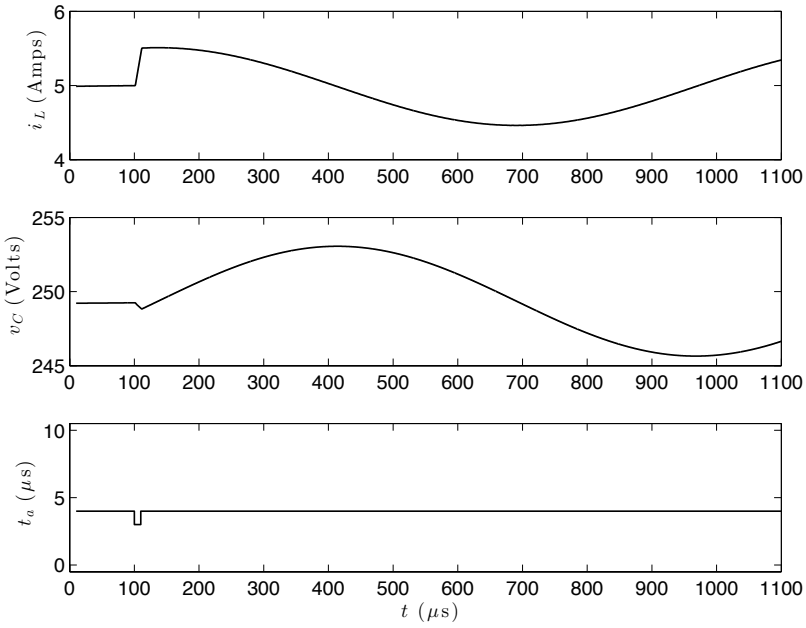

Figure 8.1: Simulated impulse response of the boost converter. The duty factor is increased by 0.1 for one cycle at 100μs.

8.2 Small-Signal Model

We can consider any variable, say v, as the sum of its steady-state component V and its *small signal component* $\tilde{v}$.

$$v = V + \tilde{v} \tag{8.1}$$

Similarly, we can write

$$i = I + \tilde{i} \tag{8.2}$$
$$d = D + \tilde{d} \tag{8.3}$$
$$\tag{8.4}$$

Using this notation, we can *linearize* a non-linear function about its operating point by substituting $X+\tilde{x}$ for each occurance of a variable x, and then dropping certain terms. We discard products of two small signal components e.g., $\tilde{i}\tilde{d}$, because the product of two small things is too small to consider. We also discard strictly steady-state terms, e.g., ID, because we are computing a small-signal component.

Applying the small-signal approximation to (7.1) and (7.2) gives

$$\frac{\Delta\tilde{i}}{t_{cy}} = -\frac{1}{L}\left(\tilde{d}V + D\tilde{v}\right) \tag{8.5}$$
$$\frac{\Delta\tilde{v}}{t_{cy}} = \frac{1}{C}\left(\tilde{d}I + D\tilde{i}\right) \tag{8.6}$$

Approximating the left hand sides as derivatives and taking Laplace transforms gives

$$\tilde{i}(s) = -\frac{1}{Ls}\left(V\tilde{d}(s) + D\tilde{v}(s)\right) \tag{8.7}$$
$$\tilde{v}(s) = \frac{1}{Cs}\left(I\tilde{d}(s) + D\tilde{i}(s)\right) \tag{8.8}$$

Substituting (8.7) into (8.8) gives

$$\begin{aligned}
\tilde{v}(s) &= \left(\frac{I}{Cs} - \frac{DV}{LCs^2}\right)\tilde{d}(s) - \frac{D^2}{LCs^2}\tilde{v}(s) \\
\tilde{v}(s)\left(1 + -\frac{D^2}{LCs^2}\right) &= \left(\frac{I}{Cs} - \frac{DV}{LCs^2}\right)\tilde{d}(s) \\
\tilde{v}(s) &= \frac{\left(\frac{I}{Cs} - \frac{DV}{LCs^2}\right)\tilde{d}(s)}{1 + \frac{D^2}{LCs^2}} \\
\frac{\tilde{v}(s)}{\tilde{d}(s)} &= \frac{ILs - DV}{LCs^2 + D^2}
\end{aligned} \tag{8.9}$$

We can write (8.9) in a form that makes the parameters of the transfer function more explicit as:

$$\frac{\tilde{v}(s)}{\tilde{d}(s)} = K_b \frac{\frac{s}{\omega_1} - 1}{\frac{s^2}{\omega_0^2} + 1} \tag{8.10}$$

where

$$K_b = \frac{V_S}{D^2} \tag{8.11}$$

$$\omega_1 = \frac{V_S}{I_L L} \tag{8.12}$$

$$\omega_0 = \sqrt{\frac{D^2}{LC}} \tag{8.13}$$

Constant $-K_b$ is the incremental DC gain of the boost converter from d to V_C, i.e., $K_b = -\frac{\partial V_C}{\partial d}$. The natural frequency of the converter is ω_0 which is reduced by D^2 from a buck with the same components. The transfer function has a zero in the right half plane with frequency ω_1. This zero corresponds to the dip in the impulse response before the ringing starts. The frequency of the zero changes as current I_L varies and thus the control parameters must be adapted to changing loads.

For our example system with $V_S = 100$V, $V_C = 250$V, $I_L = 2$A, $L = 500\mu$H, and $C = 10\mu$F we have $K_b = 625$, $\omega_0 = 5.6 \times 10^3$ (900 Hz), and $\omega_1 = 10^5$ (16kHz).

A *Bode plot* showing the magnitude and phase of $H(s)$ in (8.10) is shown in Figure 8.2. The gain from DC up to 10^3 rad/s is $K_b = 625$ (56dB) at ω_0 the phase abruptly shifts to 180 degrees and the magnitude starts rolling off at 20dB per decade. At 3×10^4 rad/s the phase starts shifting by another 90 degrees due to the zero at $\omega_1 = 10^5$ rad/s. The unity-gain frequency is 2.2×10^5 rad/s.

8.3 Controller Design

To stabilize the system we will use a PD controller $P + Rs$ to accomplish two modifications to the Bode plot of Figure 8.2. First, we use a small proportional gain $P < 1$ to reduce the DC gain. This reduces the unity-gain frequency of the system — moving it well below the zero at 10^5 rad/s. Second, we add a zero at $\omega_c = P/R$ to advance the phase at the new unity gain point well above 180 degrees.

For stability we would like a phase margin of at least 60° at the unity gain frequency. To achieve this we choose $P = 5 \times 10^{-3}$ to reduce the unity-gain frequency, and $\omega_c = 6 \times 10^3$ rad/s to add phase margin. This gives $R = P/\omega_c = 8.3 \times 10^{-7}$.

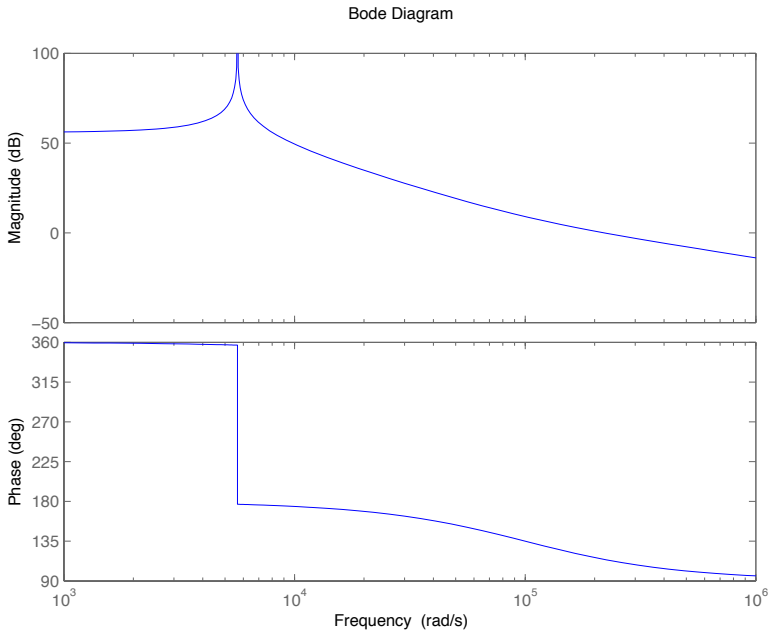

Figure 8.2: Bode plot of transfer function (8.10)

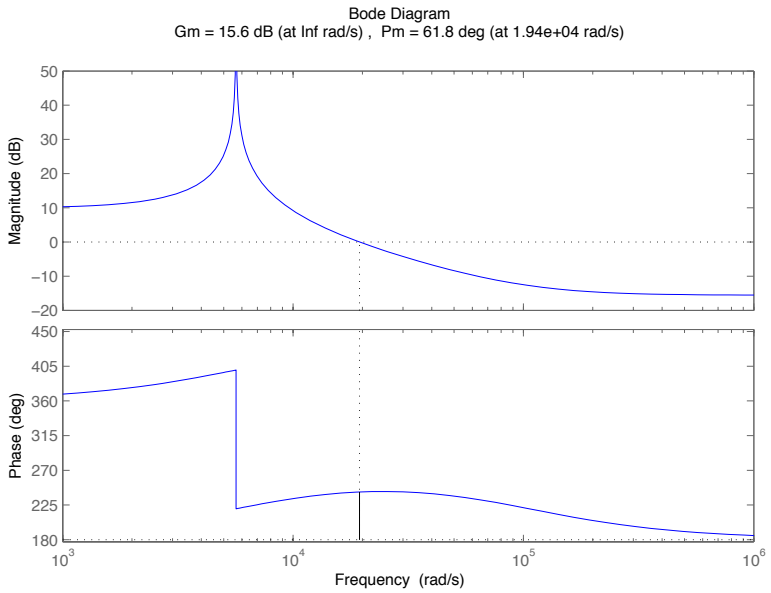

Figure 8.3: Bode plot of transfer function (8.14)

Adding a PD controller to (8.10) gives an open-loop transfer function of

$$\begin{aligned} H(s) &= K_b \frac{(P+Rs)\left(\frac{s}{\omega_1}-1\right)}{\frac{s^2}{\omega_0^2}+1} \\ &= \frac{\left(\frac{K_b R}{\omega_1}\right)s^2 + \left(\frac{K_b P}{\omega_1} - K_b\right)s - K_b P}{\frac{s^2}{\omega_0^2}+1} \end{aligned} \tag{8.14}$$

A Bode plot of this transfer function with our choice of P and R is shown in Figure 8.3. We have reduced the unity-gain frequency to 2×10^4 rad/s and have a 62° phase margin.

Figure 8.4 shows simulated waveforms from the boost regulator with this PD controller. The figures shows the response to a step in V_{ref} from 250V to 240V at 400μs and from 240V to 260V at 800μs. In both cases the converter takes about 200μs (20 PWM cycles) to respond without overshoot.

8.4 Exercises

Small signal analysis of teh buck-boost converter

Add ESR of capacitor

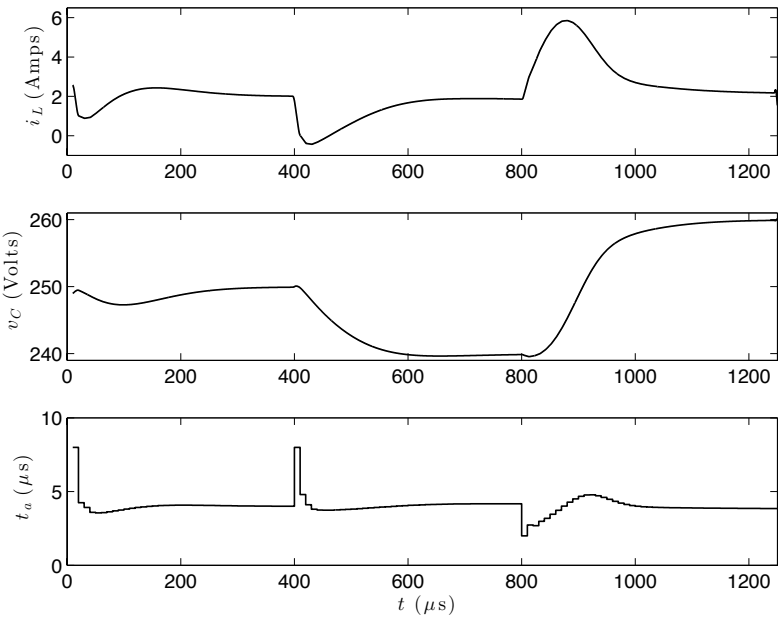

Figure 8.4: Simulated response of boost converter with PD controller.

Add load resistance

Chapter 9

Isolated Converters

The buck, boost, and buck-boost converters discussed in previous chapters are the widely used in applications when the input and output can share a supply terminal and where the step-up or step-down ratio is *small*. Whenever either of these conditions is not true, an *isolated* converter is required. Isolated converters, as the name imply isolate the input and output ports so that they may *float* relative to one another. In some cases the absolute input and output voltages may differ by thousands of volts. Most isolated converters use a transformer at their core to provide isolation.

Transformer-based converters are also used when the step-up or step-down ratio is *large* — particulary when large amounts of power are being converted. The use of a transformer in these cases enables a more economical and efficient converter. The switch of a buck converter must handle both the high input voltage of the converter and its high output current. With a transformer we can design a converter where a switch that handles the high input voltage only sees the low input current and a switch that handles the high output current only sees the low output voltage. For a converter with a $N:1$ step-down ratio this can reduce the cost of the switches by a factor of N^2. If isolation is not required, a high-ratio converter may employ an *autotransformer*.

In this chapter we start our discussion of isolated converters by introducing the ideal transformer. We go on to discuss the limitiations of real transformers — that include *magnetizing* and *leakage* inductance. Finally we show a full-bridge converter that employs a transformer in an isolated analog of the buck converter.

9.1 The Ideal Transformer

Most isolated converters use *transformers*. A schematic symbol for a transformer with two windings[1] is shown in Figure 9.1. The symbol is two back-to-back *coils*

[1] While a transformer may have any number of windings we will most commonly deal with two-winding transformers.

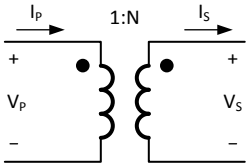

Figure 9.1: An ideal transformer with turns ratio $1 : N$ maintains the relationships $V_S = NV_P$ and $I_S = \frac{I_P}{N}$.

or *windings*. While the same winding symbol is used for an inductor, there is no inductance in an ideal transformer. The transformer is a four-terminal device, the left two terminals are connected to the *primary* winding of the transformer. The right two terminals are connected to the *secondary* winding.

The dots in the symbol indicate the polarity of the windings. When a voltage V_P is applied to the primary so that the terminal labeled with the dot is positive (as shown), a voltage V_S will appear at the secondary with the dotted terminal positive as well. When the primary is driven so that current I_P flows *into* the dotted terminal, current I_S flows *out of* the dotted terminal of the secondary.

A transformer is characterized by its turns ratio, $N_P : N_S$. Physically this is the ratio of the number of *turns* in the primary and secondary windings. Electrically, this ratio determines the ratio of voltages and currents in the primary and secondary. For an ideal transformer, the primary and secondary voltages and currents are related by

$$V_S = \frac{N_S}{N_P} V_P \tag{9.1}$$

$$I_S = \frac{N_P}{N_S} V_P \tag{9.2}$$

An ideal transformer with a $1 : N$ turns ratio, as depicted in Figure 9.1 steps up the voltage by N and steps down the current by N so that the power into the primary is equal to the power out of the secondary

$$P_S = V_S I_S = \frac{N_S}{N_P} V_P \frac{N_P}{N_S} I_P = V_P I_P = P_P \tag{9.3}$$

A transformer does not have a preferred direction. Current, and power, can flow in either direction. If you reverse the primary and secondary of Figure 9.1 you have a transformer with a $N : 1$ turns ratio.

Similarly the voltage of a transformer is not necessarily set by its primary terminal, or even by a single set of terminals. A transformer always maintains

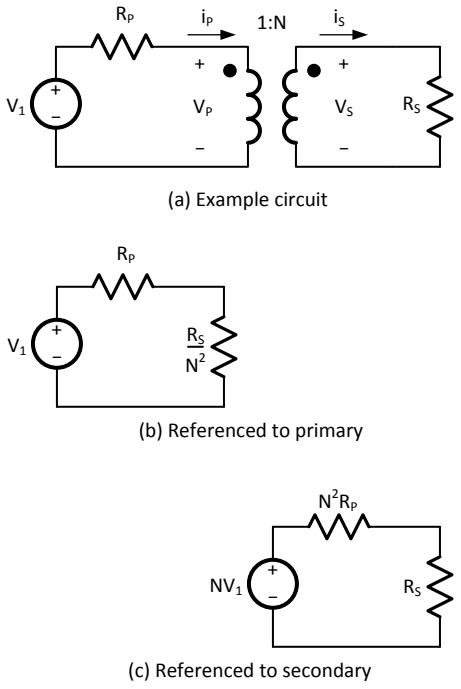

Figure 9.2: (a) A circuit using a transformer. (b) The equivalent circuit referenced to the primary. (c) The equivalent circuit referenced to the secondary.

a voltage ratio between its terminals that is set by its turns ratio. However, in general a simultaneous set of equations must be solved to determine what the voltage is.

Figure 9.2(a) shows an example of a circuit using an ideal transformer. The primary is driven by a voltage source V_1 in series with a resistor R_P. The secondary is loaded by a resistor R_S. To solve for the voltages in currents in this circuit we can either write and solve a set of equations, or we can eliminate the transformer by *reflecting* components from one side of the transformer to the other side.

Using the properties of ideal transformers and Ohm's law we can write:

$$I_P = \frac{V_1 - V_P}{R_P} \tag{9.4}$$

$$V_P = \frac{1}{N} V_S \tag{9.5}$$

$$I_P = N I_S \tag{9.6}$$

$$I_S = \frac{V_S}{R_S} \tag{9.7}$$

We can solve these four equations in four unknowns for V_P, V_S, I_P, and I_S.

An alternate approach is to reflect the secondary resistor R_S to the primary as shown in Figure 9.2(b). If we reflect an impedance Z from winding A of a transformer to winding B, the reflected impedance value becomes

$$Z_R = \left(\frac{N_B}{N_A}\right)^2 Z \tag{9.8}$$

In this case we reflect R_S from a seconeary with $N_S = N$ to a primary with $N_P = 1$ and have $R_{SR} = \frac{R_S}{N^2}$. Intuitively impedance is transformed as the square of the turns ratio N because voltage is proportional to N, current is proportional to $1/N$, and impedance is V/I which varies as N^2.

We can get the same result by combining (9.5), (9.6), and (9.7) to write

$$I_P = \frac{N^2 V_P}{R_S} \tag{9.9}$$

This reflection applies to all impedances whether they are resistances $Z = R$ inductances $Z = Ls$, or capacitances $Z = \frac{1}{Cs}$. Thus inductances and resistances increase when reflected to a winding with more turns and capacitances decrease.

Applying the same technique we can reflect the primary circuit to the secondary as shown in Figure 9.2(c). Here, in addition to applying (9.8) we must also apply a source reflection rule

$$V_R = \left(\frac{N_B}{N_A}\right) V \tag{9.10}$$

$$I_R = \left(\frac{N_A}{N_B}\right) I \tag{9.11}$$

Using any of these methods we can solve the original circuit and observe that

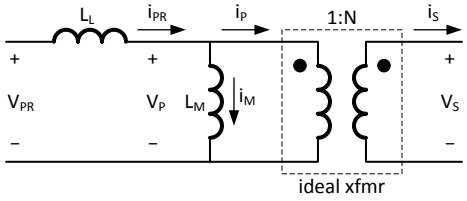

Figure 9.3: A real transformer includes a *magnetizing inductance* L_M that prevents it from handling DC voltage and a *leakage inductance* L_L that limits the upper end of its frequency range.

$$V_S = \frac{NV_1R_S}{R_S + N^2R_P} \tag{9.12}$$

$$I_S = \frac{NV_1}{R_S + N^2R_P} \tag{9.13}$$

$$V_P = \frac{V_1R_S}{R_S + N^2R_P} \tag{9.14}$$

$$I_P = \frac{N^2V_1}{R_S + N^2R_P} \tag{9.15}$$

$$= \frac{V_1}{\frac{R_S}{N^2} + R_P} \tag{9.16}$$

An ideal transform is itself a perfect isolated voltage converter. If we apply an input voltage V_P across the primary and set the turns ratio $N_P : N_S$ appropriately, we can get any desired voltage V_S out of the secondary. If it were possible to economically fabricate ideal transformers there would be little need for any other voltage converter.

9.2 A Real Transformer

Unfortunately it is not possible to fabricate ideal transformers. As shown in Figure 9.3, all real transformers have a *magnetizing inductance* that prevents them from passing DC voltage. They also have *leakage inductances* that add parasitic elements to the circuit that in practice limits the upper frequency at which the transformer can operate. In Chapter 23 we will examine the physical reason for these parasitic inductances.

Figure 9.3 shows both inductances in the primary circuit. When convenient the magnetizing can be reflected to the secondary using (9.8). However, there is

only one L_M per transformer. There is not a separate magnetizing inductance per winding.

There is a separate leakage inductance for each winding, and if the secondary leakage inductance is reflected to the primary, it appears to the right of the magnetizing inductance. However, because the leakage inductance is typically small compared to the magnetizing inductance, we typically combine the reflected leakage inductance of the secondary with the leakage inductance of the primary and treat them as a single element as shown in Figure 9.3.

9.2.1 Magnetizing Inductance

Considering just the magnetizing inductance, we see that if we apply a DC voltage to the real transformer the magnetizing current I_M will increase without bounds. If we assume the current is initially zero, it will be given by:

$$I_M(t) = \frac{V_P t}{L_M} \tag{9.17}$$

For most transformers, which have a core made from a ferromagnetic material, at some time t_s the magnetizing current will saturate the core[2]. At this point the inductance will drop dramatically and the current will increase even more steeply. Unless there is some means to limit the current some circuit component is likely to be destroyed.

The saturation magnetizing current I_{MS} limits any real transformer to handle a maximum number of *volt-seconds* X_{VS} of input before the input polarity must be reversed to reduce the magnetizing current.

$$X_{VS} < I_{MS} L_M \tag{9.18}$$

If the magnetizing current oscillates (with a triangle waveform) about zero, the transformer can actually handle twice this number of volt-seconds: half to get back to zero and half to reach the maximum current in the new direction. Thus, the minimum frequency a transformer can operate at is given by

$$\begin{aligned} f_{\min} &= \frac{V_P}{2X_{VS}} \\ &= \frac{V_P}{2I_{MS}L} \end{aligned} \tag{9.19}$$

The higher the applied voltage, the higher the frequency the transformer must operate at to avoid saturation.

In addition to operating at $f > f_{\min}$ it is important to keep the volt-seconds across any transformer exactly balanced to avoid saturation. If, over time, more volt seconds are applied in one direction than the other the magnetizing current will slowly creep up and eventually saturate the transformer.

[2]While air-core transformers don't saturate, they still have an upper current limit set by thermal constraints and by the maximum current that can be handled by the circuit.

9.2.2 Leakage Inductance

The leakage inductance governs how quickly current in the transformer can switch directions and hence sets an upper limit on operating frequency for many applications. Consider a transformer operating with a square-wave primary voltage with magnitude $\pm V_P$ and carrying current $\pm I_P$.

When the primary voltage switches polarity, from V_{PR} to $-V_{PR}$ (to avoid saturating the magnetizing inductance), the amount of time taken for the current to reverse from I_P to $-I_P$ is given by.

$$t_r = \frac{2L_L I_P}{V_{PR}} \tag{9.20}$$

During this period the transformer is not doing useful work — i.e., transferring power from the input to the output. Thus, it is not practical to operate right at the minimum period given by $2t_r$. Some additional time is required for the transformer to transfer power.

The time taken to reverse current limits the upper operating frequency of the transformer to

$$\begin{aligned} f_{\max} &= \frac{1}{2t_r} \\ &= \frac{V_{PR}}{4L_L I_P} \end{aligned} \tag{9.21}$$

From (9.19) and (9.21) we see that the operating frequency range of a real transformer is bounded below by magnetizing inductance and above by leakage inductance.

$$\frac{V_P}{2I_{MS}L_M} \leq f \leq \frac{V_{PR}}{4L_L I_P} \tag{9.22}$$

When characterizing a real transformer the magnetizing inductance can be measured by measuring the inductance of one winding while the other winding is left open. The magnetizing inductance is just the inductance of one winding if it were wound on the same core by itself. The leakage inductance can be measured by measuring the inductance of one winding while the other winding is shorted. Shorting the secondary also shorts the primary (and the magnetizing inductance) causing all of the applied voltage to drop across the leakage inductance.

9.3 A Full-Bridge Converter

Figure 9.4 shows a full-bridge converter that converts an input voltage V_1 to an isolated output voltage V_2 using a transformer. The converter can be thought

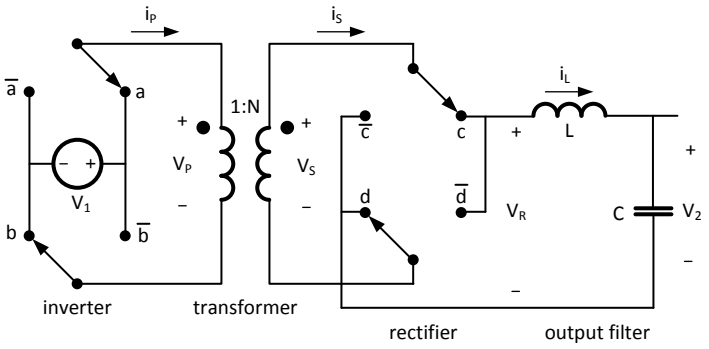

Figure 9.4: A full-bridge converter. An inverter converts a DC input source V_1 to an AC square wave to pass it through a transformer that steps the voltage up (or down) and provides isolation. A rectifier converts the transformer output back to DC. An output filter smooths the transformer output and facilitates regulation.

of as four separate components as labeled across the bottom of the figure. DC input voltage source V_1 is first converted to an AC square-wave voltage on V_P by an *inverter* composed of switches **a** and **b**. The relative phase of the square waves controlling switches **a** and **b** determines the duty factor of the square wave on V_P. Next, the AC voltage on V_P is stepped up by a factor of N by the transformer. The AC output of the tansformer on V_S is converted back to a DC voltage by the rectifier composed of switches **c** and **d**. Finally, an output filter smooths the output, converting a rectified square-wave to a steady DC voltage.

Waveforms illustrating operation of the full-bridge converter are shown in Figure 9.5. The waveforms show operation for a transformer with magnetizing inductance but no leakage inductance.

Initially, at t_0, switches $\overline{\mathbf{a}}$, **b**, **c**, and **d** are closed and current $-I_{MP}$ flows in the magnetizing inductance. The connection to $\overline{\mathbf{a}}$ and **b** shorts the primary winding making the voltage across the transformer zero. The connection through switches **c** and **d** causes I_S to be equal to I_L. Inductor current I_L is ramping down with slope $-V_2/L$ because with the transformer shorted, $-V_2$ appears across the inductor.

At time t_1 switches **c** and **d** switch to positions $\overline{\mathbf{c}}$ and $\overline{\mathbf{d}}$ causing the direction of secondary current to reverse so that $I_S = -I_L$. In the absence of leakage inductance this current reverses instantly. With leakage some finite time — and different switch positions — are needed to deal with the energy stored in the leakage inductance during current reversals. The rectifier switches can be switched at any point in the interval (t_0, t_2) because the transformer is shorted

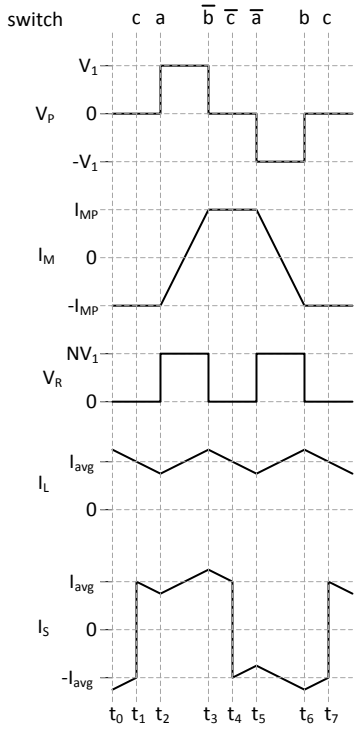

Figure 9.5: Waveforms showing operation of the full-bridge converter of Figure 9.4.

during this entire interval.

At time t_2, switch **a** switches to position **a** applying $+V_1$ across the primary. This causes the magnetizing current to start ramping up with slope V_1/L_M. The applied voltage is stepped up on the seconary with $V_S = NV_1$. This voltage causes the current in the output inductor to start ramping up with slope $(NV_1 - V_2)/L$.

At time t_3, switch **b** switches to position $\overline{\mathbf{b}}$ again shorting the primary winding. With no applied primary voltage the magnetizing current levels off at I_{MP}. Shorting the primary also causes the secondary to be shorted causing the output current to ramp down with slope $-V_2/L$ — the same as at time t_0.

The transitions at t_4, t_5, and t_6 are identical to the transitions at t_1, t_2, and t_3 but with the transformer polarity reversed. A negative voltage is applied across the primary. This gives a negative voltage across the secondary, and the rectifier is switched in the opposite direction to give a positive voltage V_R across the output filter. The rectified volage V_R across the output filter is the same at t_4, t_5, and t_6 as at t_1, t_2, and t_3.

The output filter acts just like the inductor and capacitor in a buck converter. The rectified secondary voltage V_R is a square wave with magnitude NV_1 and duty factor $D = (\pi - \phi_{ab})/\pi$, where ϕ_{ab} is the phase difference between the operation of switches **a** and **b**. Thus, as with a buck converter, the output voltage in periodic steady state is given by

$$V_2 = DNV_1 \tag{9.23}$$

The transformer steps up the voltage by a factor of N and the output filter acts like a buck converter and steps it down by the duty factor D.

9.4 Effect of Leakage Inductance

Figure 9.6 shows waveforms illustrating the operation of the full-bridge converter of Figure 9.4 including the effects of leakage inductance. This operation requires operating switches **c** and **d** so that both sides of each switch are closed simultaneously, shorting both the transformer secondary and providing a current path for the output filter. This state is denoted by part of the **c** waveform where the signal is shown both high and low — indicating that both the **c** and $\overline{\mathbf{c}}$ switches are closed at the same time (and similarly for switch **d**).

Consider the sequency of events starting at t_1 when switch **a** turns on and both sides of switches **c** and **d** are turned on. The input voltage V_1 is applied across the primary by switches **a** and **b** and the secondary is shorted by switches $\overline{\mathbf{c}}$ and **d**. This applies the full input voltage V_1 across the leakage inductance L_L and the transformer current ramps upward from its original negative state. This is reflected in the rise of I_S with slope V_1/NL_L. The slope of I_P (not shown) is V_1/L_L and the secondary current is stepped down by $1/N$. With switches $\overline{\mathbf{c}}$ and **d** also closed, a voltage of $-V_2$ is across the output inductor and its current ramps down with slope $-V_2/L$.

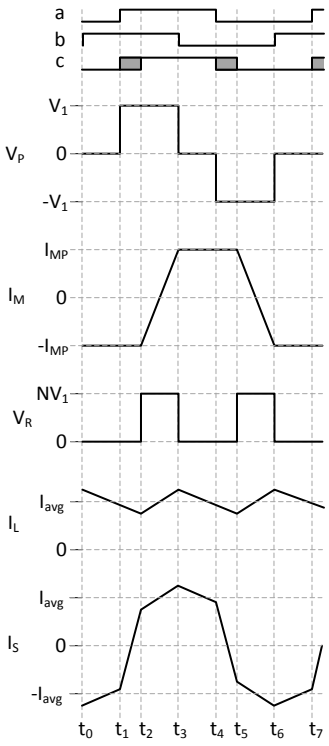

Figure 9.6: Waveforms showing operation of the full-bridge converter of Figure 9.4 including the effects of leakage inductance.

With all four rectifier switches closed, the transformer current is momentarily decoupled from the output inductor current. This decoupling is necessary to allow time for the current in the leakage inductance to reverse. During this *commutation* period the transformer current rises through zero until at t_2 $I_S = I_L$. At this point switches $\overline{\mathbf{c}}$ and $\overline{\mathbf{d}}$ open and the transformer current and output inductor current are coupled with both positive. The delay due to leakage inductance t_L from t_1 to t_2 is given by

$$t_L = \frac{NL_L\hat{I_L}}{V_1} \tag{9.24}$$

where $\hat{I_L}$ is the average output inductor current during this period.

At t_2 only a small fraction of the input voltage is being dropped across L_L, so most of the voltage appears across the transformer causing V_S,and hence V_R to approach NV_1. This causes the output current, and secondary current because they are coupled, to ramp up with slope $(NV_1 - V_2)/L$. Also at t_2 the transformer voltage appears across the magnetizing inductance and the magnetizing current I_M ramps upward with slope V_1/L_M.

At time t_3 switch **b** moves to position $\overline{\mathbf{b}}$ and the transformer primary is shorted. This is reflected instantly across the transformer, since there is no delay due to L_L and V_S and V_R fall to zero at the same time. With no transformer voltage the magnetizing current levels at I_{MP} and the output inductor current I_L starts ramping downward with slope $-V_2/L$.

At time t_4 the process repeats in the opposite direction. All four rectifier switches are closed — shorting the secondary and decoupling the secondary and output currents. At the same time switch **a** switches to position $\overline{\mathbf{a}}$ applying $-V_1$ across the leakage inductance. After the leakage current ramps down to match $-I_L$ switches **c** and **d** open unclamping the secondary and re-coupling $-I_S$ with I_L.

Each half-cycle, a voltage pulse with width $t_P = t_3 - t_1 = t_6 - t_4$ is applied to the primary. The first t_L of each pulse is used to turn the current in the leakage inductance around. The remaining $t_S = t_P - t_L = t_3 - t_2$ transfers energy to the output inductor. It is the duty factor of the secondary $D_S = 2t_S/t_{\text{cy}}$ that determines the output voltage according to (9.23).

9.5 Final Comments

Because the full-bridge converter is really just a buck converter with an isolated step-up (or down) input stage, controlling the full-bridge converter is essentially the same as controlling the buck. With a few adjustments to account for the transformer ratio and for the leakage delay t_L, the methods of Chapter 5 can be applied to the full-bridge converter. We leave the details as an exercise.

In a practical full-bridge converter isolation of the control circuit is an issue. The controller is typically referenced to the secondary because it senses the

secondary voltage. The signals controlling the primary-side switches are passed through isolators to cross the unknown input-output voltage shift.

At first glance it might appear that timing of the switches would be difficult since the secondary switches need to change state exactly at the point when the current in the leakage inductor matches the current in the output inductor. In practice this critical timing is performed automatically by diodes. Even in low-voltage systems that use MOSFETs for the secondary switches (*synchronous rectification* Section **??**) diodes associated with the MOSFETs perform the initial switching. The MOSFET then turns on to reduce the voltage drop.

9.6 Exercises

transformer with more than two windings
- solve Figure 9.2 using equations
- solve some other xfmr examples
- matlab sim of bridge converter
- control of bridge converter
- current-fed bridge

Chapter 10

The Flyback Converter

The *flyback* converter is interesting because it uses the magnetizing inductance of the transformer for energy storage. This is fundamentally different than the other isolated converters for which the magnetizing inductance is a parasitic element. The productive use of the magnetizing inductance allows a flyback converter to be built with a single magnetic component that serves both as an isolating transformer and as the energy storage element. In contrast most other isolated converters require a separate storage inductor (like the output filter inductor in Figure 9.4) in addition to the transformer.

10.1 Flyback Topology and Operation

Figure 10.1 shows a schematic diagram of a flyback converter. The converter is quite simple consisting just of a transformer, two switches, and an output capacitor. Because of its importance we draw the magnetizing L_M inductance explicitly. This is not a separate component but rather part of the transformer.

Operation of the flyback converter is controlled by the two switches which close alternately. When switch **a** is closed input supply V_1 is connected across L_M and current i_M ramps up. In this phase of operation i_M flows around a loop that includes switch **a** and the input supply. Because switch **b** is open no current flows through the ideal transformer.

When switch **a** opens and switch **b** closes current i_M flows in a loop between L_M and the primary winding. This induces a current in the secondary winding that charges output capacitor C. During this period voltage V_2 is applied across the secondary, $-V_2/N$ is reflected to the primary, and i_M ramps down.

Figure 10.2 shows waveforms for the flyback converter illustrating how the magnetizing current i_M alternates between the primary and secondary circuits. When switch **a** is closed, primary current i_P is equal to magnetizing current i_M both of which ramp upward with slope $\frac{V_1}{L}$ because the input supply is applied across L_M.

When switch **b** is closed (and **a** is open) i_P is zero and i_M flows through

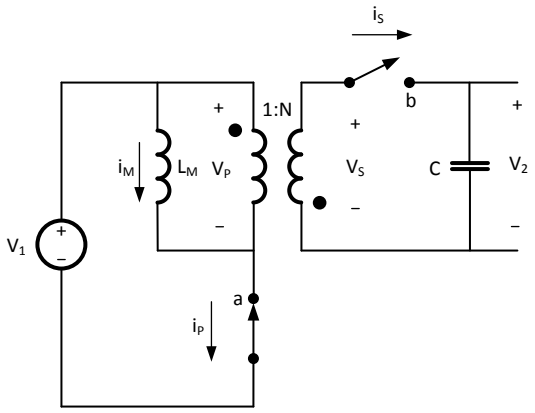

Figure 10.1: A flyback converter stores energy in the magnetizing inductance of the transformer L_M.

the secondary circuit. During this phase $i_S = \frac{i_M}{N}$. The reflected and inverted secondary voltage $-\frac{V_2}{N}$ appears across L_M causing i_M to ramp downward with slope $-\frac{V_2}{NL}$. Secondary current, being scaled by $1/N$ ramps with slope $-\frac{V_2}{N^2L}$.

In the periodic steady state the volt-seconds across L_M must balance and we have:

$$DV_1 = (1-D)\frac{V_2}{N} \tag{10.1}$$

Where D is the duty factor of switch **a**. Rewriting this gives the relation:

$$V_2 = \frac{ND}{1-D}V_1 \tag{10.2}$$

This is similar to the relationship for the buck-boost converter (6.6) but multiplied by the transformer turns ratio.

When the converter is not in the periodic steady state, the change in magnetizing current and capacitor voltage each cycle are given by:

$$\Delta i_i = \frac{t_{cy}\left(DNV_1 - (1-D)v_i\right)}{NL} \tag{10.3}$$

$$\Delta v_i = \frac{t_{cy}\left(1-D\right)\hat{i}_i - Ni_{\text{Load}}}{NC} \tag{10.4}$$

Converting to Laplace transforms gives:

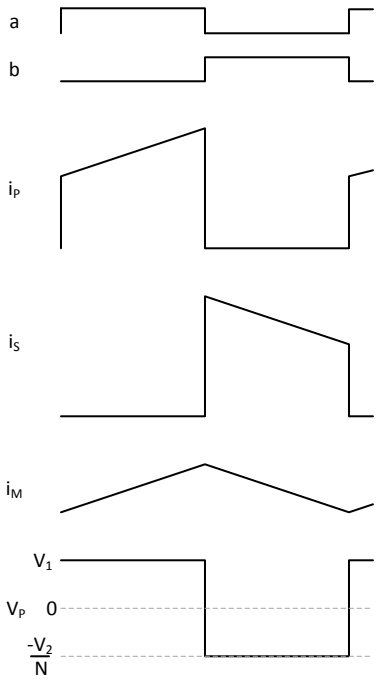

Figure 10.2: Waveforms showing operation of the flyback converter of Figure 10.1. When switch **a** is closed input supply V_1 is across L_M and magnetizing current i_M ramps up. When **a** opens and **b** closes L_M discharges its energy through the transformer into capacitor C and i_M ramps down.

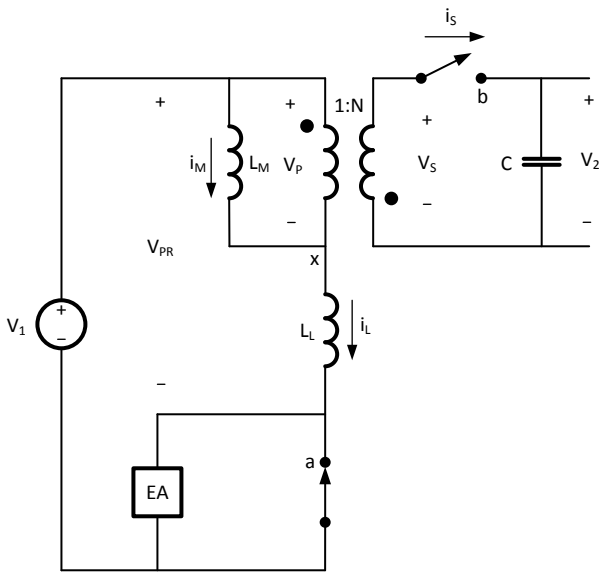

Figure 10.3: Flyback converter of Figure 10.1 with leakage inductance L_L shown. Energy absorbing device EA is used to absorb energy stored in L_L when switch **a** opens.

$$i(s) = \frac{V_1 N D(s) - (1 - D(s))v_i(s)}{NLs} \tag{10.5}$$

$$v(s) = \frac{(1 - D(s))i(s) - N i_{\text{Load}}(s)}{NCs} \tag{10.6}$$

10.2 Effect of Leakage Inductance

The leakage inductance of the transformer complicates the operation of the flyback converter delaying the transfer of current between the primary and secondary and adding losses as energy stored in the leakage inductance must be *dumped*. Figure 10.3 shows the flyback converter with the leakage inductance shown explicitly. Waveforms showing operation with finite leakage inductance are shown in Figure 10.4.

With leakage inductance primary current i_P cannot immediately jump to the level of the magnetizing current i_M when switch **a** closes. Instead i_P ramps up with slope $\frac{V_1 + V_2/N}{L_L}$. As more of i_M is diverted into i_P, i_S ramps down. During this period of commutation both switches **a** and **b** are closed. Finally,

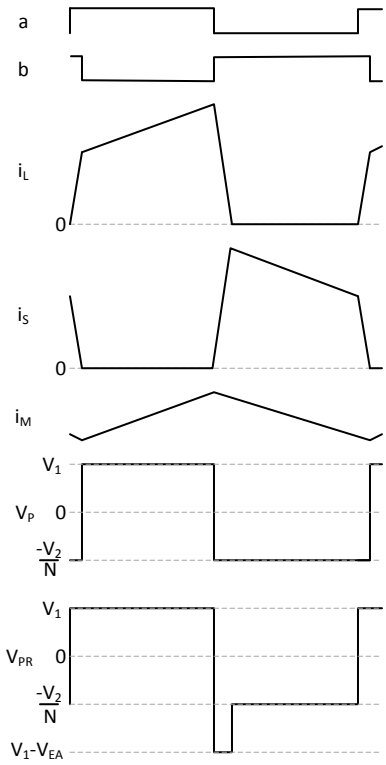

Figure 10.4: Waveforms showing operation of flyback converter of Figure 10.3. Leakage inductance introduces delay into *commutation* of current from primary to secondary.

when i_S reaches zero, switch **b** opens and the flyback operates as in Figure 10.2 until switch **a** opens. Because the input voltage V_1 is not applied across the magnetizing inductance until i_L reaches i_M the effective pulse width is reduced by the delay of the ramp.

When switch **a** opens the leakage current is directed into an *energy absorber* (EA). The voltage across EA must be greater than $V_1 + V_2/N$ to decrease i_L. In practice the energy absorber is often the MOSFET used for switch **a** operating in its avalanche breakdown region (Section **??**). The energy lost by dumping i_L into EA is at least the leakage energy $E_L = i_L^2 L_L/2$ itself and may be greater. Some flyback converters use methods to recover this leakage energy. One such method is described in Section 10.3.

As i_L ramps down — dumping its energy into EA — the balance of i_M (scaled by $1/N$) flows into i_S. When i_L reaches zero, i_S reaches i_M/N.

The bottom two traces of Figure 10.4 show the ideal primary voltage V_P and the real primary voltage V_{PR}, which includes the drop across the leakage inductance. These two traces differ during the periods where i_L is ramping rapidly to commutate the current between primary and secondary. When switch **a** first turns on V_P is $-V_2/N$ while V_{PR} is V_1. This applies $V_1 + V_2/N$ across L_L to ramp i_L upward. When switch **a** turns off V_P is again $-V_2/N$ while V_{PR} is $V_1 - V_{EA}$ applying $V_{EA} - V_1 - V_2/N$ in the opposite direction across L_L to ramp i_L downward.

10.3 Diagonal Flyback

The energy stored in the leakage inductance can be recycled to the input supply by adding three switches to the flyback converter as shown in Figure 10.5. A second **a** switch is added to isolate the top of the primary from the input supply when **a** turns off. Two **c** switches are added to connect the primary *backwards* across the input supply to recover leakage energy.

During the middle of the waveforms in Figure 10.4 when switch **a** turns off and **b** turns on, switch **c** also turns on just long enough for the leakage energy to be returned to the input supply. During this period, with switch **b** closed, V_2 appears across V_S and is reflected as $-V_2/N$ across V_P. With the bottom **c** switch closed this places point **x** at V_2/N. For the leakage current to reduce we must have $V_1 > V_2/N$. Then the voltage across the leakage inductance is $V_1 - V_2/N$ and i_L ramps down with slope $\frac{NV_1 - V_2}{NL_L}$. When i_L reaches zero switches **c** turn off.

The **c** switches are typically implemented as a pair of diodes with the cathodes up. This makes their timing automatic. When switches **a** open the flow of leakage current forward biases diodes **c** turning them on. When the leakage current reaches zero (i.e., when all of the magnetizing current is flowing to the secondary) the diodes turn off.

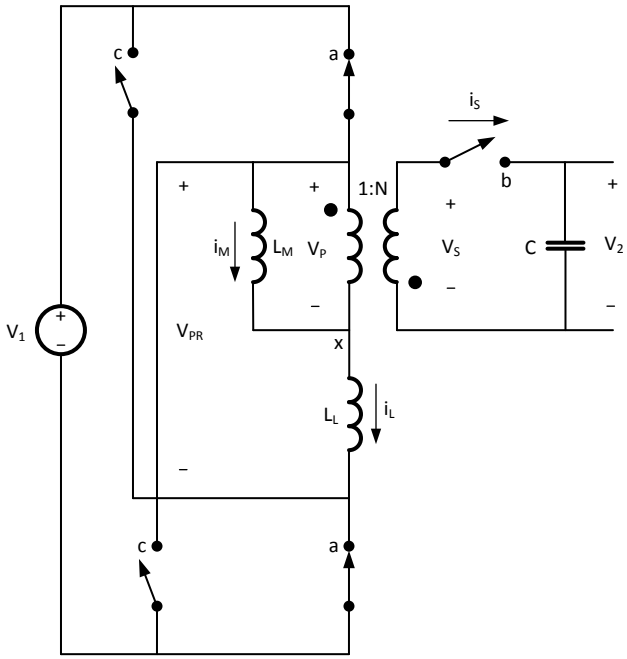

Figure 10.5: A diagonal flyback converter recycles energy in the leakage inductance to the input supply.

10.4 Primary-Side Sensing

One advantage of the flyback converter is that the output voltage can be sensed on the primary side. This allows the converter to be controlled entirely from the primary side — without the need for an optocoupler, or other isolation device, to communicate control signals across the transformer.

When switch **b** is closed, output voltage V_2 appears directly across the secondary and is reflected so that $-V_2/N$ appears across the primary. This scaled version of the output voltage can be sensed after the leakage current has stopped flowing. Alternately an auxiliarly winding can be added to the transformer allowing the secondary voltage to be sensed any time switch **b** is closed.

10.5 Exercises

give some component values calculate slopes and give waveforms

matlab simulation

controller

Chapter 11

The Forward Converter

The forward converter, like the full-bridge converter, is an isolated version of a buck converter. Unlike the full-bridge, however, the forward converter operates on only half of the transformer's B-H curve and requires a third winding (called a *recovery winding* on the transformer to discharge the magnetizing inductance each cycle.

11.1 Forward Topology and Operation

Figure 11.1 shows a schematic diagram of a forward converter. The converter uses a transformer with three windings: the primary **P**, the secondary **S**, and the recovery winding **R**. Winding **R** returns the magnetizing energy to the input supply. Four switches control operation of the converter. Switches **a** and **b** apply the input voltage V_1 across the primary and recovery windings respectively — in opposite directions. Switch **c** connects the secondary winding to inductor L. Switch **d** supplies current to inductor L when switch **c** is open. Switch **a** is usually realized with a MOSFET. Switches **b-d** may be realized with diodes since they switch on whenever the voltage across the switch is in the direction of the switch's *arrow*.

Waveforms showing operation of the converter, ignoring leakage inductance, are shown in Figure 11.2. The converter operates in two phases. During the first phase switches **a** and **c** are closed. During this phase input voltage V_1 is applied across the primary. The transformed voltage NV_1 appears across the secondary and is applied to inductor L via switch **c** causing i_L to ramp up. During this phase, the inductor current flows in the secondary winding $i_S = i_L$, and the transformed current flows in the primary $i_P = Ni_L$. The voltage applied to the primary also causes the magnetizing current i_M to ramp up, starting at zero at the beginning of the phase. In typical applications i_M is much smaller than i_L. The first phase must end before i_M reaches saturation.

During the second phase switches **a** and **c** are open. Thus no current flows in the primary or secondary: $i_P = i_S = 0$. Switch **d** is closed to supply inductor

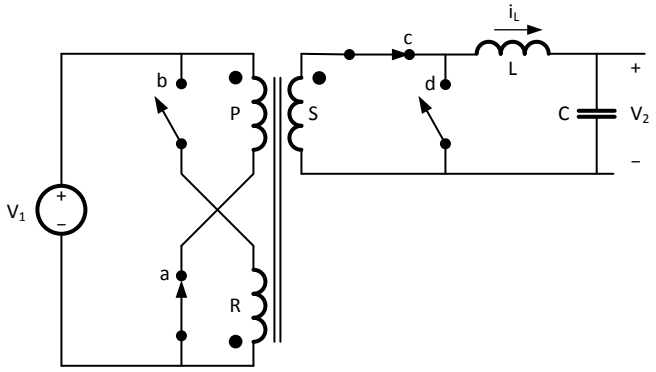

Figure 11.1: A forward converter uses a recovery winding **R** to return the energy in the magnetizing inductance to the power supply V_1 each cycle.

L with current. The voltage across the inductor is $V_L = -V_2$ causing i_L to ramp down. Switch **b** is closed to apply the input voltage across the recovery winding in the reverse direction (note the dots in Figure 11.1). During this phase the voltage across the recovery winding is $V_R = -V_1$ and this transformed voltage appears across the primary $V_P = -\frac{N_P}{N_R}V_1$, and the primary-referenced magnetizing inductance. This causes the current in the magnetizing inductance to ramp down. When this current reaches zero the second phase ends and switch **b** opens. The voltage across the primary adds to the supply voltage resulting in a voltage across open switch **a** of $V_a = V_1\left(1 + \frac{N_P}{N_R}\right)$.

When i_M reaches zero, an optional third phase may be inserted where only switch **d** is closed. During this *dead* phase no current flows in any of the transformer windings. Current continues to flow in the loop formed by switch **d**, L, and C.

The winding ratio of the primary and recovery windings determines a maximum duty factor for the forward converter. This minimum duty factor is needed to ensure volt-second balance — i.e., that a net zero volt-seconds are applied to the magnetizing inductance over the course of a cycle. Referenced to the primary, the volt-seconds applied during a cycle are $V_1 t_a - V_1 \frac{N_P}{N_R} t_b$. Thus, to achieve volt-second balance we must have:

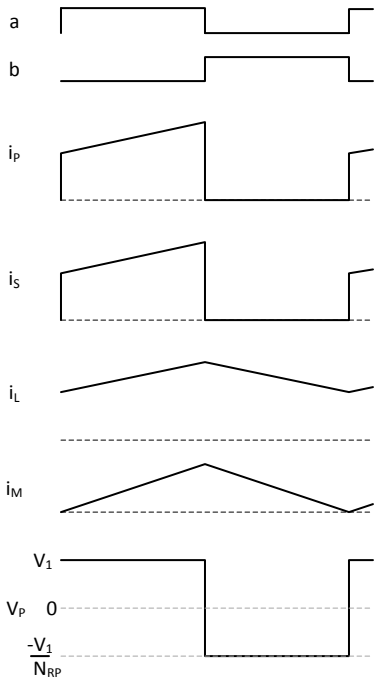

Figure 11.2: Waveforms showing operation of the forward converter.

$$
\begin{aligned}
t_a &= \frac{N_P}{N_R} t_b \\
D_{\max} &= \frac{t_a}{t_a + t_b} \\
&= \frac{N_P}{N_P + N_R}.
\end{aligned} \tag{11.1}
$$

For the typical case where $N_P = N_R$, $D_{\max} = 0.5$. Most converters operate with equal turns on windings P and R. Decreasing N_R to increase duty factor results in very high voltages across switch **a** during the second phase.

Smaller duty factors can be accomodated by adding a dead phase after the magnetizing current goes to zero. However, the duty factor cannot be larger than $D_{\max}$ or the magnetizing current won't be fully reset at the end of each cycle. Over time it will slowly increase until it reaches saturation — with catastrophic consequences.

It is worth pointing out that the cost of the recovery winding is negligible. This winding carries only the magnetizing current which is typically a small fraction (usually less than 10%) of the primary current. Hence it consumes only a small fraction of the copper on the primary side of the transformer.

11.2 Comparison of Converters

Like the flyback converter (Chapter 10), in the forward converter the primary and secondary each only conduct during half the cycle. However, unlike the flyback — where the primary conducts during the first half cycle and the secondary conducts during the second — in the forward converter both the primary and secondary conduct during the first half cycle and the recovery winding conducts during the second half cycle.

The forward converter is typically used at *medium* power levels between 100W and 500W. Below 100W a flyback converter is usually preferred for economy — because it uses only a single magnetic component. Above 500W a full-bridge converter is usually preferred because it makes better use of the core's B-H curve — using both sides —and because it makes better use of the transformer's copper — flowing current on both half cycles. In the middle, the forward converter is a useful option.

The forward converter requires two magnetic components, but they can be optimized better than the flyback transformer since the transformer isn't required to store energy. Also the current requirement for the forward magentics are half that of the flyback's transformer. Because the flyback only supplies current during half the cycle (depending on duty factor), the average current in the windings when they are conducting is twice that of the inductor in the forward converter which supplies current over the entire cycle.

The transformer in a forward converter is bigger than a comparable full-bridge converter because it uses only half of the B-H curve, and because it

uses the copper only half the time. This is offset at lower power levels by the advantage of requiring fewer switches than the full-bridge.

Of course these power numbers are only guidelines. Depending on parameters and constraints the crossover points between the differnet converter topologies may vary. However, the order of the three topologies — flyback at the lowest power levels, forward at middle levels, and full-bridge at the highest levels — usually holds across parameters and constraints.

11.3 Exercises

affect of leakage inductance.

allow negative magnetizing current equivalence to center-tapped bridge

matlab simulation without leakage

with leakage

spice using ideal switches

control the forward converter

Chapter 12

Discontinuous Conduction Mode

Up to now we have been discussing converters, isolated and non-isolated, that operate in *continuous conduction mode* (CCM). That is, their inductor current flows *continuously* over the entire switching cycle. In this chapter we introduce a converter that operates in *discontinuous conduction mode* (DCM) where the inductor current flows during part of the cycle and then is zero for the remainder of the cycle. We illustrate DCM in the context of a buck converter. However, it can be applied to any type of converter. We will see that a buck converter operating in DCM has very different dynamics — and hence requires a different control law — than the same converter operating CCM.

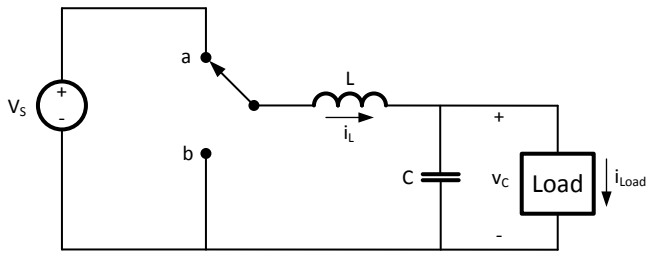

Figure 12.1: A buck converter — repeated from Figure 2.3.

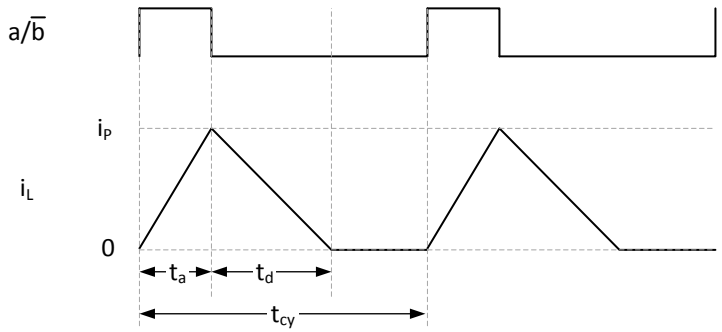

Figure 12.2: Waveforms for a buck converter operating in discontinuous mode. Each cycle current i_L starts at zero, ramps up during the time t_a switch **a** is on, and then ramps down to zero at time t_d.

12.1 The Buck Converter Operating in DCM

Waveforms for the buck converter of Figure 12.1 operating in discontinuous conduction mode are shown in Figure 12.2. Each cycle starts with $i_L = 0$. When the switch is in position **a** current ramps up with slope $\frac{V_S - v_c}{L}$. At time t_a the current has reached its peak value:

$$i_P = \frac{t_a \left(V_s - v_c\right)}{L} \tag{12.1}$$

At this point the switch moves to position **b** and i_L ramps down with slope $\frac{-v_c}{L}$ reaching zero after time:

$$t_d = \frac{t_a \left(V_s - v_c\right)}{v_c}. \tag{12.2}$$

So the width of the base of the triangular current pulse is

$$t_a + t_d = \frac{t_a V_s}{v_c}, \tag{12.3}$$

and the average current over the cycle is the area of this triangle divided by the cycle time:

```
if(i_l < 0) % in discontinuous mode no negative current
  i_l = 0 ;
end
```

Figure 12.3: To simulate a buck converter in discontinuous mode we add this `if` statement to the code of Figure 4.1 to prevent i_L from going negative.

$$\hat{i}_L = \frac{t_a^2 V_S (V_s - v_c)}{2Lt_{\rm cy} v_c}, \tag{12.4}$$

$$= \frac{t_{\rm cy} D^2 V_S (V_s - v_c)}{2Lv_c}. \tag{12.5}$$

In contrast to the buck operating in continuous condition mode (CCM) (Figure 3.1), in DCM i_L starts each cycle at zero. (12.5) shows that the average current over a cycle $\hat{i}_L$ varies directly with D^2, the square of duty factor. In contrast, in CCM, the derivative of the current varies linearly with D.

Discontinuous mode results when switch **b** is realized as a diode (or a MOSFET controlled to act like a diode) at low currents and low pulse widths.

12.2 Simulation

To simulate the buck converter in discontinuous mode, we simply add the `if` statement shown in Figure 12.3.

[More]

12.3 Control

To develop a control law for the buck operating in DCM we linearize (12.5) by defining a control variable $x = D^2$ giving:

$$\hat{i}_L = \frac{t_{\rm cy} x V_S (V_s - v_c)}{2Lv_c}. \tag{12.6}$$

If we assume v_c is a constant V_C for the purposes of control, this is just a simple linear dependance of $i_L = Kx$. The voltage just integrates this:

$$v(s) = \frac{\hat{i}_L(s)}{Cs} \tag{12.7}$$

$$= \frac{x(s)K}{s} \tag{12.8}$$

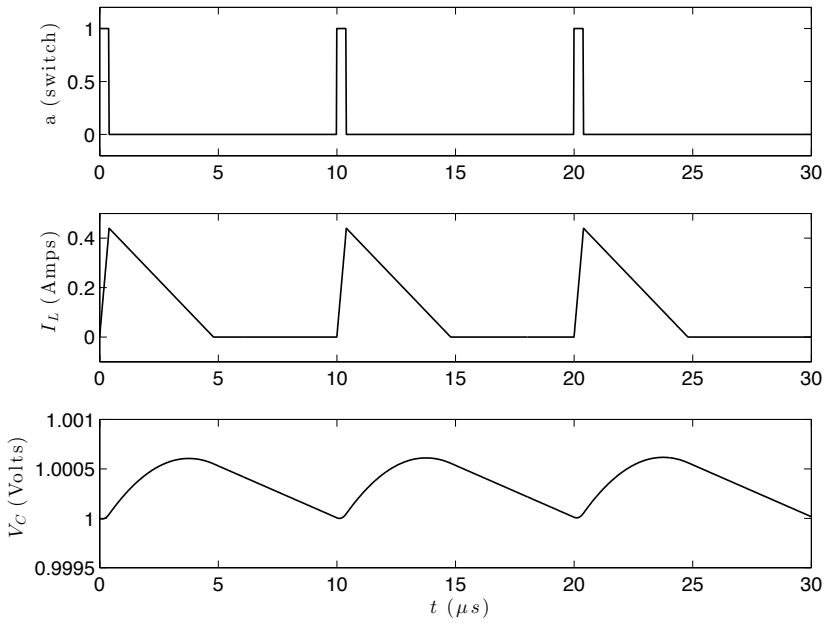

Figure 12.4: Simulated waveforms showing **a** top, i_L middle, and v_c bottom for three cycles of a buck converter operating in discontinuous conduction mode.

Where gain constant K is given by

$$K = \frac{t_{\text{cy}} V_S \left(V_s - V_c\right)}{2LCV_c}. \tag{12.9}$$

Note that because gain K depends strongly on V_c we may need to change our control parameters as V_c changes.

To control this simple first-order system we can use a PI controller where the control law is

$$x(s) = \left(P + \frac{Q}{s}\right) e(s) \tag{12.10}$$

The open-loop response is

$$H(s) = \frac{KPs + KQ}{s^2}, \tag{12.11}$$

and the closed loop response is

$$G(s) = \frac{H(s)}{1 + H(s)} \tag{12.12}$$

$$= \frac{KPs + KQ}{s^2 + KPs + KQ} \tag{12.13}$$

So the natural frequency and damping factor are given by

$$\omega = \sqrt{KQ} \tag{12.14}$$

$$\zeta = \frac{KP}{2\omega} \tag{12.15}$$

We can solve these equations for our controller gains:

$$Q = \frac{\omega^2}{K} \tag{12.16}$$

$$P = \frac{2 * \zeta * \omega}{K} \tag{12.17}$$

Using the parameters from Chapter 5 and choosing $\omega = 3 \times 10^4$ and $\zeta = 1$ we get $P = 0.91$ and $Q = 1.36 \times 10^4$.

The Matlab code for our linearized PI controller is shown in Figure 12.5. This code determines the voltage error `er` and integrates this error to give an integral error `ier`. These two error terms are combined using the proportional and integral gains `P` and `Q` to generate linear control variable `x`. After limiting `x` we take its square root to give duty factor and multiply by `t_c` to give pulse width.

The simulated response of the controller is showin in Figure 12.6.

[More]

```
% linearized PI control law - run once per cycle
er = v_ref - v_c ;
ier = ier + er*t_c ;
x = (er*P + ier*Q) ;
x = max(0, min(x_max, x)) ;
t_a = t_c*sqrt(x) ;
```

Figure 12.5: Matlab implementation of linearized PI controller for a discontinuous mode buck converter.

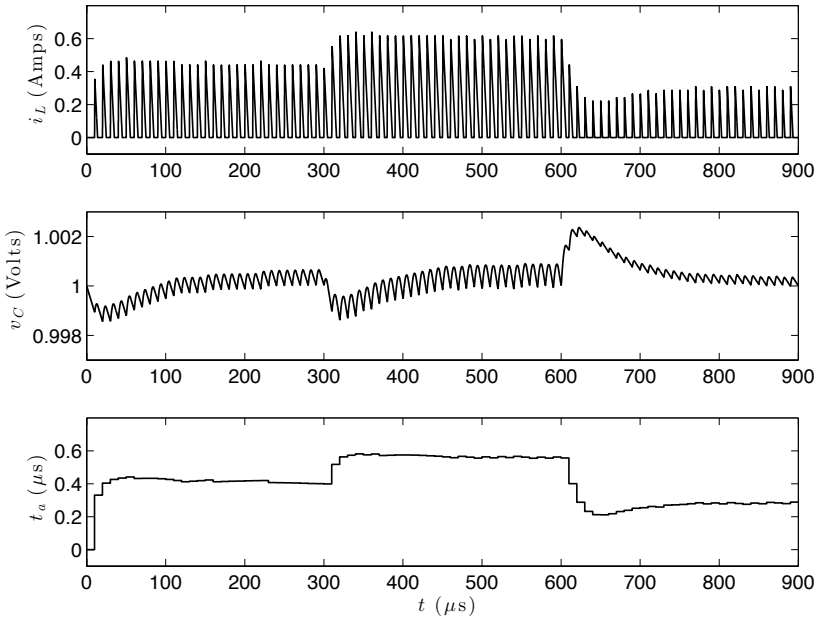

Figure 12.6: Simulated waveforms showing control of a buck converter in discontinuous mode. Waveforms show i_L (top), v_C (middle), and t_a (bottom). i_{Load} is initially 0.1A, steps up to 0.2A at $30\mu s$, and steps down to 0.05A at $60\mu s$.

12.4 Exercises

boost dcm

flyback dcm

inverter that is dcm at bottom of waveform, ccm over the top do a matlab simulation

Chapter 13

Capacitive Converters

Most switched voltage converters use inductors as their primary energy storage element. This is because the voltage at the terminals of an inductor can be switched quickly. An inductor holds its current constant but can have its voltage switched to transfer energy from one voltage to another.

We can also use capacitors to build voltage converters by charging a capacitor at one voltage and then switching its terminals to change a parallel connection to a series connection or vice-versa. With this approach we are limited to rational ratios of input to output voltage. Also, even with ideal switches capacitive converters have losses.

13.1 A Capacitive Voltage Doubler

The circuit shown in Figure 13.1 is a capacitive voltage doubler. Each cycle capacitor C_1 is charged to the input voltage and then *lifted up* so that its top terminal is at twice the input voltage. This *doubled* voltage is stored on output capacitor C_2 where it is drained off by the load before being replenished during the next cycle.

When the two switches are in position **a** input source V_S is connected across capacitor C_1 and capacitor C_1 is charged to V_S. Because real switches have finite resistance R_S this takes a finite amount of time (given by a $\tau = R_S C_1$ time constant) and even if t_a is many τ the voltage on C_1 doesn't quite reach V_S.

When the switches move to position **b** C_1 which is now charged to V_S has its negative terminal connected to the positive terminal of the supply. That puts C_1 and V_S in series. The voltage across the series combination of V_S and C_1 is momentarily $2V_S$.

The upper switch in position **b** connects the series combination of V_S and C_1 in parallel with output capacitor C_2 and current will flow to redistribute charge until $V_S + V_{C1} = V_{C2}$.

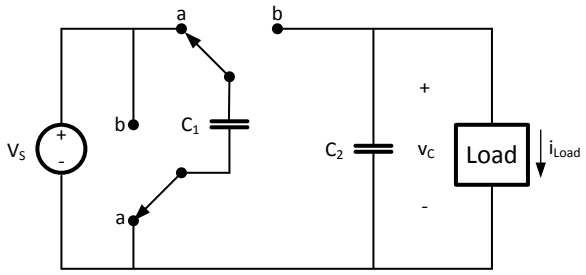

Figure 13.1: A capacitive voltage doubler. When the two switches are in position **a** capacitor C_1 is charged to input voltage V_S. When the switches move to position **b**, the bottom of C_1 is connected to the top of the input source making the top of C_1 momentarily at double the input voltage $2V_S$.

13.2 Analysis of the Capacitive Voltage Doubler

In this section we analyze the capacitive voltage doubler to determine its output voltage and its efficiency. Unlike the inductive converters we have examined up until now - which are perfectly efficient until we start considering component losses (e.g., in swtiches and inductors), capacitive converters have losses even with ideal components (i.e., switches that switch instantly with no resistance and lossless capacitors).

The capacitive converter is best understood by looking at charge transfer. Over the course of a switching cycle $t_{\rm cy}$ charge $Q_{\rm cy} = I_{\rm Load} t_{\rm cy}$ is transferred to the load. All of this charge must be transferred from the input source to C_1 when the switches are in position **a**. Thus, the voltage on C_1 at the instant the switches move to position **a** is

$$\begin{aligned} V_b &= V_S - \frac{Q_{\rm cy}}{C_1} \\ &= V_S - \frac{I_{\rm Load} t_{\rm cy}}{C_1} \end{aligned} \tag{13.1}$$

The energy dissipated in the switch during this transfer is that needed to charge C_1 through $V_S - V_b$.

$$\begin{aligned} E_b &= \frac{1}{2} C_1 \left(V_S - V_b\right)^2 \\ &= \frac{Q_{\text{cy}}^2}{C_1} \\ &= \frac{\left(I_{\text{Load}} t_{\text{cy}}\right)^2}{C_1} \end{aligned} \tag{13.2}$$

While the switch is in position **a** the load continues to discharge C_2. At the end of the cycle, just before switching to position **b** the voltage on C_2 is

$$V_c = V_b - \frac{I_{\text{Load}} t_a}{C_2} \tag{13.3}$$

Where t_a is the time in switch position **a**.

When the switch moves to position **b** capacitor C_2 which is charged to V_c shares charge with C_1 which is charged to V_S. The resulting voltage V_a is given by adding the voltage drop during t_b to V_b.

$$V_a = V_b + \frac{I_{\text{Load}} t_b}{C_1 + C_2} \tag{13.4}$$

Which is also equal to the voltage given by the charge sharing

$$V_a = \frac{C_1 V_S + C_2 V_c}{C_1 + C_2} \tag{13.5}$$

During this charge sharing the energy dissipated in the switch is given by

$$\begin{aligned} E_a &= \frac{C_1 C_2 \left(V_S - V_c\right)^2}{C_1 + C_2} \\ &= \frac{C_1 C_2 \left(\frac{I_{\text{Load}} t_{\text{cy}}}{C_1} + \frac{I_{\text{Load}} t_a}{C_2}\right)^2}{C_1 + C_2} \\ &= \frac{I_{\text{Load}}^2 \left(C_2 t_{\text{cy}} + C_1 t_a\right)^2}{C_1 C_2 \left(C_1 + C_2\right)} \end{aligned} \tag{13.6}$$

So the total energy lost each cycle is

$$\begin{aligned} E_{\text{cy}} &= E_b + E_a \\ &= \frac{\left(I_{\text{Load}} t_{\text{cy}}\right)^2}{C_1} + \frac{I_{\text{Load}}^2 \left(C_2 t_{\text{cy}} + C_1 t_a\right)^2}{C_1 C_2 \left(C_1 + C_2\right)} \\ &= I_{\text{Load}}^2 \left(\frac{t_{\text{cy}}^2}{C_1} + \frac{\left(C_2 t_{\text{cy}} + C_1 t_a\right)^2}{C_1 C_2 \left(C_1 + C_2\right)}\right) \end{aligned} \tag{13.7}$$

So we see that the inherent *energy* loss of the capacitive converter increases quadratically with load current I_{Load} and is inversely proportional to the capacitance. The energy increases quadratically with cycle time t_{cy}; so *power* increases linearly with t_{cy}.

One way to think about this power loss is that its just the load current times the instantaneous voltage drop between V_S and V_2. With ideal circuit elements, if we pick C_1 and C_2 so that average voltage drop is about 10%, the converter will be about 90% efficient[1] This is in contrast to an inductive converter which, with ideal elements, is 100% efficient. Because of this inherent efficiency disadvantage, capacitive converters are rarely used when efficiency matters unless inductors are not available.

13.3 Exercises

Capacitive voltage halver

Capacitive voltage negater

Multiply by 3/2

[1] A real converter, however will have many other losses and will not be anywhere near this efficient.

Chapter 14

The Grid

In this chapter we discuss the power *grid* that is used to connect power generators to end users in much of the developed world. The grid both to distributes power and controls its quality. The distribution involves physically moving energy from the point of generation to the point of use. The control involves adjusting generators, switches, and occasionally loads to maintain constant voltage and frequency. The grid motivates the need for power converters that operate with AC inputs or outputs which are the topics of Chapters 15 and 16.

14.1 Structure of the Grid

Figure 14.1 illustrates the overall structure of the power grid. A number of power plants generate power. The power from these plants is distributed over long distances by transmission lines. Within smaller areas distribution lines transmit the power to individual customers that consume the power. In the US, the entire grid operates at 60Hz AC and is synchronized[1]. In other parts of the world, a 50Hz frequency is used. All points in the system operate at the same frequency and (modulo transmission delays) the same phase.

Historically AC was chosen for power distribution because it can be converted from one voltage to another using just a transformer. Before the availability of modern power electronics, a motor generator pair was needed to convert DC from one voltage to another. With modern power electronics, DC is now an attractive alternative to AC power transmission and HVDC lines are being used for many new projects — despite the need to convert to AC to connect to the rest of the grid.

[Figure of daily power demand]

[1] Southern California originally used a 50Hz system and did not convert to 60Hz until 1948

Figure 14.1: The grid connects power plants to consumers of electricity via a network of transmission and distribution lines.

The main challenge of the grid is to economically generate energy while adapting to variation in load and in some types of generation. Load varies with the time of day and with weather. Peak power is often on a hot summer afternoon when everyone is running their air conditioners. In the early evening lights go on. Late at night, power use is at its daily minimum. Solar and wind generators also vary with weather and time of day. Against this backdrop of variabilty, the most economical power plants operate at near constant output, are slow to change output, and are expensive to turn on and off. Generating plants that can respond rapidly to transients are much more expensive to operate.

14.2 Generation

Power plants come in various shapes and sizes — and with different variablity and response times. The bulk of the energy in the US is generated by large fossil fuel plants that range from 100s of MW to a few GW of capacity. A typical plant burns fossil fuel to generate steam that turns a turbine that in turn turns an electric generator. A large generator may output up to a GW (30kA at 33kV) or more.

Large fossil fuel plants may take hours to start up and come to full power. Also, due to the adverse effects of temperature cycling, it is expensive to turn them on and off. Thus, they cannot respond quickly to changes in load. The large plants are supplemented by small *peaking* plants, often small jet engines or diesel generators, that can be quickly brought on line to handle instantaneous demand. A typical system also keeps a certain amount of *spinning reserve* — power plants that are warmed up and spinning synchronous with the grid but have reserves of power that can be accessed quickly.

Wind farms may be up to several GW in capacity with individual turbines being a few MW each. In 2015 4.7% of the energy in the US was generated by wind. Utility-scale solar plants are typically a few MW in capacity. In 2015 only 0.6% of the energy in the US was generated by solar (EIA). Wind and solar generate clean energy but their output may vary due to weather (and daylight for solar). Thus, they must be backed up by either energy storage or by fossil fuel plants.

14.3 Transmission and Distribution

Power is transmitted by very high voltage three-phase transmission lines. To reduce resistive losses, the output of a power plant is converted from an intermediate voltage (10-40kV) to a very high voltage (100-500kV) for transmission. The 30kA output of a 1GW generator at 33kV drops to 2kA at 500kV — making it much easier to transmit. The I^2R resistive power losses drop quadratically as voltage is increased.

Close to the point of use, lower-voltage (10-30kV) distribution lines are used to distribute power to customers. Power in your neighborhood is most likely

distributed in this voltage range on poles and/or underground lines. A *substation* houses the transformers used to convert from the high-voltage transmission voltage to the lower-voltage distribution voltage.

Small pole-mounted or underground transformers are used to convert from the distribution voltage to the customer voltage. For US residential customers these are single-phase transformers that are fed from two-phases of the three-phase distribution voltage and generate center-tapped 240V AC for a small group of customers. 120V applicances are connected between the center tap (AC neutral) and one of the two lines (AC hot).

AC voltage conversion in the grid is performed by transformers. These transformers are identical to the ones described in Chapter 9. However, in most power systems, three-phase transformers are used. Because these transformers operate at 60Hz, they are very large and typically have steel cores. To first approximation, the volume of a transformer is inversely proportional to frequency so to handle the same amount of power, a 60Hz transformer must be 1,000× the volume of a 60kHz transformer in a switching power converter.

If the load applied to the grid is capacitive or inductive, the voltage and current in the grid will be out of phase. With such a load, the *power factor* (Chapter 16), the ratio of the power transmitted to the product of the RMS voltage and current, will be less than one. For sinusoidal voltage and current we have:

$$PF = \frac{P}{V_{\text{rms}} I_{\text{rms}}} \tag{14.1}$$

$$= \cos(\theta). \tag{14.2}$$

Where θ is the phase difference between the voltage and the current.

A low power factor is bad because the *imaginary power* causes resistive losses in the transmission and distribution lines without delivering any power to customers. To improve power factor in the presence of inductive loads — which are characteristic of motors used in industrial machinery — many utilities switch capacitor banks onto the line.

14.4 Control of the Grid

The grid is primarily controlled by modulating the generators at each power station to maintain voltage and frequency. When the load increases, the increased current translates to increased torque on the shaft of a generator. This torque causes the generator to slow — dropping its output frequency. The generator controller reacts by increasing its power to counteract the torque and maintain a constant frequency. The generator can also control its output voltage — to maintain a constant voltage in the presence of load variations.

Grid transients may occur due to generation or transmission problems in addition to load variation. For example a generator may unexpectedly go off line due to a failure, or a transmission line supplying power from a distant

generator may go down. Other generators and transmission facilities must react to carry the power removed by the failure. When the generator is operating at full capacity and a transient occurs requiring additional power additional generating resources must be brought on line quickly to maintain frequency and voltage. The control problem is also complicated because many generators may respond simultaneously to the same transient. Depending on how this response is coordinated the most economical generator may not be the one that picks up the incremental load.

In some cases the grid controller may respond to transients by reducing load. Most often this is done via an agreement with the consumer where, in exchange for a lower rate, they agree to allow their power to be cut off at certain times. Sometimes, this is done by dropping whole sections of the grid off line to avoid having the entire grid fail.

spot prices

14.5 Future of the Grid

problems caused by variability

problems caused by geographic

distant wind and solar

distributed solar

microgrids and nanogrids - connected, can island - neigborhood or city scale

14.6 Exercises

If you could design a *grid* from scratch using modern components what would it look like. What frequencies and voltages would you use for transmission and distribution? Justify your choices.

Chapter 15

Inverters

Up to now we have dealt with voltage converters that take a DC input voltage and produce a DC output voltage. In many cases, however, we have an AC input voltage, and AC output voltage, or both. AC is widely used for power transmission because it allows voltages to be converted using just a transformer.

It turns out that the same converters we have been using to handle DC voltages can also be used to handle AC voltages if we simply consider an AC voltage to be a *slowly*-varying DC voltage. For most power distribution which uses 50 or 60 Hz AC this is a safe assumption for a converter that operates at 100s (or even 10s) of kHz. For a 100kHz converter, for example, a 60Hz sine wave traverses only 0.216 degrees (0.0038 radians) during one switching cycle. At the steepest point of the sine wave the voltage varies by only 0.19% of the peak-to-peak value between cycles. From the perspective of a modern switching converter, AC really is just slowly-varying DC.

In this chapter we describe inverters, voltage converters that produce an AC output. Inverters are used both to drive AC power distribution lines and to generate waveforms for driving AC induction and brushless permanent-magnet motors Chapter **??**. The topology we describe is similar to the buck converter of Chapter 2 except that we modify the converter so it can produce negative as well as positive outputs because our slowly varying DC output takes on negative as well as positive values.

In Chapter 16 we introduce the concept of *power factor* and describe how to build a converter with an AC input that maintains unity power factor — i.e., over an AC cycle it draws an input current that is proportional to the input voltage.

15.1 Basic Inverter

Figure 15.1 shows a basic single-phase inverter that takes a DC input voltage V_1 and generates an AC output voltage V_2. The circuit is basically a buck converter (Figure 12.1) with a full-bridge of switches between the input source

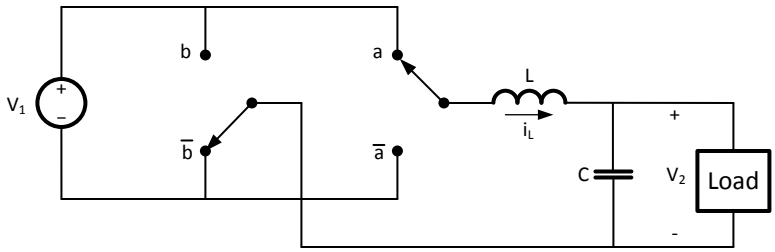

Figure 15.1: A single-phase inverter takes power from a DC source V_1 and generates an AC voltage V_2 into a load. A full-bridge of switches **a** and **b** select the polarity of the output and pulse-width-modulate the input voltage into the output LC filter.

and the inductor to reverse the polarity when needed. The bridge is similar to the circuit used in the full-bridge converter (Figure 9.4) just drawn differently.

During the positive half of the output cycle, switch **b** remains in the $\overline{\mathbf{b}}$ position — connecting the bottom terminal of the load to ground (the negative terminal of V_1) — and switch **a** switches each cycle to generate a pulse width proportional to the instantaneous desired output voltage, just like the buck converter. During the negative half of the output cycle, switch **a** remains in the $\overline{\mathbf{a}}$ position — connecting the top terminal of the load to ground and switch **b** switches to generate the desired negative voltage across the load.

To first approximation, to generate a sine wave $\sin(\omega t)$, each switching cycle we generate a pulse with width $t_p = t_{\rm cy}|\sin(\omega t)|$. If $\sin(\omega t) > 0$ we apply the pulse to switch **a**, holding **b** in the $\overline{\mathbf{b}}$ position. Otherwise we appy the pulse to switch **b** holding switch **a** in the $\overline{\mathbf{a}}$ postion.

Because the inverter, like the single-polarity buck converter of Chapter 2 is inherently unstable, control of the inverter is not quite this simple. Transients can excite oscillitory modes of the LC output filter. We require a controller, such as the PD controller described in Chapter 5 to damp these oscillations.

The inverter of Figure 15.1 is not limited to generating sine waves. It can generate arbitrary band-limited wave forms. For an arbitrary function $f(t)$, each cycle we apply $t_p = t_{\rm cy}|f(t)|$ to the appropriate switch.

15.2 Inverter Simulation

Figure 15.2 shows a simulation of an inverter generating a 1V 60Hz sine wave into a 1Ω resistor load from a 1V DC source. The switching frequency of the inverter is $f_{\rm cy} = 60kHz$, $L = 10\mu$H, and $C = 630\mu$F. A zoomed version of this simulation is shown in Figure 15.3.

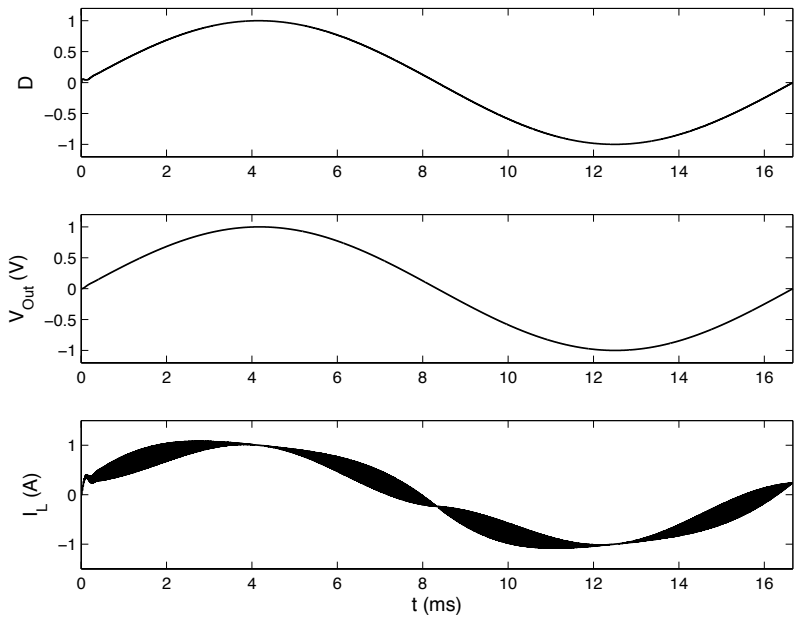

Figure 15.2: Simulation waveforms for an inverter generating a 60Hz sine wave with a 60kHz PWM frequency driving a 1Ω load. From top to bottom the traces show: duty factor, capacitor voltage, and inductor current.

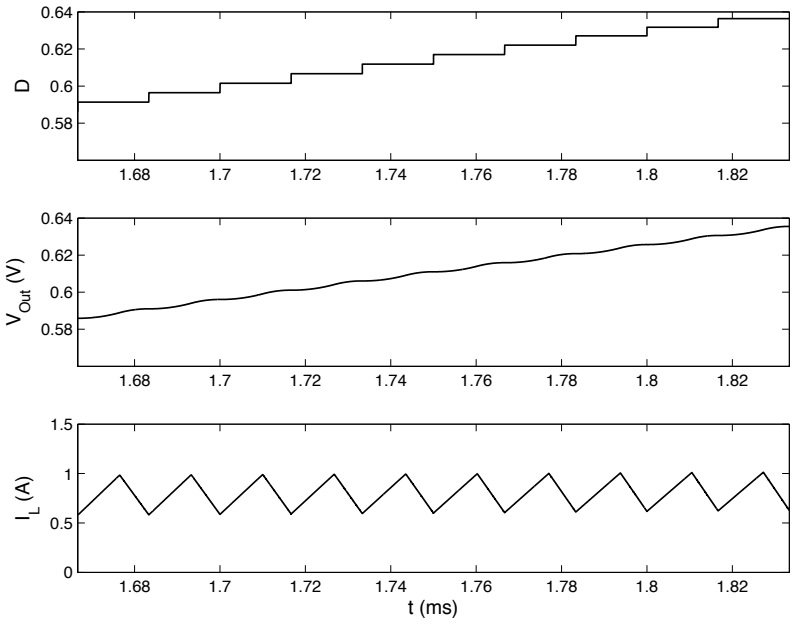

Figure 15.3: Zoomed view of Figure 15.2 showing the region from 1.66ms to 1.84ms.

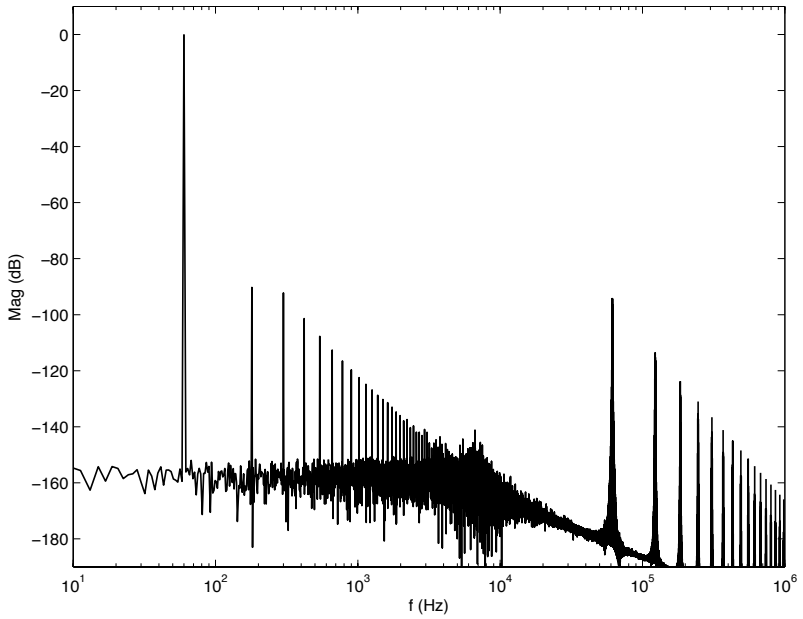

Figure 15.4: Spectrum of inverter output shows dominant peak at 60Hz. The peak at the switching freqeuncy is attenuated by the LC filter to -94dB.

The duty factor, shown in the top trace of each figure tracks the sine wave. Note, however, that this duty factor is not generated open loop, but is rather the output of a PD controller where $\sin(\omega t)$ is the reference input to the controller. The center trace of each figure shows the voltage waveform accurately traces the desired sine wave to about 1% accuracy. There is a very small voltage ripple at the 60kHz switching frequency. The bottom trace shows that the inductor current slightly leads the capacitor voltage — in order to charge and discharge the capacitor. It has a ripple that peaks at nearly 0.4A but is almost zero at the voltage peak - where the duty factor is 1.

In applications where an inverter generates a sine wave we one measure of the quality of the inverter is its spectral purity. Figure 15.4 shows the spectrum out of an inverter. The figure was generated by running the inverter simulation for 32 cycles and taking an FFT of the resulting waveform across C. We have increased the value of C to 6.7mF to give an LC filter with a cutoff frequency of $f_c \approx 600$Hz. We set the 0dB reference at the magnitude of this 1V tone. The noise floor at -155dB is determined by the length of the simulation (and the FFT) and does not reflect any property of the inverter. The corner where this noise floor intersects the roll-off of the 60Hz harmonics does *not* reflect the cutoff frequency of the filter.

There are two sets of lines of interest in this spectrum, the harmonics of the 60Hz output tone, and the harmonics of the 60kHz switching frequency. The 60Hz tone dominates the spectrum. Its first two odd harmonics (the third harmonic at 180Hz and the fifth at 300Hz) are about 90dB down. These harmonics could be reduced by tighter regulation of the output voltage. The seventh and higher odd harmonics are attenuated by the LC output filter.

The line at the 60kHz switching frequency is attenuated 80dB by the filter so it has a magnitude of -94dB. Both even and odd harmonics of the fundamental appear in this group of lines and this chain of harmonics rolls off at the 40dB per decade of the LC filter.

In this chapter we have considered an inverter that is operating as a synthesizer, generating its own sine wave independently. In some cases we wish to build *grid-tied* inverters that generate an output waveform synchronized with the AC grid. To build a grid-tied inverter we regulate not the output voltage — that is determined by the grid — but rather the output current. Such inverters generate a current that is proportional to the instantaneous grid voltage. This is an example of power-factor control which is the topic of the next chapter.

15.3 Exercises

matlab simulation of an inverter

inverter that generates a triangle wave

derive PD controller parameters for inverter with L = 10μH and C=2000μF.

Design output filter to attenuate lines further

Chapter 16

Power Factor Correction

When connecting a device to an AC line it is often required that the device current be proportional to the line voltage. This is true whether the device is driving current into the AC line — e.g., a grid-tied inverter — or drawing current from the line — e.g., an off-line voltage regulator. The requirement for current-voltage proportionality is often described in terms of *power factor* — the ratio of real to apparent power. Intuitively, power factor is the fraction of current that is at the same frequency and in phase with the voltage waveform.

Power factor can vary from unity for two reasons: phase shift and harmonic distortion. Historically, low power factor was due to phase differences between the voltage and current waveforms due to inductive or capacitive loads. For modern electronics the problem is more often due to harmonic distortion — current waveforms with frequency content well above the 60Hz fundamental frequency.

To manage power factor, we again view the AC line as a slowly varying DC source (or load) and regulate our current out of (or in to) the line to be proportional to the voltage. In many power converters this proportional current control is performed in a separate *power factor correction* input stage. For a grid-connected inverter, the inverter is controlled to provide a proportional output current.

In some rare cases we are required to control current in a manner that generates out-of-phase current and voltage. For example, utility-scale inverters are often required to generate either an inductive (lagging) or capacitive (leading) phase shift, purposely generating a power factor less than unity. Providing such phase shifts requires handling currents flowing opposite the normal direction during some fraction of the cycle.

16.1 Power Factor

Power factor is defined as the ratio of real power to apparent power averaged over an integral number of AC cycles. Apparent power here is the RMS value of

current multiplied by the RMS value of voltage. If the instantanous voltage and current at time t are $v(t)$ and $i(t)$ respectively we can calculate power factor as:

$$\begin{aligned} PF &= \frac{P}{V_{\text{rms}} I_{\text{rms}}} \\ &= \frac{\int v(t) i(t) dt}{\sqrt{\int v^2(t) dt} \sqrt{\int i^2(t) dt}} \end{aligned} \tag{16.1}$$

If an AC voltage source drives a load composed of passive linear circuit elements (resistors, capacitors, and inductors), then the current waveform will be a sine wave at the line frequency and the power factor is determined by the phase angle θ between the voltage wave and the current wave, $PF = \cos(\theta)$.

For example, if our load is an 2.7mF capacitor in parallel with a 1Ω resistor, then the admittance of the load is $Y = 1/R + j\omega C$, then for a 60Hz input source ($\omega = 2\pi \times 60 = 377$ rad/s) our admittance is $Y = 1 + j\eth$. which has a magniude of $\sqrt{2}$ and an angle $\theta = pi/4$. In this case the power factor is $PF = 2^{-1/2} = .707$.

For switching power converters we are usually less concerned about phase shift in the line-frequency component of the current than we are with higher-frequency components of current. For example, some older power supplies rectified the AC input voltage and then connected the rectifier output to a filter capacitor. Such power supplies draw current only during the part of the AC cycle when the input waveform is above the capacitor voltage. During the rest of the cycle the input current is zero.

Consider the simplified current waveform shown in the top part of Figure 16.1. The current is $\pm$1A during 20% of the cycle and 0 during hte remainder of the cycle. The sinusoidal voltage waveform with RMS value of 1 is shown (dotted) for comparison. The power spectrum of the current waveform is shown in the bottom panel of Figure 16.1. The spectrum shows significant power in the odd harmonics of the current signal.

To compute the power factor of this signal (with the phase relationship to the voltage signal as shown), we numerically perform the integrals in (16.1) and compute

$$\begin{aligned} P &= 0.28\text{W} \\ I_{\text{rms}} &= 0.45\text{A} \\ V_{\text{rms}} &= 1.0\text{V} \\ PF &= 0.62 \end{aligned} \tag{16.2}$$

Such a low power factor would be unacceptable in most power converters. Most regulations require a power factor very close to one. Some regulations explicitly limit the power at each harmonic of 60Hz.

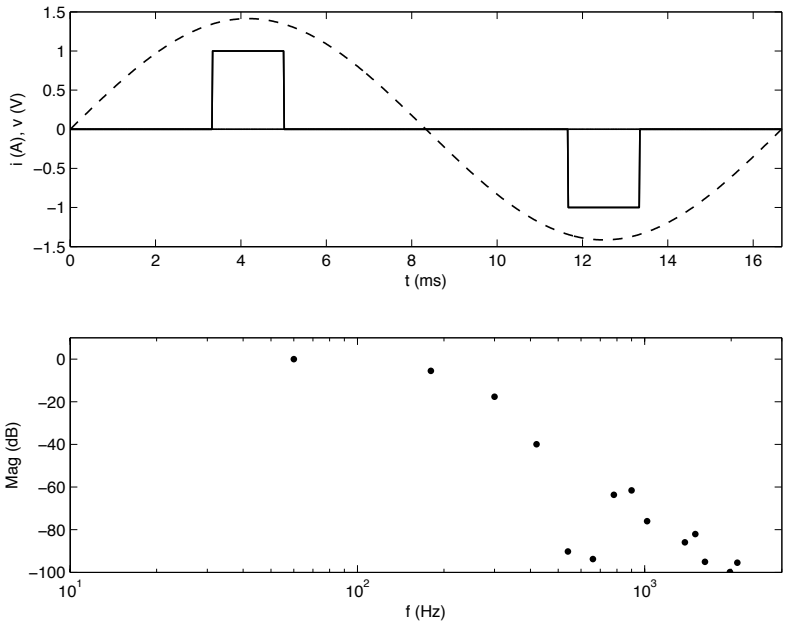

Figure 16.1: Example non-sinusoidal current waveform is ±1A during 20% of the cycle and zero during the remainder of the cycle. The top trace shows the voltage (dotted) and current (solid) waveforms. The bottom trace shows the power spectrum of the current waveform normalized to its fundamental magnitude.

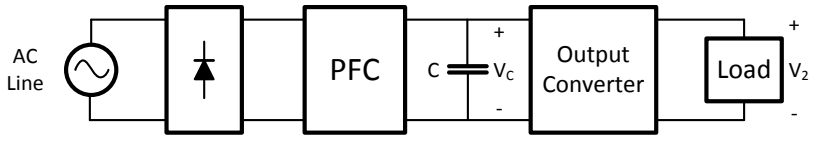

Figure 16.2: A two-stage PFC corrected converter. A PFC input stage converts a rectified AC input to an intermediate DC voltage V_C while drawing sinusoidal input current to maintain unity power factor. A second converter converts from V_C to the output voltage V_2 without the constraint of sinusoidal input current.

16.2 Power Factor Correction

Many voltage converters use a power-factor-correction (PFC) input stage to tailor their input current profile to match the AC voltage waveform. As shown in Figure 16.2 the converter consists of two stages. A bridge rectifier — denoted with a diode symbol — converts the input AC waveform to rectified AC, $v = |A\sin(\omega t)|$. The input PFC stage takes this rectified AC (120V or 240V 50 or 60Hz) input and generates an intermediate DC voltage V_C (often 400V DC) on a storage capacitor C. This intermediate node is often called a *DC link*. The input PFC stage regulates its input current to follow a sinusoidal profile to maintain unity power factor.

The output stage converts the intermediate voltage V_C to the output voltage (AC or DC) using one of the converter topologies previously discussed. The output converter almost always uses an isolated converter to isolate both output terminals from the AC line.

The storage capacitor C stores energy during the high periods of the AC cycle to provide energy to the output stage during the low periods of the cycle. The intermediate node can tolerate a fairly high voltage ripple on V_C compared to the ripple allowed on output voltage V_2. Thus, a smaller capacitor can be used than if this energy were stored on the output filter capacitor.

The voltage and current into the PFC stage are of the form

$$v = A_v \sin \omega t$$
$$i = A_i \sin \omega t$$

Thus, the power into the PFC stage is

$$\begin{aligned} p &= A_p \sin^2 \omega t \\ &= \frac{A_p}{2}\left(1 + \sin 2\omega t\right) \end{aligned} \tag{16.3}$$

Figure 16.3 illustrates the energy that the DC-link capacitor C must store. The top panel shows the waveform for the voltage and current. The bottom panel shows the instantaneous power into the DC-link (solid line) and out of the DC-link (dotted line). During the period from $\omega t = \pi/4$ to $3\pi/4$ the capacitor accumulates energy that is $2/\pi$ (0.637) of the energy out of the DC-link during the same period. This is the ratio of the area under a cosine wave from 0 to π to the area of a rectangular box bounding the cosine during the same period.

As an example, consider a PFC-corrected supply that provides 1kW of power and must maintain the DC-link voltage between 350V and 500V while operating off a 60Hz AC supply. During the 4.17ms of the quarter sine wave from $\pi/4$ to $3\pi/4$ the DC-link outputs 4.17 J of energy and the capacitor must store $(2/\pi)(4.17) = 2.65$ J of energy. We can solve for the required capacitance by writing the expression for the change in capacitor energy:

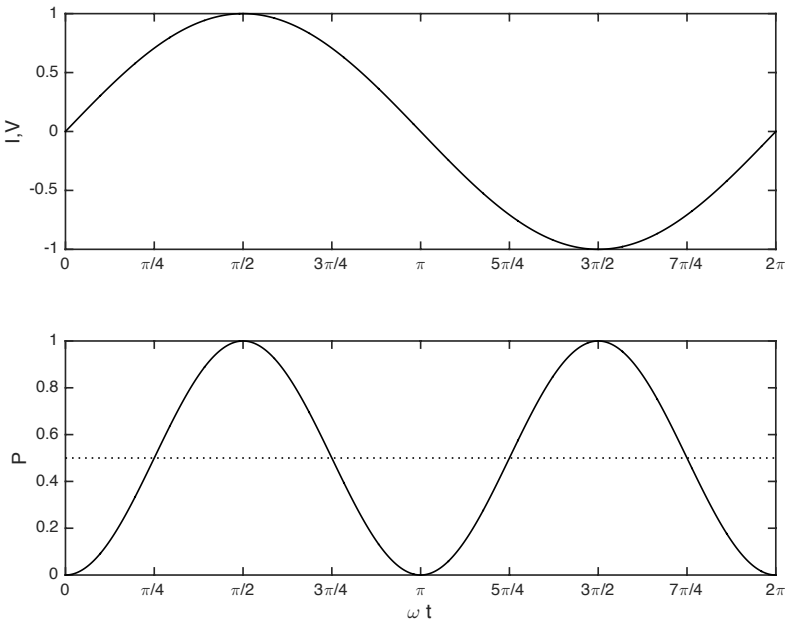

Figure 16.3: Waveforms for power-factor correction. Voltage and current into the converter are shown in the top panel. Power in to (solid line) and out of (dotted line) the converter are shown in the bottom panel. The power-factor corrector must store the energy between the two curves in the bottom panel from $\pi/4$ to $3\pi/4$ and return this energy between $3\pi/4$ and $5\pi/4$.

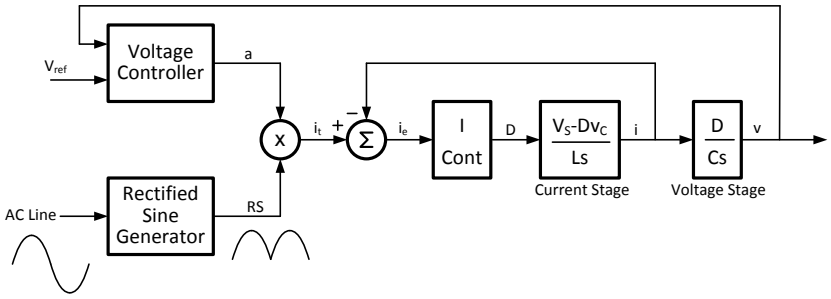

Figure 16.4: Block diagram of a PFC controller using nested voltage and current control similar to Figure 7.1. The current is regulated to follow a rectified sine wave synchronized to the AC line. The rectified sine wave is scaled by the voltage control loop to maintain a desired voltage level v_{ref} on V_c. The voltage loop must be operated slowly enough to avoid distorting the sine wave.

$$\begin{aligned}
\Delta E &= \frac{1}{2}C\left(V_1^2 - V_2^2\right) = 2.65\text{J} \\
C &= \frac{2\Delta E}{(V_1^2 - V_2^2)} \\
&= \frac{2(2.65)}{500^2 - 350^2} \\
&= 41.6\mu\text{F}
\end{aligned} \tag{16.4}$$

Even though we are swinging the capacitor through less than a third (30%) of its voltage range (150V out of 500V), we are able to make use of just over half (51%) of the peak energy storage in the capacitor because of the square-law relationship between voltage and energy.

The PFC stage of Figure 16.2 is often implemented using a boost converter (Section 6.1). An input filter is added to isolate the power line from switching-frequency noise generated in the boost. The boost converter is well suited to take the rectified AC input voltage, which varies slowly between 0 and 340V (for 240V AC), and boost it to an intermediate voltage higher than the peak AC voltage.

The PFC input stage is controlled with two loops as shown in Figure 16.4. The inner loop is a current control loop (Chapter 7) that takes as input a target current and modulates D the duty factor of switch **a** to drive the average inductor current to equal the target current. The duty factor D varies widely over the cycle as the input voltage changes.

The outer control loop sets the target current to be equal to a rectified sine wave synchronized to the AC line multiplied by a demand factor a. This causes

the current waveform to be sinusoidal — giving near unity power factor — with the amplitude of the sine wave being set to satisfy the power demand of the output stage.

The demand signal is derived from the intermediate voltage, for example by using a PID controller to drive the intermediate voltage to a target value. It is important that the demand signal not track the normal variation in the intermediate voltage over the AC cycle or harmonics will be introduced into the current waveform — degrading power factor. This is accomplished by making the bandwidth of this controller substantially lower than 60Hz. One approach, that responds to transients faster than a PID controller, is to set the demand based on the average power drawn over the last AC half cycle.

The rectified sine wave can be generated by simply dividing down a rectfied version of the AC line voltage. However, a more accurate approach is to synchornize a sine-wave generator to zero crossings of the AC line. This avoids having the current waveform repeat the harmonic distortion of the voltage waveform.

16.3 Simulation of a PFC Input Stage

Figure 16.5 shows simulated waveforms for a PFC input stage using the controller shown in Figure 16.4. The input is a rectified 240V AC signal. The output V_C is loaded with an 800Ω resistor that is switched to a 1.6kΩ resistor at 33.3ms. The boost converter operates at 100kHz with L=2mH and C=50μF. The current controller is a PI control loop as described in Chapter 7 with $\zeta = 1$ and $\omega = 2\pi \times 10^4$ (10kHz). The voltage controller is also a PI controller with $\zeta = 1$ and $\omega = 75$ (12Hz). The multiplier signal a out of the voltage controller is only allowed to change at zero crossings of the AC signal.

The controller is initialized with $a = 1$ and during the first half cycle the current is a half-sine wave witha peak of 1A. This is not sufficient to meet the current demand and V_C falls from its target value of 450V. The voltage controller responds by setting $a = 1.56$ for the second half cycle. This stabilizes the fall of V_C and the integral control loop ultimately returns it to 450V at which point $a = 1.47$.

After the fourth half cycle, at t =33.3ms, the load resistance is doubled. This causes V_C to rise during the fifth half cycle as the multiplier a remains at 1.47. The voltage controller responds in half-cycle six by setting $a = 0.696$ and the integral controller returns V_C to 450V with the multiplier settling at $a = 0.725$.

Figure 16.6 shows a transient in the opposite direction. At 33.3ms the output load resistor is halved from 800Ω to 400Ω . Immediately following the transient the voltage on V_C drops dangerously low. If this voltage drops below 350V, the output converter may not be able to meet its specifications. To handle extreme load transients, it may be necessary to update the multiplier value a in the middle of a half cycle. This will distort the current waveform for that half-cycle but that is usually preferable to dropping the output voltage below

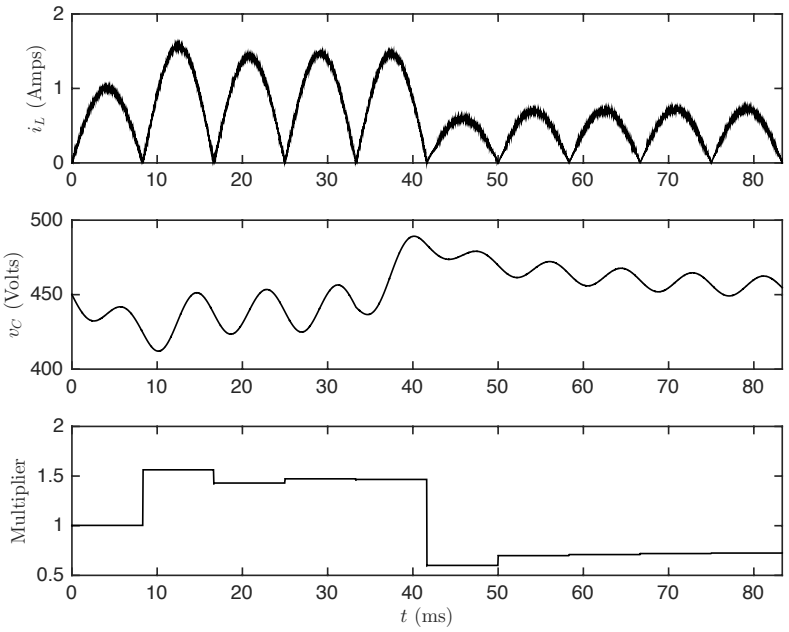

Figure 16.5: Waveforms from a Matlab simulation of a PFC input stage. The top panel shows the inductor current i_l after smoothing. The middle panel shows the voltage V_C. The bottom panel shows the multipler a produced by the voltage regulator (Figure 16.4) to scale the target current i_t sine wave. The load resistor is doubled at 33.3ms.

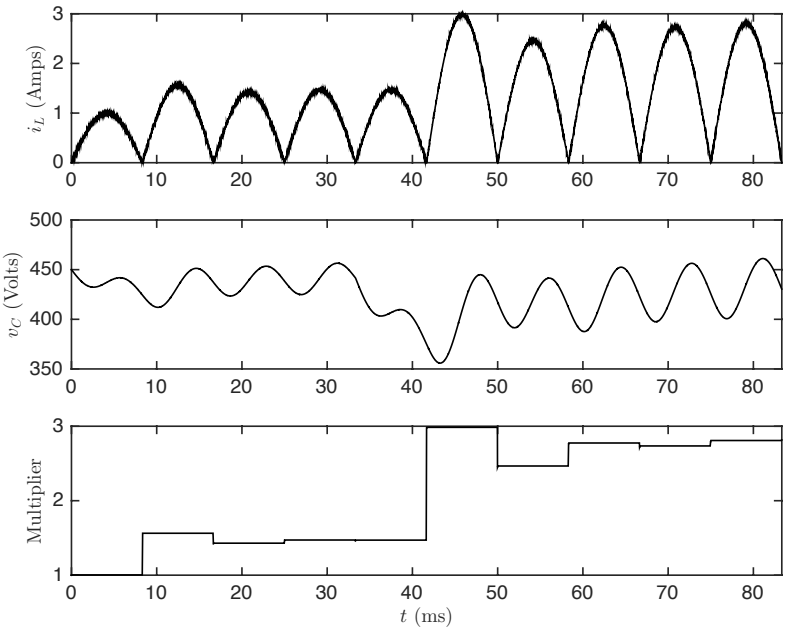

Figure 16.6: Waveforms from a Matlab simulation of a PFC input stage where the load resistor is halved at 33.3ms.

specification. Another alternative is to use a larger storage capacitor.

16.4 Grid-Tied Inverters

A grid-tied inverter is a special-case of a PFC correction circuit where the power (and the current) flows from the converter to the AC line. However, other than reversing the direction of power flow, operation is the same as a PFC circuit on a power supply input (Figure 16.2). The circuit topology is that of Figure 15.1. Rather than regulating the voltage to be sinusoidal as in Chapter 15, the current out of the grid-tied inverter is regulated to be proportional to a sine-wave synchronized with the line frequency. The magnitude of the sine wave is controlled to match the power being input to the inverter.

16.5 Three-Phase Systems

To achieve unity power factor in a three-phase AC system, the current in each of the three phases must be regulated to be proportional to a sine wave synchronized to that phase. The advantage of three-phase systems is that very little energy storage is required after the PFC input stage. This is because the net power delivered from all three phases is constant over the AC cycle. Thus, there is no need to store energy during the high part of the AC cycle and return it during the low part of the cycle.

16.6 Exercises

calculate power factor of a few linear loads

calculate power factor of some waveforms - e.g., square and triangle wave

like Figure 16.1 but 10% duty factor

30% duty factor

offset pulse from $\pi/2$ to $2pi/3$

develop a matlab simulation, including controller, of a grid tied inverter. Assume that the input is a solar panel.

Part II. Components

Semiconductor Switches
- FETS, IGBTs, and Diodes
- DC characterisitcs
- AC characteristics
- Charge storage

Controlling switches
- Switches in pairs
- Switch drivers
- High-side drivers
- Dead-time control

Magnetic components
- Magnetic circuits
- Field H and B, Flux, Reluctance, and permeability
- Inductance
- Magnetic materials - steel, ferrite, air
- Losses

Designing Inductors
- Basic design procedure
- Optimization

Designing Transformers
- Magnetizing and leakage inductance
- Design procedure
- Optimization

Batteries
- Basics
- Simple model
- Charge/discharge curves
- Temperature effects
- Battery management
- Battery charging

Capacitors
- Types of capacitors
- Impedance Spectrum
- Models

Chapter 17

Switches

Up to this point we have analyzed green electronic systems with ideal switches that have zero resistance when on, can block an unlimited voltage when off, and switch from on to off, and vice versa, in zero time. Unfortunately such ideal switches are not available to build real systems. The transistors and diodes we use for actual switches have a number of non-ideal properties including non-zero on resistance, finite breakdown voltage, and non-zero switching time. These non-ideal properties lead to conduction losses — due to on-resistance — and switching losses — due to finite switching time — in real green electronic systems.

17.1 DC Characteristics of Real Switches

Figure 17.1 shows the I-V characteristics of four real switches. Each curve shows the current that flows through the corresponding switch (y-axis) as a function of the voltage applied across the switch (x-axis). The I-V curve for an ideal switch is a vertical line up the y-axis.

The two MOSFETs[1] have a linear I-V characteristic that can be described as an *on resistance* R_{on}. The 60V FET has a very low R_{on} of 1.5mΩ. Unfortunatley to get such a low on-resistance requires bulding the FET in a manner that causes it to break down when voltages greater than 60V are applied across its terminals. For FETs of a given cost, R_{on} increases quadraticaly with voltage. The 600V FET has 10x the breakdown voltage of the 60V FET but pays for this higher breakdown voltage with a 100mΩ R_{on}. For low voltage applications, below about 400V, the low R_{on} of MOSFETs make them the switch of choice.

The diode and IGBT have non-linear I-V curves. For these devices, very little current flows until a forward voltage of about 0.7V is reached. Then the current increases rapidly with increasing voltage. To first approximation we can model an on diode or IGBT as a voltage source and a resistor — approximating

[1] MOSFET is an acronym for metal-oxide-semiconductor field-effect transistor, we often refer to them as just FETs.

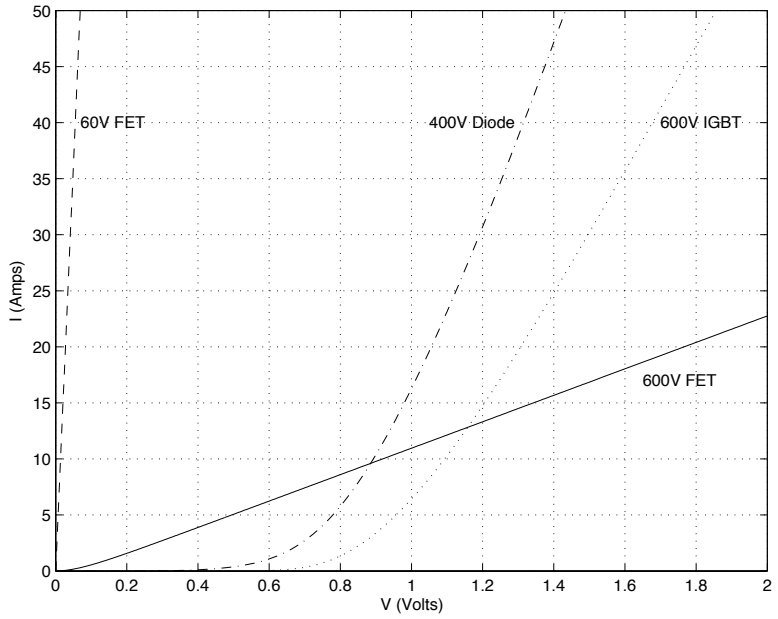

Figure 17.1: I-V Characteristics of four real switches. From right to left along the 20A line: a 60V MOSFET (IRLB 3036), a 400V diode (STTH20R04), a 600V IGBT (FGH40N60), and a 600V MOSFET (FCB36N60N). The two MOSFETs have a linear I-V characteristic than can be characterized by an on-resistance. The diode and the IGBT have a non-linear I-V characteristic than can be approximated by a voltage source and a resistance.

the I-V curve with two lines. The IGBT in Figure 17.1 for example can be modeled as an 0.9V source in series with a 20mΩ resistor. The low incremental resistance of IGBTs relative to MOSFETs make them attractive for high-voltage — above about 400V — applications. At 600V for example, our example IGBT has lower drop — and hence lower conduction losses for currents above 13A.

The fixed diode drop of an IGBT or diode make them inappropraite for use at very low-voltages where such a drop would cause excessive losses. Dropping 0.9V in a 1V power supply would result in 47% loss. Even in a 10V supply the 8% loss would be considered unacceptable. Such supplies use *synchronous rectification* to replace diodes with MOSFETs that have a much lower voltage drop.

17.2 AC Characteristics of FETs and IGBTs

While Figure 17.1 shows the static characteristics of real switches Figure 17.2 shows the dynamic characteristics. This figure shows the non-zero switching time of these real devices and its effect on switching losses.

The top panel of this figure shows a 600V MOSFET switching 20A through 500V in the boost configuration shown in Figure 17.3. At the start of time the MOSFET is off. At 10ns the gate of the MOSFET is driven to 15V with a 30ns rise time and a 2.2-Ohm gate-drive resistor. When the rising gate reaches the threshold voltage V_T of the MOSFET, it turns on very quickly with the current rising to 20A in about 1ns. The voltage then falls through 500V over about 10ns — limited by the parasitic drain capacitance of the FET.

The bottom panel shows the power dissipation of the MOSFET — the product of voltage across the device with current through the device. At the point where the current reaches 20A but the voltage is still 500V the FET is dissipating 10kW. Fortunately this enormous power dissipation lasts only a brief period of time and the total energy of this switching event is just 50μJ.

When the MOSFET turns off (starting at 160μs) the voltage rises to 500V over about 10ns. Only when the voltage reaches 500V can the diode in the boost converter start conducting causing the current to abruptly fall to zero. The bottom panel shows the power dissipated during this transient. The total turn-off switching energy is 52μJ.

The total switching loss over one cycle of operation for the boost converter is 102μJ — the sum of the 50μJ turn-on energy and the 52μJ turn-off energy. If this converter were running at 100kHz, the switching loss would be $P_{\rm sw} = fE_{\rm sw} = 10.2$W.

The waveforms and numbers in Figure 17.2 are unusually optimistic because we are neglecting parasitic inductance in the circuit which slows the current transient and also be cause we are assuming an ideal diode. Most real diodes have reverse recovery characteristics, or at least substantial capacitance, that causes the current to overshoot for a considerable period before the voltage starts to fall.

The middle panel of Figure 17.2 shows the AC characteristics of an IGBT

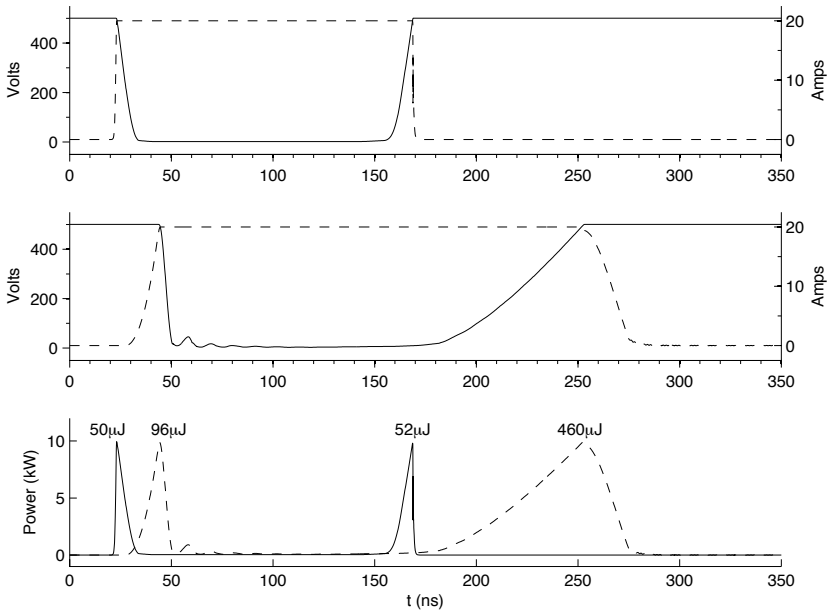

Figure 17.2: Switching characteristics of a FET and an IGBT switching 20A through 500V in a boost configuration. The top panel shows current (dashed line and right scale) and voltage (solid line and left scale) of a 600V MOSFET turning on and then turning off. The middle panel shows the same waveforms for a 600V IGBT. The bottom panel shows the instantaneous power dissipated by the FET (solid line) and IGBT (dashed line) with the energy (the area under the curve) shown for each switching event.

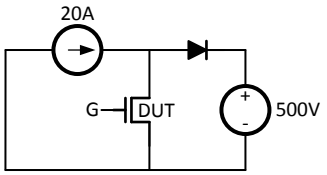

Figure 17.3: Test configuration for the MOSFET and IGBT switching tests shown in Figure 17.2. The device under test (DUT) is connected in a boost configuration with a 20A current source and a nearly ideal diode connected to 500V.

we see that the IGBT is slower than the FET for both turn-on and turn-off. Its turn-on time is about twice that of the FET — hence its turn-on switching loss of 96μJ is about double the FET's turn-on loss. The real issue with IGBTs is their turn-off time. The IGBT turn off time is an order of magnitude longer than that of the FET and its 460μJ turn-off loss is correspondingly higher.

Slow turn off is characteristic of bipolar transistors in saturation. The power handling part of an IGBT is a bipolar transistor and time is required to clear the minority carriers from the collector before it will turn off. IGBTs come in various speed grades that trade voltage drop against turn off time. For slow switching applications IGBTs are constructed with very low forward voltage drop — by allowing the bipolar transistor to go deep into saturation. They pay for the low voltage drop with very long turn-off times - often more than 1μs. For higher-frequency applications IGBTs are constructed with a higher forward voltage drop to keep the bipolar transistor out of saturation and they switch considerably faster as a result.

Chapter 18

Metal-Oxide Field Effect Transistors (MOSFETs)

Power MOSFETs are the workhorse switches for power electronic systems that operate at 200V or below. They switch on and off quickly and have a low, resistive, voltage drop when on. At voltages above 400V, IGBTs are usually preferred over MOSFETs because for a given cost they can handle more current.

To withstand high forward voltages when on, power MOSFETs are typically realized using a vertical DMOS structure. Inherent in this structure is a *body diode* that conducts when the MOSFET is reverse biased. This makes the MOSFET a *one-quadrant switch*. It only controls current in the first quadrant of the V-I plane. It is always on in the third quadrant. If a four-quadrant switch is needed, two MOSFETs can be connected back-to-back.

The avalanche breakdown characteristic of the body diode acts to limit the reverse voltage across the MOSFET during transients. This device, for example, can absorb the energy stored in the leakage inductance in a simple flyback converter (Figure 10.1). As long as the energy and power absorbed during avalanche breakdown stay within limits the device can be used in this manner with no adverse effects.

18.1 Simple Model

Figure 18.1(a) shows the schematic symbol for a power MOSFET. The body diode is explicitly shown because it is an important circuit element. The MOSFET has three terminals: gate, source, and drain. The *gate* is the control terminal which controls current flow between the *source* and the *drain*. Unlike a planar MOSFET, the source and drain terminals are *not* symmetrical and cannot be interchanged. The source terminal of the MOSFET is internally connected to the *body* terminal of MOSFET so we do not show this terminal separately.

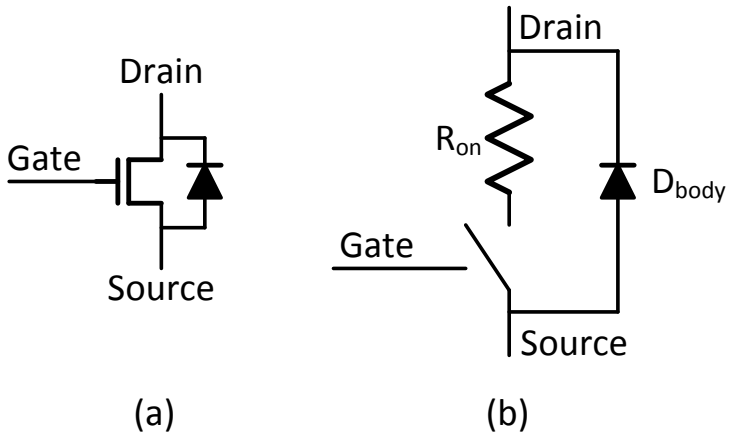

Figure 18.1: (a) Schematic symbol for a MOSFET with the body diode explicitly shown. (b) Simple switch-resistor model for a MOSFET.

A simple model of the power MOSFET is shown in Figure 18.1(b). The MOSFET behavior is modeled by a switch and a resistor. When the MOSFET is *fully-on* the source and drain are connected by a resistor with value R_{on}. When the MOSFET is *off* the source and drain are disconnected. The body diode has its own characteristics that we will visit in Chapter 19.

When the voltage between the *gate* and the *source*, V_{GS} exceeds the *threshold voltage* of the device V_T, the MOSFET is *on*. When $V_{GS} = V_{\text{on}}$ the MOSFET is *fully-on* and has resistance R_{on} between the source and drain. When $V_{GS} < V_T$ the MOSFET is off.

During the turn-on and turn-off transient we approximate the conductance of the FET with a straight line. If the turn-on time is t_{on} the conductance at time $t < t_{\text{on}}$ is approximately

$$G(t) = \frac{t}{t_{\text{on}} R_{\text{on}}} \tag{18.1}$$

In Section 18.3 we see that this linear current ramp is due to the linear build up of charge in the channel during turn on.

18.2 Structure

Figure 18.2 shows the physical structure of a typical power MOSFET. The source terminal is connected to both the P+ *body* region and the N+ drain

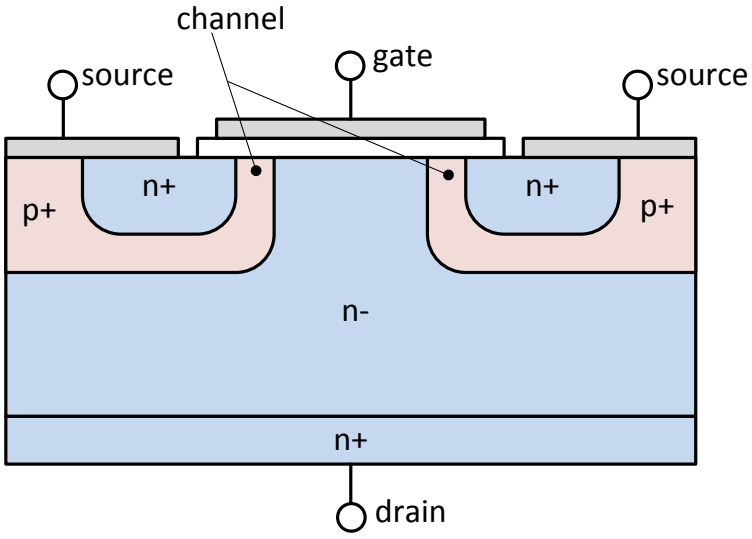

Figure 18.2: Structure of a D-MOS FET. The drain terminal on the bottom of the device connects to a lightly doped (N-) region to drop V_{DS} when the device is off. The insulated gate applies a field to invert the $P+$ silicon in the channel to form a conducting region when $V_{GS} > V_T$. The source terminal connects both to the $N+$ source terminal and to the $P+$ body region forming the body diode.

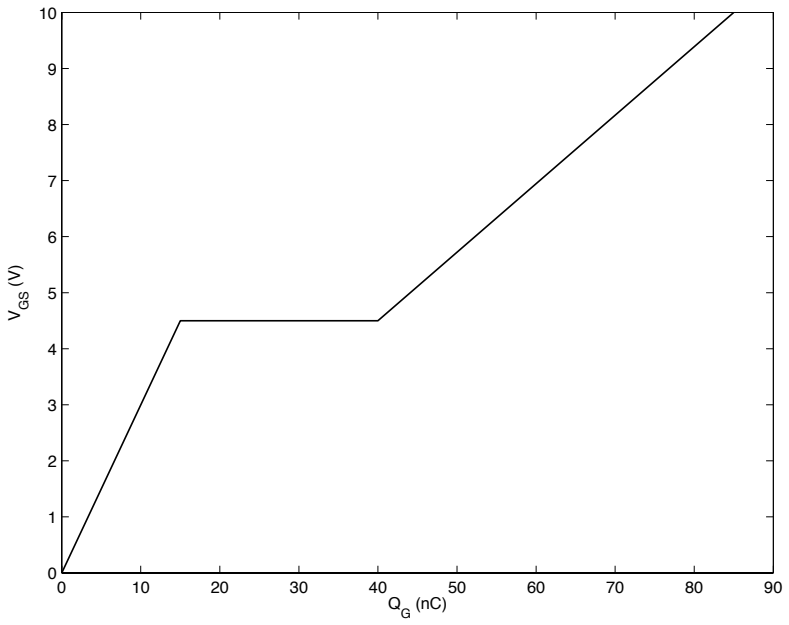

Figure 18.3: Gate voltage (V_{GS}) as a function of gate charge Q_G. This curve shows how the incremental capacitance (the inverse slope of the curve) changes with voltage. The gate has a fixed capacitance (linear curve) until $V_{GS} = V_T = 5$V. At this point the incremental capacitance becomes very large, V_{GS} remains nearly constant as charge is added, until the channel is fully inverted. The gate then takes on a different fixed capacitance as V_{GS} increases above V_T.

terminal. The gate terminal applies a field to the *channel*, the portion of the P+ region under the gate. The drain terminal is connected to the far side of the vertical N- region.

When the gate is driven above the threshold voltage, the P+ region under the gate becomes *inverted*, acting like an N region and resistively connecting the source and drain N-regions. When the gate is low, the junction between the P+ body region and the N- drain region is reverse biased. The reverse voltage is dropped across a depletion region that forms mostly in the lightly-doped N-region. If too large a reverse voltage is applied, the depletion region will grow completely across the channel connecting the source and drain and resulting in breakdown.

18.3 Dynamics

The speed of a MOSFET is governed by the time required to transfer charge to or from the channel via the gate terminal. The MOSFET can be turned on (and off) as quickly as charge can be moved in to (or out of) the channel. Once the channel inverts, all charge carriers transferred into the channel act to carry charge from drain to source. Hence, the more charge in the channel, the more carriers and the higher the conductance of the channel. When the gate terminal is fully on the channel is fully inverted and no further charge is transferred, at this point the conductance of the channel is $G_{\text{on}} = 1/R_{\text{on}}$.

The speed of charge transfer, and hence of MOSFET switching, is governed both by the internal resistance and inductance of the gate and source terminals and by the characteristics of the gate-drive circuit that completes the circuit between gate and source. We describe gate-drive circuits in more detail in Chapter 22.

The capacitance seen between a MOSFET's gate and source terminals is highly variable. For this reason MOSFET manufacturers specify a total gate charge Q_G in addition to specifying C_{GS}. The charge Q_G is more useful for understanding switching dynamics.

Figure 18.3 illustrates the variation in capacitance (the inverse slope of the V-Q curve) as a function of gate charge for a 600V 90mΩ MOSFET (FCB36N60N). As the gate charges from $V_{GS} = 0$ to V_T the capacitance is fixed. This is the capacitance from the gate terminal to the source and drain terminals. When V_{GS} reaches V_T, at 15nC of charge the channel inverts and the incremental capacitance becomes very low. During this period the drain is rising and the charge applied to the gate is consumed charging the Miller capacitance C_{DG}. V_{GS} holds nearly constant as the charge ramps from 15nC to 40nC. At this point the drain stops rising and the incremental capacitance becomes a different fixed value.

If the gate is being driven with a 1A current source Figure 18.3 shows the voltage on the gate as a function of time. The current source deposits 1nC on the gate every 1ns. With this drive the transistor woudl remain off for the first 15ns. Then, when $V_{GS} = V_T = 4.5$V its conductance begins to linearly ramp from 0 to 11S over the next 25ns, 0.44S/ns. If the MOSFET is in the configuration of Figure 17.2 with $V_{DS} = 500$V, current ramps at (500)(0.44)=222A/ns. In practice finite inductance in the circuit slows the current ramp to something considerably slower.

18.4 Parasitic Elements

Because a properly driven MOSFET can turn on blindingly fast, the parasitic circuit elements can have a first-order effect on the performance of the MOSFET and its switching losses. Figure 18.4 shows the major parasitic circuit elements associated with a typical MOSFET. Table 18.4 lists typical parasitic values for a 600V 100mΩ power MOSFET.

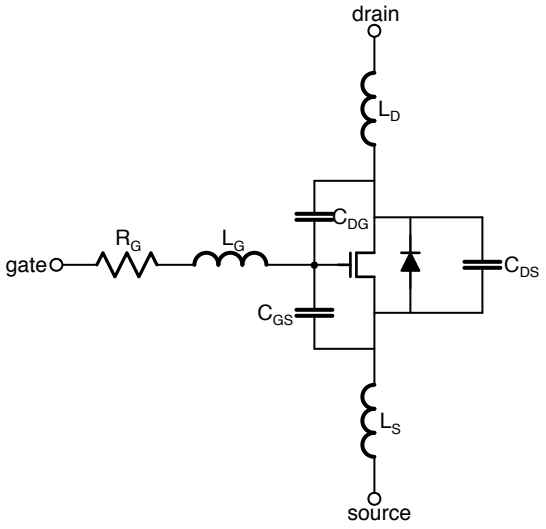

Figure 18.4: MOSFET parasitic circuit elements. At high operating frequencies these parasitic circuit elements have a large effect on the performance and switching losses of a MOSFET.

Symbol	Value	Units	Description
C_{DS}	200	pF	Drain-source capacitance
C_{DG}	70	pF	Drain-gate (Miller) capacitance
C_{GS}	3600	pF	Gate-source capacitance
Q_G	86	nC	Total gate turn-on charge
Q_{GD}	35	nC	Gate-drain turn-on charge
L_S	7	nH	Source inductance
L_D	3	nH	Drain inductance
L_G	7	nH	Gate inductance
R_G	1.5	Ω	Gate resistance

Table 18.1: Values of parasitic circuit elements for a typical 600V 100mΩ power MOSFET in a TO-220 package.

Parasitic inductance, which is largely due to the package, limits how fast current can ramp in the gate-source and source-drain circuits. While we list all three lead inductances, for circuit analysis we are concerned with two sums. The gate-source inductance $L_{GS} = L_G + L_S$ determines how fast the gate driver can ramp current to charge the gate. Given L_{GS}=14nH for our typical FET, a 10V gate driver can produce a current slope of at most 0.7A/ns. In practice the gate driver and circuit board will add more inductance and some of the 10V will drop across the resistance of the FET and the gate driver so we may be lucky to get 0.5A/ns.

The drain-source inductance $L_{DS} = L_D + L_S$ gets added to other inductances in the loop supplying current during the switching transient to limit current ramp during turn-on and turn-off. Our typical FET has L_{DS} = 10nH. Parasitic inductances in the diode, capacitor, and circuit board traces that complete the loop can easily triple this bringing the total loop inductance to L_{loop} =30nH. With 500V across 30nH the current slope is limited to 17A/ns. This is an order of magnitude slower than the FET itself is capable of.

Energy stored in the loop inductance causes overshoot and ringing. Consider the circuit of Figure 17.3 but with 30nH of inductance around the loop that includes the FET, diode, and voltage source. As the drain falls to ground during turn-on, the current in the loop inductance will exceed 20A. The excess current will ring out in an LRC circuit formed by L_{loop}, R_{on}, and C_{DS}, and the energy represented by this excess current will be lost. During turn-off 20A in L_{loop} represents 12μJ of stored energy that will be lost and can lead the drain to overshoot.

Minimizing inductance is of first-order importance. Newer packages, such as *thinPAK* have L_{DS} as low as 3nH which, in conjuction with careful board layout and parallel arrays of bypass capacitors, greatly speeds switching and reduces losses.

The Miller capacitance C_{DG} is multiplied by the difference in drain vs gate voltage swing. This has two ill effects. First, due to the Miller effect, C_{GD} accounts for over a third of the total gate charge slowing gate turn-on. Charge Q_{GD} is C_{DG} multiplied by the 500V swing. More critical, however, is that this capacitive feedback can lead to oscillations. During turn-on with a slow gate driver capacitive feedback can turn the FET back off. This in turn causes the drain to rise turning the FET back on, and so on. A similar phenomena can occur during turn-off where the rising drain turns the FET back on. In either case the resulting oscillations can greatly increase switching losses.

The drain-source capacitance C_{DS} along with the capacitance of the diode must be charged and discharged as the drain rises and falls. This both leads to CV^2 energy losses and increases switching losses by slowing the drain transition. At 500V, the energy stored on 200pF is 50μJ.

The gate-source capacitance is not really a parasitic element since its critical to the operation of the FET. However this large capacitance complicates the design of the gate drive circuit. Because the capacitance varies with voltage, as explained in Section 18.3, it is more useful to design the gate driver to provide the charge Q_G than to drive C_{GS}.

18.5 Avalanche Energy

Inductor current flowing into the drain node of a MOSFET as it turns off can charge the drain node above $V_{DS\max}$. This occurs, for example, in a flyback converter (Figure 10.3) where the energy from the leakage inductance of the transformer is transferred to the drain node. When V_{DS} rises above $V_{DS\max}$ the reverse-biased body diode of the MOSFET is driven into *avalanche breakdown* and begins to conduct.

This avalanche breakdown is not destructive to the MOSFET as long as the total energy of the event is within limits. The FCN36N60 FET, for example, can handle up to 1.8J of energy in a single avalanche breakdown event and up to 3.12mJ in repetitive events. This avalanche breakdown serves a useful purpose in applications such as the flyback converter where it serves as the energy-absorbing device in Figure 10.3.

18.6 Typical MOSFETs

The key data-sheet parameters of six representative MOSFETs are shown in Table 18.2. The part numbers are listed in the first row of the table. The first four columns are silicon MOSFETs with voltage ratings of 20V, 60V, 200V, and 600V respectively. They illustrate how key MOSFET parameters vary with voltage. The fifth column is a 200V gallium-nitride (GAN) FET. It is instructive to compare the properties of the 200V GAN FET with the 200V Si FET. The final column is a 1200V silicon-carbide (SiC) FET. Comparing this device to the 600V silicon FET highlights its properties.

The second row lists the maximum drain-source voltage supported by the FET. As voltage increases, R_{on}, tabulated in the third row increases approximately quadratically. To see this trend, the quantitity $V_{DS\max}^2/R_{\text{on}}$ is listed in the second to last row. Most of the devices have a value near 2MW. The 600V Si FET is a bit higher at 4.4MW. The 1200V SiC FET stands out at 58MW highlighting the advantage of SiC. SiC devices switch very high voltages with very low resistance — making them superior to IGBTs in high-voltage applications except for price.

The fourth row of the table lists gate charge Q_G which remains relatively constant between 70 and 90nC for the Si FETs across the voltage range. The GAN FET has a value more than an order of magnitude lower at 5nC. This very low gate charge allows GAN FETs to be switched very quickly. Thus, properly used they have very low switching losses. Driving the gate of a GAN FET is tricky, however, current devices require the gate voltage to be kept between -5 and +6V vs ± 20V for Si FETs. The SiC FET also has a relatively low Q_G for the amount of power it is switching. To see this, the quantity $V_{DS\max}^2/R_{\text{on}}Q_G$ is tabulated in the last row. The GAN FET and SiC FET have values nearly an order of magnitude higher than the Si FETs for this figure of merit.

The output capacitance C_{oss} is listed in the fifth row of the table. This value is

Device	20V	IRLB3036	IRFB4227	FCB36N60N	EPC2010	C2M0025120	Units	Description
$V_{DS\max}$	20	60	200	600	200	1200	V	Maximum V_{DS}
R_{on}		1.9	20	81	18	25	mΩ	On resistance
Q_G		91	70	86	5	161	nC	Gate charge
C_{oss}		1020	460	80	270	220	pF	Output capacitance $C_{SD} + C_{GD}$
$I_{D\max}$		195	65	36	12	90	A	Maximum continuous drain current
I_{DM}		1100	260	108	60	250	A	Maximum pulsed drain current
E_{AS}		290	140	1800			mJ	Single-event avalanche energy
$P_{\max}$		380	330	312		463	W	Maximum power dissipation
V^2/R		1.9	2.0	4.4	2.2	58	MW	Figure of merit
V^2/RQ_G		21	29	52	440	360	mV/s	Figure of merit

Table 18.2: Key parameters of six field-effect transistors.

describe the parts
walk through the parameters

18.7 Exercises

Chapter 19

Diodes

Diodes are switches that control themselves by allowing current to flow in one direction but then switching off when the current attempts to flow in the other direction. They are widely used in high-voltage converters and play a key role in low-voltage converters (which use MOSFETs as synchronous rectifiers).

The schematic symbol for a diode is shown in Figure 19.1. The *arrow* in the symbol indicates the direction of current flow. In an ideal diode, current flows only in one direction, from the *anode* to the *cathode*. Substantial reverse current only flows if the reverse voltage exceeds V_R, the maximum reverse voltage or *blocking* voltage of the diode.

At first glance a diode is deceptively simple — a two terminal device that passes current in one direction. The dynamic behavoir of a diode, however, is often a major concern in the design of a converter or motor drive. When the voltage across a conducting diode reverses, the diode does not immediately turn off. Instead, it conducts current in the opposite direction for a brief period of time. The timing of this *reverse recovery* and its *softness* greatly affect system performance. A diode with a fast reverse recovery time stops conducting quickly — before a large reverse current has built up. A diode with a *hard* or *snappy* recovery abruptly switches its current to zero when it recovers — possibly caus-

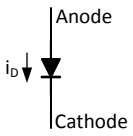

Figure 19.1: Schematic symbol for a diode. Current i_D nominally flows only from the *anode* to the *cathode* and not in the reverse direction.

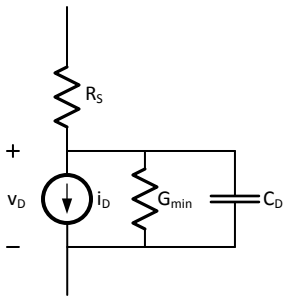

Figure 19.2: Simple DC model for a diode consists of a current source i_D, a series resistance R_S, and a parallel conductance, $G_{\min}$.

ing damaging voltage transients across stray inductances. A diode with a *soft* recovery, on the other hand, more gradually reduces its reverse current to zero — reducing inductive voltage transients. Some diodes are also slow to turn on. This *forward recovery* behavior can also cause issues.

19.1 Simple Model

A simple large-signal model for a diode operated above its maximum reverse voltage ,show in Figure 19.2, consists of a dependent current source, two resistances, and a variable capacitance. The series resistance R_S dominates the voltage drop of the diode at high currents. The parallel conductance $G_{\min}$ determines the current when reverse biased and at low positive voltages. It is included primarily to aid convergence of simulations. The DC diode current i_D depends exponentially on the voltage across the current source v_D:

$$i_D = I_S \left(\exp\left(\frac{v_D}{N V_T} \right) - 1 \right) \tag{19.1}$$

Where V_T is the *thermal voltage* which is given by:

$$V_T = \frac{KT}{q} \tag{19.2}$$

At room temperature (300 °K), $V_T = 25.85$mV. Thus each time the diode voltage increases by 28.85N mV, current increases by a factor of e. Because V_T increases linearly with temperature, one might think that diode current would decrease as temperature increases. This is not the case because the saturation

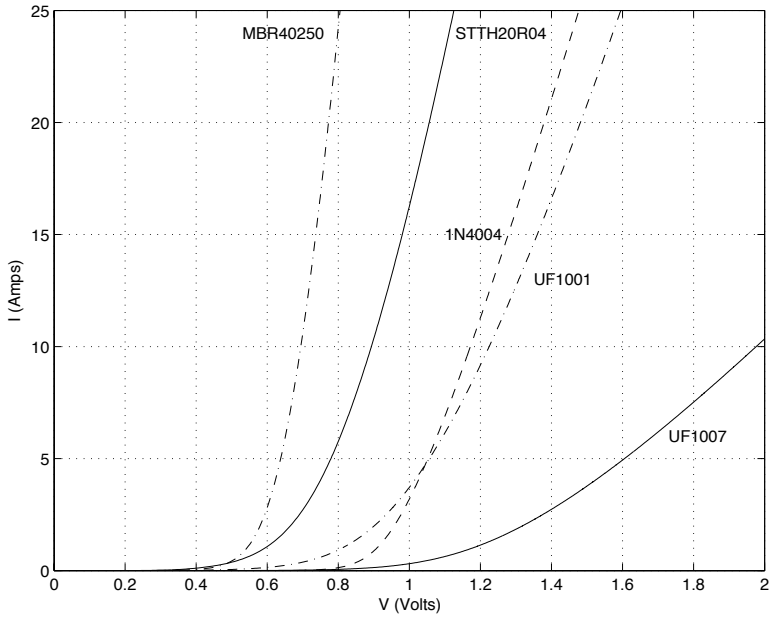

Figure 19.3: I-V Characteristics of the diodes from Table 19.1. These curves illustrate how one sacrifices low V_F for fast switching time or high V_R.

current I_S has a stronger temperature dependence.

$$I_S = I_0 \left(\frac{T}{T_0} \right)^{X/N} \tag{19.3}$$

$$\tag{19.4}$$

Where T is the diode temperature (in °K), T_0 is not nominal temperature (typically 300 °K), I_0 is the saturation current at T_0, N is the ideality parameter of the diode, and X is a parameter of the diode in question. For typical silicon junction diodes $X = 3$ and for Schottky diodes, $X = 2$.

Curves showing the first quadrant of this DC I-V relationship for several typical diodes are show in Figure 19.3. As voltage is increased, the current increases exponentially until the voltage drop across the series resistance exceeds the voltage drop across the current source. At currents above this point, the V-I characteristic is dominated by the series resistance R_S. The slope of this portion of the curve is $1/R_S$.

The V-I curve of the 1N4004 diode, plotted in Figure 19.3 are derived from parameters $I_S = 3.70\text{nA}$, $N = 1.774$ and $R_S = 17.56\text{m}\Omega$.

The diode capacitance C_D in Figure 19.2 dominates the transient response

of the diode. It is the sum of junction capacitance and the diffusion capacitance:

$$C_D = C_j + C_d \tag{19.5}$$

Where the junction capacitance dominates at zero and reverse bias and decreases with increasing reverse bias:

$$C_j = \frac{C_{j0}}{\sqrt{1 - \frac{V}{V_0}}} \tag{19.6}$$

The diffusion capacitance dominates during forward bias and accounts for minority carriers stored in the diode:

$$C_d = \frac{\tau_T}{NV_T} i_D \tag{19.7}$$

Where τ_T is the transit time of the diode. The diffusion capacitance accounts for the reverse recovery time of the diode.

19.2 Key Parameters

Diodes come in many different varieties. In this book we are concerned only with *power* diodes, sometimes referred to as *rectifiers*, not small signal diodes, RF diodes, varactors, or other sundry types. Even power diodes come in many flavors depending on the voltage, current, and speed required.

The key data-sheet parameters of some representative power diodes are shown in Table 19.1. The I-V characteristics of these diodes are plotted in Figure 19.3. The first diode, a STTH20R04 is the diode shown in Figure 17.1. This is a 400V, 20A diode. That is, it will block 400V across the diode in the reverse direction and will carry an *average* forward current of 20A when biased in the forward direction. Forward current specifications need to be read carefully. This 20A diode can handle *non-repetitive surge current* of up to 150A and RMS forward current up to 50A. The STTH20R04 is specified as having a forward voltage of $V_F = 1.5$V at 25°C and 20A. This is a rather pessimistic maximum. Figure 19.3 shows the diode dropping 1.05V under these conditions. This voltage drop is reduced by 0.35V at 125°C which leads to *current hogging* as described in Section 19.5 below.

The STTH20R04 is a fast diode — with a reverse recovery time of 35ns — and has a moderately soft recovery with $S = 0.3$. (The softness factor will be explained in more detail in Section 19.3 below.)

The second diode in the table is a 1N4004. This is a slow, 1A, 400V diode often used to rectify 60Hz AC. The 1N4004 is rated for *average* forward current I_F=1A, as are the next two diodes in the table. However, all three of these 1A diodes can handle non-repetitive surge currents of up to I_{FSM}=30A.

Device	STTH20R04	1N4004	UF1001	UF1007	MBR40250	Units	Description
V_R	400	400	50	1000	250	Volts	Maximum reverse voltage
I_F	20	1	1	1	40	Amps	Maximum forward current
I_{FSM}	150	30	30	30	150	Amps	Maximum non-repetitive forward surge current
V_F	1.5	1	1	1.7	0.86	Volts	Forward voltage
t_{rr}	35	1000	50	75	35	ns	Reverse recovery time
S	0.3						Softness factor
t_{fr}	150					ns	Forward recovery time
V_{FP}	2.5					Volts	Softness factor

Table 19.1: Key parameters of five diodes.

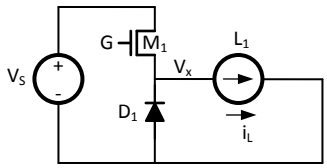

Figure 19.4: A buck converter using a diode for the *low-side* switch illustrates the problems of diode reverse recovery.

The main difference between the 1N4004 and the other diodes is its slow speed. The 1N4004 data sheet does not list a reverse recovery time, but measurements show it to be several microseconds. This slow recovery time makes the 1N4004 unsuitable for use in a switching converter operating at 10s of kHz or faster.

The next two diodes are fast 1A diodes but with very different voltage ratings. The UF1001 has a V_R of only 50V while the UF1007 can block V_R =1kV. Figure 19.3 shows that the cost of blocking voltage is increased V_F. At 1A, the UF1001 drops 0.81V vs 1.2V for the UF1007. One also pays a price in speed for increased blocking voltage. The UF1007 is slightly slower than the UF1001.

The final diode in the table is a 250V, 40A Schottky diode. Schottky diodes use a metal-semiconductor junction rather than a p-n junction. This construction leads to a very fast diode with a significantly lower V_F making them ideal for most switching applications. Unfortunately, silicon Schottky diodes are not available with high blocking voltages. Recently, Silicon-Carbide (SiC) Schottky diodes have become available with blocking voltages up to 3.3kV. However, they are quite expensive.

19.3 Reverse Recovery

When the current through a diode reverses, the diode does not instantly turn off. Current flows in the reverse direction through the diode (from cathode to anode) until the minority carriers are cleared. This, sometimes slow, process of turning the diode off is referred to as *reverse recovery*. It has a major effect on the performance and efficiency of switching power converters. To understand the effect of diode reverse recovery on power converters consider the buck converter shown in Figure 19.4. To simplify our analysis we consider the inductor to be a current source — which is nearly the case if the ripple current is small. A waveform showing diode current during the transient when M_1 turns on is shown in Figure 19.5.

Initially M_1 is off and all of the inductor current $I_1 = I_L$ is flowing in the diode. As M_1 turns on, its current ramps up — we assume a linear ramp —

and the diode current ramps down. At time t_1 MOSFET M_1 is carrying all of the inductor current $i_M = I_L$ and the diode current reaches zero, $i_D = 0$. An ideal diode would stop conducting at this point. However, a real diode continues conducting, clamping V_x to 0 until a current of $-I_{rr}$ is reached time at time t_2. At this point diode current starts ramping down and voltage V_x begins to rise. At time t_3, diode current is effectively zero[1] and the reverse recovery process is complete.

The reverse recovery time $t_{rr} = t_a + t_b = t_3 - t_1$ is the total time required for the diode to recover from the zero crossing of the current waveform until the current returns to near zero. The reverse recovery charge Q_{rr} is the shaded area in the figure. Reverse recovery time varies widely depending on the initial current slope. The steeper the slope the faster the recovery but the larger I_{rr}. Q_{rr} is roughly constant independent of current slope.

The recovery softness factor S is defined as the ratio of the absolute value of slope at the zero crossing to the maximum slope during the recovery:

$$S = \frac{\frac{-di_D(t_1)}{dt}}{\max_{t=t_2}^{t_3}\left(\frac{di_D}{dt}\right)} \tag{19.8}$$

In the figure the maximum up-slope is exactly twice the down slope so $S = 0.5$. The smaller S the larger the voltage transients induced by the diode current going to zero.

If the FET switches quickly and the diode is slow in the buck converter of Figure 19.4, the reverse recovery currents can get very large. The power dissipated in the FET by this current — while the full voltage is still dropped across the FET — dominates switching losses.

Consider the simulated current waveform shown in Figure 19.6 for the diode current in a buck converter with an IRLB3036 (60V, 2mΩ MOSFET) and an 8EWF ($t_{rr} = 140$ns) diode. The MOSFET is driven with a 4Ω gate resistor giving a very fast turn-on. As the MOSFET turns on the diode remains conducting — causing the MOSFET to effectively short power and ground. The simulated waveforms show the current exceeding 500A before the diode recovers. Such a current level would result in destruction of the FET, the diode, or both, after a number of repetitions if not immediately. The simulated waveform also shows an extremely hard recovery with the current going from 500A to zero almost instantly. This is an artifact of the SPICE diode model. The 8EWF actually has a relatively soft recovery with $S = 0.5$.

In addition to the damaging current, the waveform shown in Figure 19.6 also represents an enormous switching loss. With 50V across the FET, the 500A current results in 25kW peak dissipation. Over the 15ns recovery time the FET dissipates 170μJ due to reverse recovery current. In contrast, the loss due to the FET's finite switching time (about 10ns) at 50V and 20A (1kW peak

[1]The recovery process is considered to be complete when the current reaches 10% of its peak negative value. The current never returns entirely to zero as there is a small amount of reverse current in the steady state.

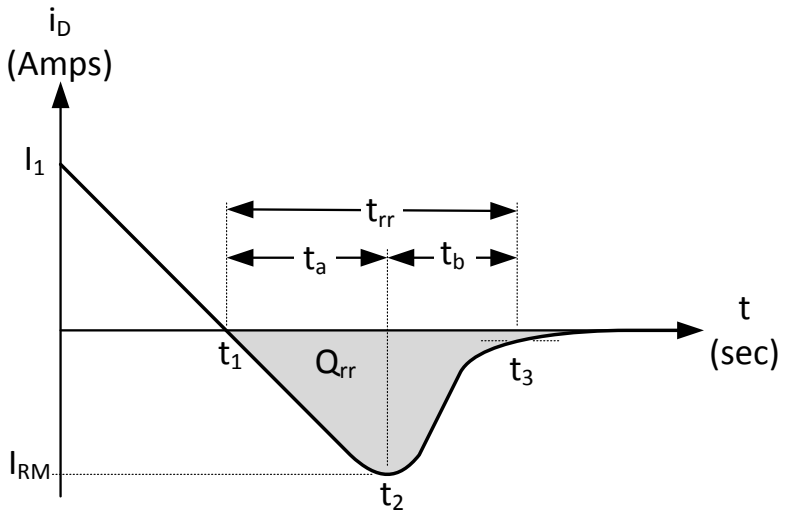

Figure 19.5: Diode current during reverse recovery.

power) is about 5μJ. The slow recovery of the diode increases the switching loss of the FET by 34×. At a 200kHz switching frequency, the reverse recovery results in 34W of switching loss compared to 1W for the FET with an infinitely fast diode. The large reverse current can be reduced by switching the FET more slowly — by using a larger gate resistor. However the switching loss due to reverse recovery is largely unaffected by FET switching speed. The only way to reduce this switching loss is to use a faster diode (or a different converter topology that avoids the problem entirely).

The recovery time in Figure 19.6 is 15ns while the 8EWF diode has $t_{rr} =$ 140ns. This discrepancy is due to the current slope. The 8EWF data sheet specifies $t_{rr} =$ 140ns with a current slope of 25A/μs. The current slope in Figure 19.6 is 40A/ns, 1600 × faster. The t_{rr} specifications from different diodes are not directly comparable unless they are measured with identical current slopes.

Figure 19.7 shows simulated diode current for the same test case as Figure 19.6 but with a modified SPICE model that more accurately models the soft recovery of the diode. The destructive current and high switching losses remain (they are real), but the current slope during recovery is more realistic. The slope is still 10^{11} A/s which can still deliver destructive voltages across even small amounts of parasitic inductance.

The magnitude of the peak reverse recovery current $I_{\rm rr}$ can be reduced by

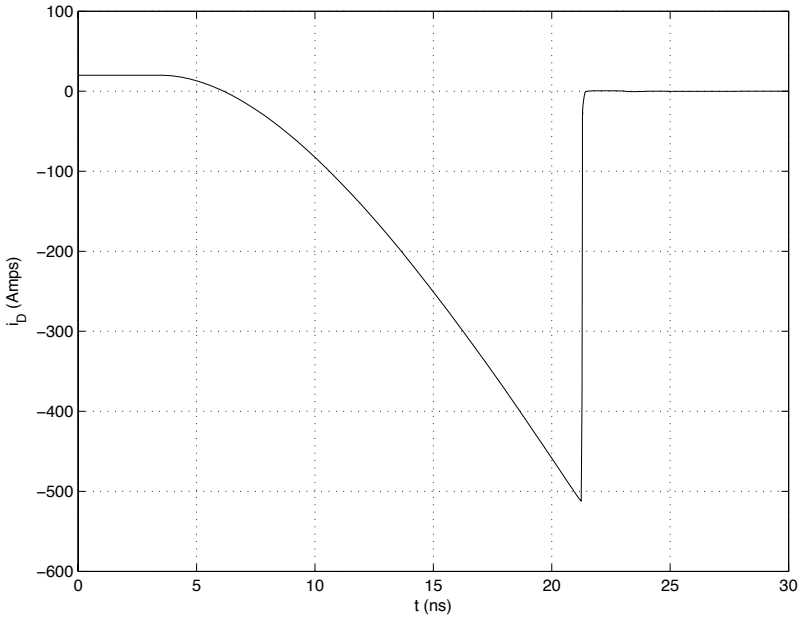

Figure 19.6: Simulation of diode current during reverse recovery for an 8EWF diode (t_{rr} =140ns).

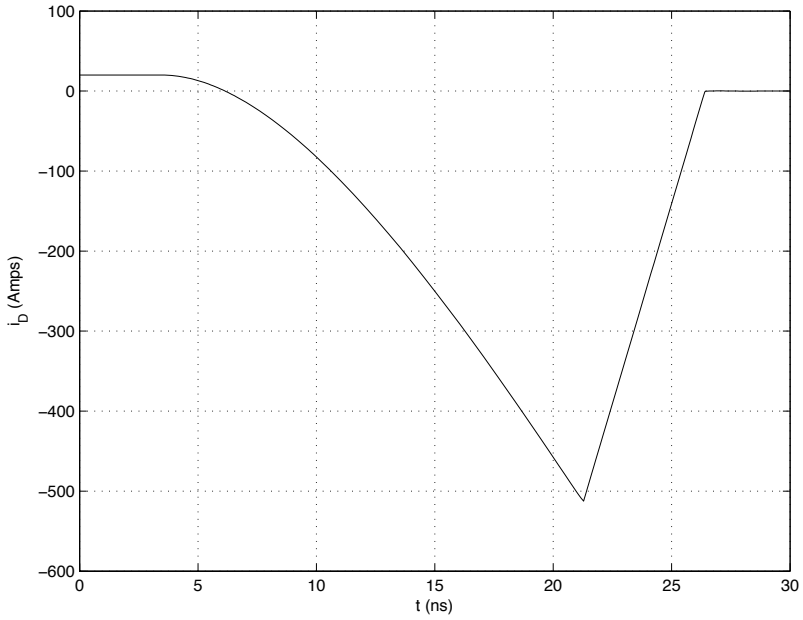

Figure 19.7: Simulation of diode current during reverse recovery for an 8EWF diode (t_{rr} =140ns) using a SPICE model that more accurately models soft recovery.

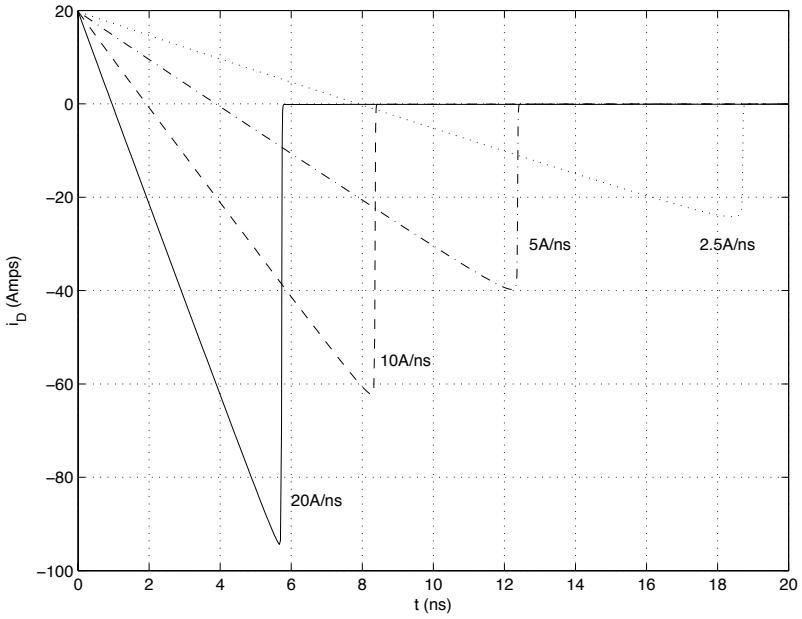

Figure 19.8: Simulation of diode current for an STTH20R04 diode (t_{rr} =35ns) for varying current slopes.

reducing the current slope when the FET turns on — typically by using a larger gate resistor. While this can prevent the destructive current shown in Figure 19.6, it does not greatly affect switching energy because the reverse recovery charge Q_{rr}, the shaded area in Figure 19.5 remains roughly constant and $E_{swrr} = Q_{rr}V_{ds}$.

The effect of current slope on reverse recovery current is illustrated in Figure 19.8. The figure shows the simulated reverse recovery current of a STTH20R04 diode for current slopes varying from 2.5A/ns to 20A/ns. (For comparison the slope in Figure 19.6 is 40A/ns.) As the currrent slope is reduced the peak reverse current is reduced nearly proportionally. However the switching energy due to reverse recovery, which is proportional to Q_{rr} decreases much more slowly. For the for curves shown in Figure 19.8 Q_{rr} is 230, 210, 180, and 140nC respectively (with the larger slope having the larger Q_{rr}).

19.4 Forward Recovery

Power diodes also take some time to fully turn on. When a large forward current is rapidly applied to a diode the resistance it encounters is initialy

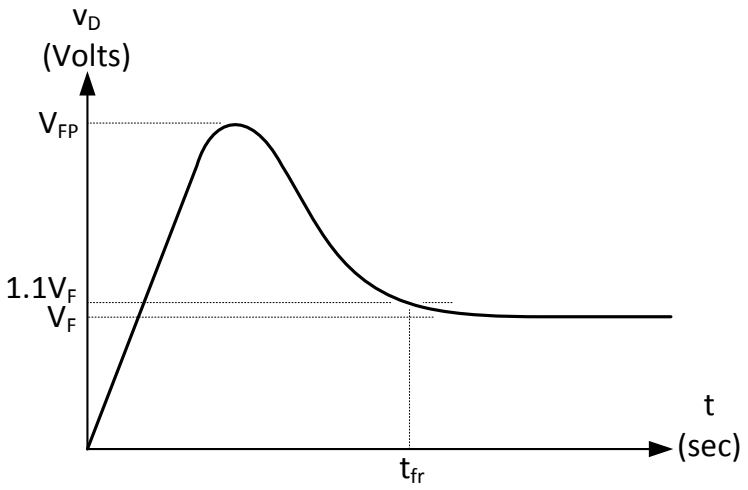

Figure 19.9: Diode voltage during forward recovery. A rapid current ramp is applied to the diode at $t - 0$. The forward voltage initially overshoots to V_{FP} before dropping to its steady-state value of V_F. The forward recovery time t_{fr} is the time for the forward voltage to drop to within 10% of its steady-state value.

higher, sometimes substantially higher, than the steady-state DC resistance. This is because the intrinsic region of the PIN diode initially has few charge carriers. This high initial resistance can result in a high initial forward voltage V_{FP} as shown in Figure 19.9. Once the density of charge carriers reaches its steady-state value, the resistance of the diode drops to its steady-state vaue as does the forward voltage.

The STTH20R04 diode specifies a forward recovery time of $t_{fr} = 150$ns and a maximum peak forward voltage of $V_{FP} = 2.5$V (typical V_{FP} is 1.7V). Compared to the steady state $V_F = 1.5$V, this represents a worst-case initial 67% increase in forward voltage.

Consider the buck converter of Figure 19.4. When transistor M_1 turns off, i_L is applied to diode D_1 with a ramp that depends on the turn-off time of M1 — which may be 10ns or less. Normally we would expect D_1 to keep V_x clamped to within V_F of ground. However, during reverse recovery V_x may drop as low as V_{FP} below ground. Depending on the circuitry attached to this node, a large negative undershoot could present a problem. Some diodes have a V_{FP} far worse than the 2.5V of the STTH20R04.

19.5 Current Hogging

Connecting multiple diodes in parallel — to handle a larger current can be problematic because diode current *increases* with temperature. Unless the thermal design keeps the parallel diodes at nearly the same temperature or *ballast* resistors are placed in series with the diodes one diode may wind up hogging most of the current — until it fails catastrophically.

This *current hogging* is an example of positive feedback. One diode gets a little warmer than the others. This causes its current to increase. This in turn causes it to dissipate more power and get even warmer. The cycle continues until the diode is very hot and carrying nearly all of the current.

Figure 19.10 shows the temperature sensitivity of a 1N4004 diode. As temperature increases the diode current at a given voltage increases significantly (or conversely the voltage at a given current decreases signficantly). Neglecting series resistance, for a typical diode, current is proportional to a power of temperature (in °K) as given by (19.1). At the rated current of 1A, increasing the temperature by 1°C decreases the voltage drop by approximately 1mV.

Referring to Figure 19.10, consider a case where four 1N4004 diodes are connected in parallel to carry an average current of 4A. If one diode becomes hot (125°C) while the others remain at room temperature (25°C), the hot diode will carry 3A of current while the cool diodes carry about 0.3A each. This will cause the hot diode to get hotter, causing it to carry even more current until it fails.

The dynamics of diode *thermal runaway* are illustrated in Figure 19.11. The figures shows simulated temperature (top panel) and current (bottom panel) for four diodes connected in parallel sharing a 20A load. One diode has a thermal resistance to ambient that is 10% higher than the other three. This

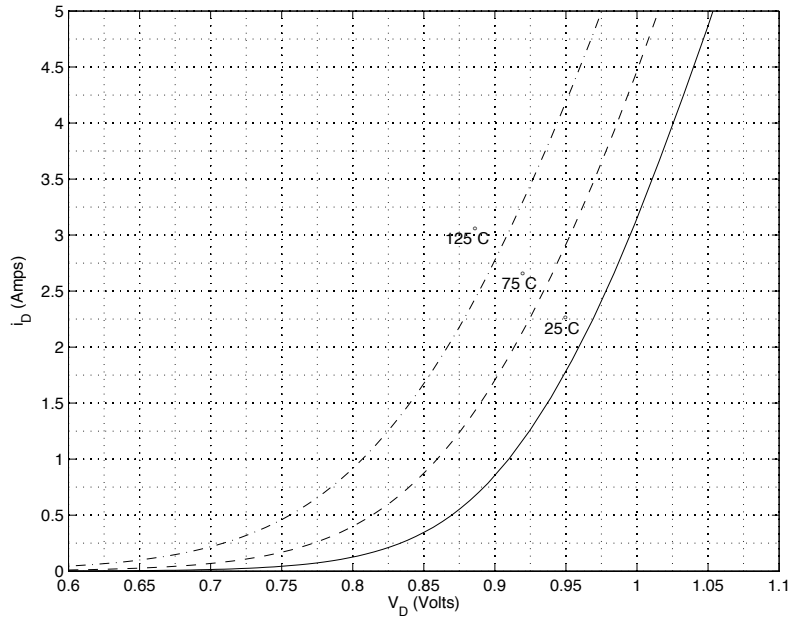

Figure 19.10: I-V characteristics of a 1N4004 diode at 25°C, 75°C, and 125°C. Current increases as temperature increases leading to *current hogging* if multiple diodes are connected in parallel.

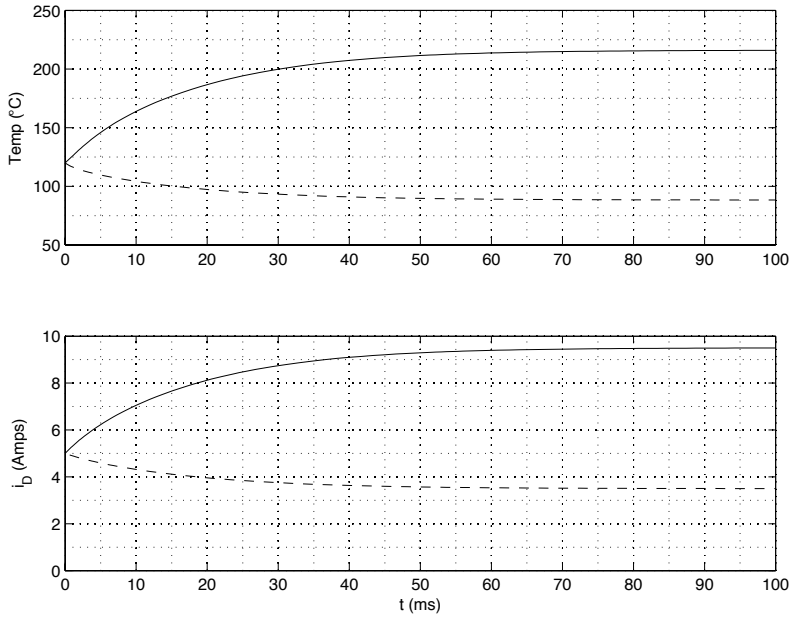

Figure 19.11: Thermal runaway caused by current hogging with four diodes connected in parallel carrying 20A and with a nominal thermal resistance of 20°C/W and a thermal time constant of 1ms. One diode has a thermal resistance that is 10% higher than the others. The top panel shows diode temperatures. Diode current is shown in the bottom panel.

10% mismatch in thermal resistance leads to a 270% mismatch in current. In the steady state the *hot* diode (solid line) is carrying 9.5A of current while the other three diodes (dotted line) carry 3.5A each.

Thermal mismatches are quite common and can be significantly larger than 10%. They may be caused by component placement, proximity to other power dissipating components, the direction of convective or conductive heat flow, and other factors.

If diodes are to be placed in parallel, the thermal design should use a physical arrangement that keeps the diode's junction temperatures (not case temperatures) closely matched, and/or a *ballast* resistor should be placed in series with each diode as shown in Figure 19.12 to reduce the current mismatch that results from temperature mismatch.

Figure 19.13 shows the effect of ballast resistors on the V-I characteristics of a typical diode as a function of temperature. The figure shows the V-I characteristics for a 1N4004 diode with and without a 0.5Ω ballast resistor at 25°C and 125°C. Without the ballast resistor, two diodes at the same voltage

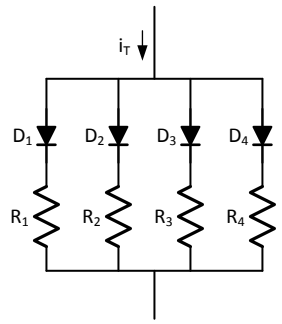

Figure 19.12: To balance current through parallel diodes, a *ballast* resistor should be inserted in series with each diode.

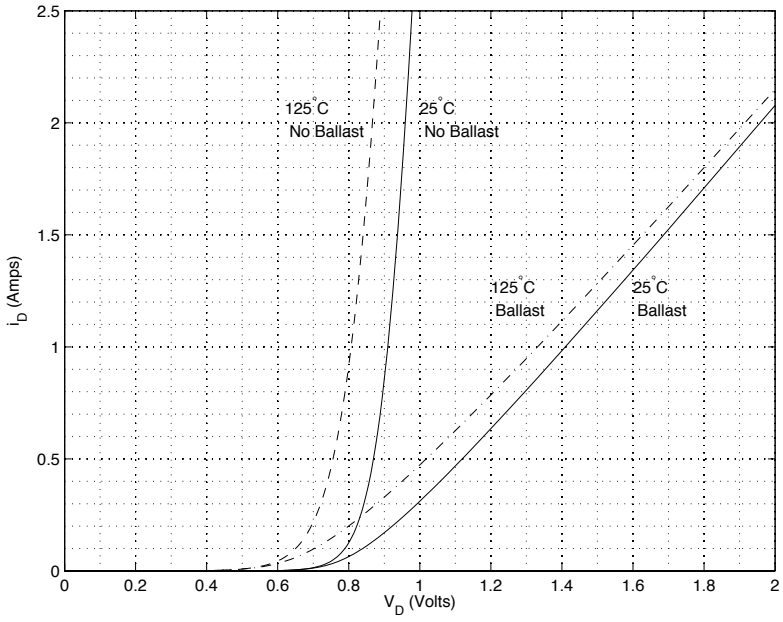

Figure 19.13: I-V Curves of 1N4004 diodes at 25°C and 125∗∘C with and without 0.5Ω ballast resistors.

but different temperatures can have vastly different currents. At v_D=0.8V, for example, the 125°C diode is carrying 0.9A while the 25°C diode carries only 0.1A.

Adding the ballast resistor greatly reduces this imbalance. At 1.1V for example the hot diode carries 0.6A while the current in the cool diode is 0.48A. The difference is narrowed from 0.8A to 0.12A. With ballast, the current difference between the hot diode and the cold diode narrows further as current increases — the lines converge and at some current will cross over. This is because resistors have a positive temperature coefficient — resistance increases, and hence current decreases, as temperature increases — naturally balancing current between parallel resistors.

The downside of providing current balance using ballast resistors is the power dissipated in the ballast resistors and their cost. In some cases the ballast resistor can be implemented with a foil trace on a printed-circuit board minimizing cost.

19.6 SPICE Models

The standard Berkeley SPICE diode model which is used in most derivatives of the SPICE circuit simulator can reasonably accurately capture the DC characteristics of a diode. However, it does a very poor job modeling the dynamics of a diode — reverse recovery and forward recovery — that are important in power electronics.

Even for the aspects of a diode that can be captured by the standard SPICE diode model, accurate simulation depends on having an accurate model. Even models from the diode manufacturers are not always accurate or may not be accurate under the test conditions (e.g., current slope) important in your application. For any critical diode application there is no substitute for a physical experiment to measure the important parameters on a real diode.

The DC characteristics of a simulated diode are determined by SPICE parameters `IS`, `RS`, and `N`. If these parameters are accurately fit, for example by using the *three-point method*, the resulting I-V curve is a fairly accurate representation of the behavior of the actual device. Similarly temperature dependence is accurately captured by parameter `XTI` which is typically 3 for silicon junction diodes and 2 for Schottky diodes.

Reverse recovery time is determined by model parameter `TT`. The model gives a rough dependence of recovery time on the current slope but does not model the *softness* of the recovery. The standard model has a very hard recovery $S = 0$ which is worse than any real diode. This is illustrated in Figure 19.5.

The standard SPICE model has no method to handle *forward recovery*. It can only simulate diodes that turn-on instantly with no initial increase in series resistance.

A subset of the SPICE model parameters for some of the diodes discussed in this Chapter are shown in Table 19.6. The first three columns I_S, R_S, and N

determine the DC I-V characteristics of the diode according to (19.1).[2] The next two columns, τ_T and C_{J0}, determine the diode capacitance and reverse recovery time according to (19.6) and (19.7). Note that the transit time for the 1N4004 4.3μs is more than 100$\times$ larger than the transit time of the STTH20R04 at 15ns. The next parameter XTI is the variable X in (19.4) that largely determines the temperature sensitivity of the diode. Finally, BV is the breakdown voltage, or maximum reverse voltage, V_R of the diode.

The first two rows are different manufacturer's models for the same diode, the 1N4004. The top row is a model from Fairchild Semiconductor while the next row is a model from Diode's Inc. The model parameters differ substantially — by two orders of magnitude in I_S — and the I-V curves they generate are quite different. Models, even from manufacturers should not be trusted until they have been verified against physical experiments.

19.7 Measurements

Any diode parameters that are critical to operation of your systme should be verified by experiment. Typically one measures the I-V curve at at least two temperatures (like Figure 19.10), the reverse recovery current profile (like Figure 19.5), and the forward recovery voltage profile (like Figure 19.9). The forward and reverse recovery curves should be measured with the current slopes that will be seen in the actual application.

The buck converter of Figure 19.4, with a series resistor added between the diode and ground for current measurement, is an ideal test set for taking the transient measurements. Current slope can be adjusted by varying the gate resistor.

[scope photos of reverse and forward recovery]

19.8 Exercises

Calculate losses due to reverse recovery and FET switching time in buck converter.

Simulate losses due to reverse recovery in a boost converter

[2] The actual SPICE parameter names are `IS, RS, N, TT, CJO, XTI` and `BV`.

Diode	I_S	R_S	N	τ_T	C_{J0}	XTI	BV
1N4004 (F)	3.70E-9	1.76E-2	1.77		1.73E-11	3	400
1N4004 (D)	7.69E-11	4.20E-2	1.45	4.32E-6	3.98E-11	3	400
STTH20R04	1.50E-5	9.72E-3	1.71	1.5E-8	1.55E-10	3	
8EWF	2.35E-12	1.37E-2	1.43	1.78E-7	5.22E-11	2	650
MBR40250	3.00E-6	3.99E-3	1.76	3.48E-8	9.00E-10	5	250

Chapter 20

Switching Losses

Designing an efficient power converter with a fixed budget requires balancing switching losses, conduction losses, and core losses in the magnetic components. The switching frequency is a key factor in balancing these losses. A fixed amount of switching *energy*, E_{sw} is dissipated during each switching cycle. Thus, switching *power*, $P_{sw} = fE_{sw}$, increases with switching frequency f. A magnetic component of a given size, on the other hand, has a loss that decreases with frequency — since the energy stored in the component is proportional to the cycle time $t_{cy} = 1/f$.

In this chapter we examine the switching losses in a half-bridge. We derive an expression for losses that accounts for the three major components of loss: (a) the finite switching time of the device, (b) the reverse recovery charge of the opposing switch, and (c) the parasitic capacitance of the switching node. We also examine the contribution of parasitic inductance to switching loss.

20.1 The Half Bridge

Power stages are usually constructed from pairs of switches that alternately connect one end of an inductor to power or ground as shown in Figure 20.1. In the figure the inductor is modeled as a current source. A pair of switches connected in this manner is called a *half bridge*. Half bridges are used, for example, in the buck converter of Chapter 2, the boost converter of Section 6.1, and the buck-boost converter of Section 6.2. Two such circuits form a *full bridge*, as needed, for example, to realize the switches in the full-bridge converter of Section 9.3.

When current flows in only one direction in the inductor, one switch of the half-bridge can be realized with a diode as shown in Figure 20.2. In the buck configuration, the low switch can be realized with a diode. Similarly, a diode may be used as the upper switch in the boost configuration. If current flows both into and out of the half bridge, as in a full-bridge converter, then both switches must be realized with active switches (such as MOSFETs or IGBTs).

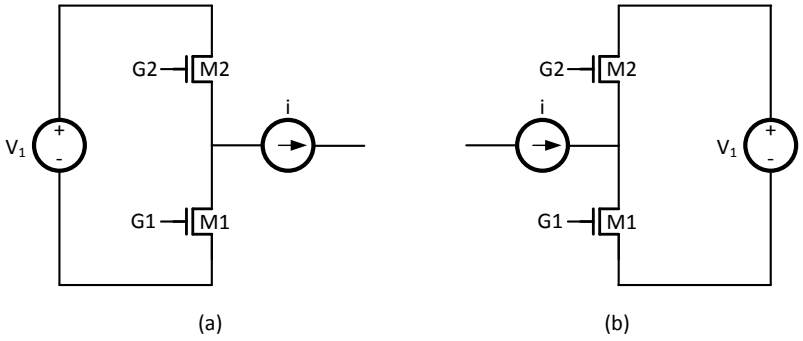

Figure 20.1: Switches are typically used in a half-bridge configuration that connects one end of an inductor (shown as a current source) to power or ground. The current may flow out of the half bridge, as in the buck configuration (a), into the half bridge, as in the boost configuration (b), or alternately into and out of the half bridge, as in a full-bridge converter (not shown).

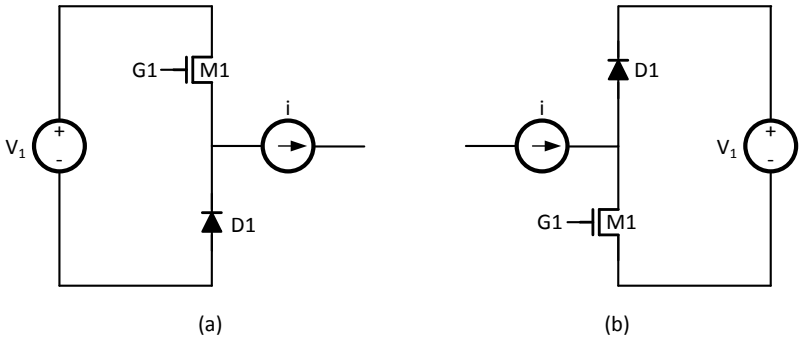

Figure 20.2: If the current in the inductor flows only in one direction, one switch of the half-bridge can be realized with a diode. (a) For the buck configuration — current flows out of the half bridge — the low switch can be realized with a diode. (b) For the boost coniguration — current flows into the half bridge — the high switch can be realized with a diode. (not shown) For situations where current flows both into and out of the half bridge, neither switch can be realized with a diode.

The switch that *opposes* the current source dissipates the bulk of the switching energy. This is the switch that cannot be replaced by a diode in the buck or boost configuration — i.e., the high switch in the buck or the low switch in the boost. In the buck configuration with both switches off, the current flowing out of the midpoint of the half-bridge will discharge the parasitic capacitance of this node, causing it to fall. It will ultimately be clamped to GND by the diode of the low switch. To pull the midpoint high, the upper switch of the buck opposes the current source - which is pulling the node low. Similarly, the current into the midpoint node of the boost causes it to rise. To pull the node low, the low switch of the boost opposes the current.

The low switch of the buck doesn't have to do any work. When the upper switch turns off, it just needs to wait long enough (see this discussion of dead time below) for the switching node to fall before it turns on. It can then switch on, and later switch off, with zero voltage across its terminals — hence dissipating almost no power. The opposing switch, on the other hand dissipates considerable power to perform its transitions — since it must fight the current source.

20.2 Turn-On Transient

The turn-on transition of the "opposing" switch takes place in three steps — labeled t_1, t_2, and t_3 in Figure 20.3. These three steps represent the three sources of switching loss: rise-time loss, reverse-recovery loss, and capacitive loss. In this section we calculate the switching loss during these three steps using a simple model in which current in the opposing device ramps linearly and Q_{RR} is independent of slope. SPICE simulations should be run to accurately calculate loss. However, this simple analytic model we present here is useful because it builds insight about the sensitivity of loss to circuit parameters.

During t_1 the current in the opposing switch rises from 0 to the inductor current I_L. We assume a linear current ramp with slope S, so this period takes $t_1 = I_L/S$. The energy dissipated during this phase is given by multipling the current and voltage curves:

$$\begin{aligned} E_1 &= \int_0^{t_1} V_{DS} I_D dt \\ &= \frac{1}{2} I_L V_{DD} t_1 \qquad (20.1) \\ &= \frac{V_{DD} I_L^2}{2S} \qquad (20.2) \end{aligned}$$

From (20.2) we see that the rise-time energy E_1 can be reduced by ramping the current faster, with a larger slope S.

At the end of the first period, all of the inductor current is flowing through the opposing switch. There is zero current flowing in the diode of the "other"

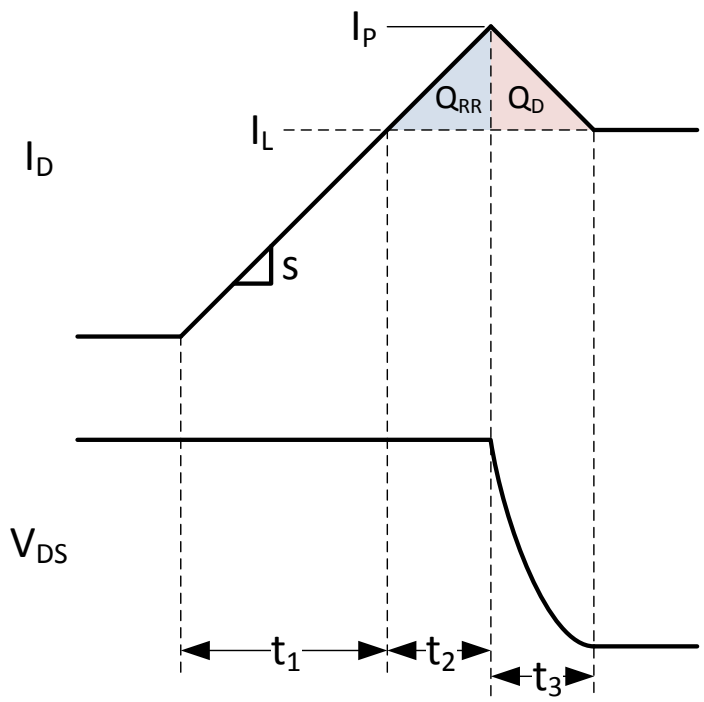

Figure 20.3: The turn-on transient of the "opposing" switch is divided into three regions. During t_1 current in the device ramps to inductor current I_L. The reverse-recovery charge Q_{RR} is cleared from the other switch in t_2. Finally, during t_3 the drain charge Q_D is discharged as the voltage across the opposing switch V_{DS} falls.

switch[1]. However, this diode will remain *on* and conducting until its reverse recovery charge Q_{RR} is cleared. Thus, I_D continues to increase and a reverse recovery current $I_{RR} = I_D - I_L$ flows backwards through the diode. This current continues until the integral of I_{RR} is equal to Q_{RR} which occurs after $t_2 = t_{RR} = \sqrt{2Q_{RR}/S}$. Here we are assuming an abrupt or hard recovery of the diode. In practice we use diodes with a *soft* recovery to avoid large inductive spikes (Section 19.3).

The energy during this reverse-recovery time is given by:

$$\begin{aligned} E_2 = E_{RR} &= \int_{t_1}^{t_1+t_2} V_{DD} I_D dt \\ &= V_{DD} I_L t_2 + \frac{1}{2} V_{DD} t_2^2 S \\ &= V_{DD} I_L t_2 + V_{DD} Q_{RR} \end{aligned} \tag{20.3}$$

The reverse recovery energy has two components. The right term $V_{DD}Q_{RR}$ is the energy needed to remove the reverse recovery charge from the diode. This term is independent of slope S or reverse recovery time t_2[2] The left term $V_{DD}I_Lt_2$ is the power dissipated by the opposing switch carrying the inductor current I_L during the reverse recovery time. Making the reverse recovery time smaller - via a steeper slope S reduces this component. Thus, we see that to first approximation reverse recovery loss is also reduced by making the slope faster.

When silicon diodes are used, the reverse recovery loss, E_2 often dominates the entire switching loss. The use of a SiC diode can nearly eliminate E_2.

During the final period of the turn-on transient the voltage across the opposing FET is discharged. At the start of this period, the diode is off and the current in the opposing switch is $I_D = I_P$. During the discharge we assume the current slopes linearly down to I_L. Thus, the duration of the discharge is $t_3 = 2Q_D/(I_P - I_L)$.

The energy during t_3 is:

$$\begin{aligned} E_3 &= \frac{V_{DD}Q_D}{2} + \frac{V_{DD}I_Lt_3}{3} \\ &= \frac{C_D V_{DD}^2}{2} + \frac{V_{DD}I_Lt_3}{3} \end{aligned} \tag{20.4}$$

The first term is the energy stored on C_D, the drain capacitance[3]. This capacitance includes the parasitic capacitance of the opposing switch, the capacitance of the cathode of the diode, and interconnect parasitics. The second term represents the energy dissipated delivering current to the inductor while the drain is discharged.

[1] Even if the "other" switch is a FET, at this point the FET is off and all current through this switch is flowing through its body diode.

[2] In practice Q_{RR} depends on S, so this term is not exactly constant.

[3] MOSFET data sheets often refer to their contribution to this capacitance as C_{oss}.

We can combine the three phases (20.2), (20.3), and (20.4) to give the overall switching energy:

$$\begin{aligned} E_{SW} &= E_1 + E_2 + E_3 \\ &= \frac{V_{DD} I_L^2}{2S} + V_{DD} I_L t_2 + V_{DD} Q_{RR} + \frac{C_D V_{DD}^2}{2} + \frac{V_{DD} I_L t_3}{3} \end{aligned} \tag{20.5}$$

The model we have presented here has made a number of approximations. The current ramp in the opposing switch isn't perfectly linear. Also, most diodes have a soft recovery — blurring phases 2 and 3. However, this model is accurate enough to estimate the effect of certain design changes on switching loss.

For example, suppose we are considering using a FET with a lower R_{on} to reduce conduction losses. We can accurately esimate the increase in the switching loss due to increased C_D using (20.3). We can also estimate the switching loss due to reverse recovery using Q_{RR} and (20.4). In choosing the size of our gate drive resistor we can consider the effect of rise time on switching loss given in (20.2).

[Scope photo]

[SPICE results for Si and SiC diodes]

20.3 Turn-Off Transient

Significantly less energy is dissipated during turn-off than during turn-on and to first approximation the turn-off transient can be ignored.

Figure 20.4 shows the turn-off transient for the case of a slow switch. In this case, the current source charges the switching node before the switch current reaches zero. This can happen when using a slow switch on a switching node with low capacitance. In the figure, the current where the voltage reaches V_{DD} is denoted I_1. Integrating the product of voltage and current over the interval gives:

$$E_{\text{off}} = V_{DD} t_r \left(\frac{I_L}{6} + \frac{I_1}{3} \right). \tag{20.6}$$

The fast-switch case is illustrated in Figure 20.5. Here the switch current ramps to zero before the current source fully charges the switching node. This can happen with a fast switch, a low-capacitance switching node, and/or a low inductor current, I_L. In the figure the current ramps to zero in time t_c when the swtich-node voltage is V_1. Integrating the product of voltage and current here gives:

$$E_{\text{off}} = \frac{V_1 I_L t_c}{6}. \tag{20.7}$$

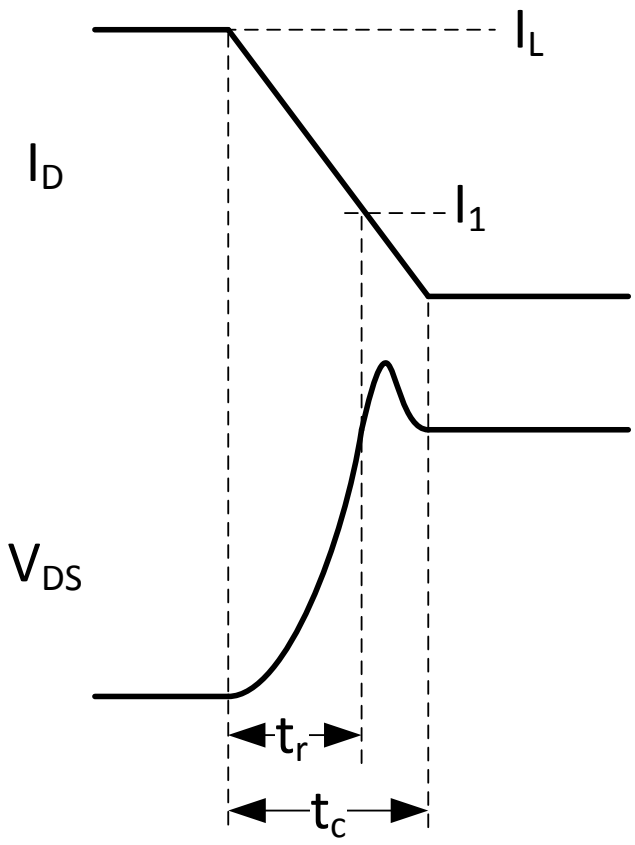

Figure 20.4: The turn-off transient of the "opposing" switch in the case where V_{DS} rises before I_D reaches zero.

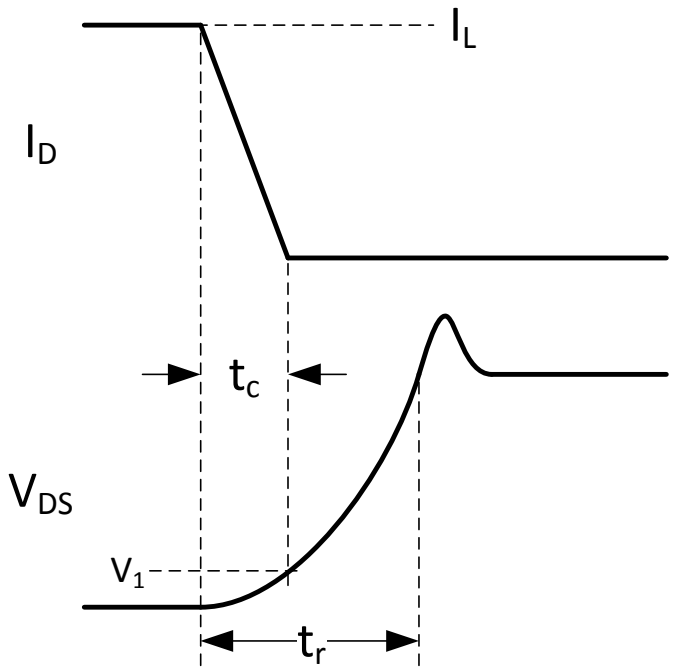

Figure 20.5: The turn-off transient of the "opposing" switch in the case where I_D reaches zero before V_{DS} finishes rising.

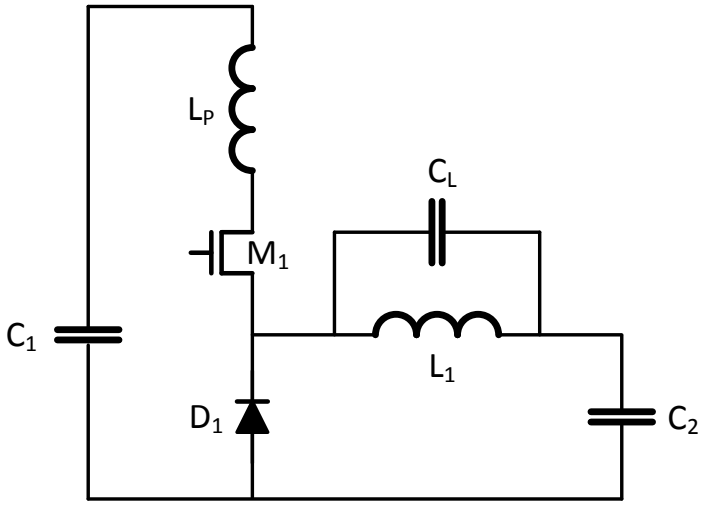

Figure 20.6: Energy stored in parasitic elements L_P and C_L is lost each switching cycle — adding to switching losses.

The energy in the fast-switch case, Figure 20.5, is strictly less than that of the slow-switch case. While turn-off energy is small compared to turn-on energy, it can be made smaller with a faster switching time.

20.4 Parasitic Losses

Parasitic circuit elements can add significantly to switching loses because energy stored in these elments each cycle is generally lost. Figure 20.6 shows a buck converter with the two most significant parasitics, L_P and C_L, explicitly shown. Inductor L_P represents the parasitic inductance around the loop formed by M_1, D_1, and C_1. This inductance includes the series inductance of C_1, the source and drain inductance of M_1, the anode and cathode inductance of D_1, and the inductance of printed circuit board traces connecting these components. With small, surface-mount FETs and diodes fed by multiple small, parallel surface-mount capacitors and careful PCB layout, it is possible to reduce L_P to 10nH or less. On the other hand, with leaded components and careless layout it is not unusual to see parasitic loop inductances as high as 100nH.

Component packaging is a critical factor in parastic inductance. An Infineon IPL65R099C7 FET in a *ThinPAK* surface-mount package has source + drain

capacitance of 3nH. In contrast, an equivalent FET in a leaded TO247 package has source + drain capacitance of 8nH. Surface-mount diodes and capacitors have similar inductance advantages. Placing multiple components in parallel is also effective in reducing inductance. For capacitors in particular, placing several small surface-mount capacitors in parallel gives lower inductance than using one large capacitor. However, the PCB layout needs to be done carefully to avoid adding trace inductance in series with the capacitors. Placing capacitors between the diode and FET usually works well.

Parasitic inductance causes loss because all of the energy stored in this inductance at the peak current level, I_P in Figure 20.3, is ultimately dissipated as the energy in the inductor *rings out*. For example, suppose $I_P = 40A$ and $L_P = 30$nH. The energy lost each cycle due to parasitic inductance is:

$$E_{LP} = \frac{I_P^2 L}{2} \tag{20.8}$$

$$= \frac{(40A)^2(30nH)}{2} \tag{20.9}$$

$$= 48\mu\text{J} \tag{20.10}$$

Reducing the parasitic inductance to 10nH would cut E_{LP} to 16μJ.

Reducing parasitic inductance also helps reduce electromagnetic interference (EMI). As the energy in L_P rings out with C_D it generates high-frequency oscillations that can radiate from the converter — interfering with other circuits.

The series capacitance of the inductor, C_L, is also a source of loss. Because C_2 is large, C_L is to first approximation in parallel with the drain capacitance C_D. Thus, an increased C_L increases E_3 above in (20.4).

One method to reduce C_L is to realize the inductor using several small inductors in series. The small inductors typically have lower parallel capacitance than a large inductor, and placing them in series reduces the inductance further.

Chapter 21

Anatomy of a Switching Cycle

To illustrate some of the pragmatics of using real switches we examine one cycle of operation of a high-voltage IGBT half-bridge driving a current-source load. This scenario is representative of motor controllers, inverters, and high-voltage voltage-converters where the inductive load is essentially a current source.

As we walk through the cycle we will examine the effects of finite switching time and finite on-resistance. We look at the effects of key circuit parasitic elements including the gate-to-drain *Miller* capacitance and of the supply loop inductance. We also examine the effects of the diode reverse-recovery time and the *snappiness* of the diode recovery.

[Where do I show a low-voltage cycle?]

21.1 An IGBT Half-Bridge

Figure 21.1 shows a half-bridge implemented with IGBTs. The half bridge drives a current-source load I_L. IGBTs Q_H and Q_L connect node X to either the positive supply V_s or ground respectively. *Free-wheeling* diodes D_H and D_L provide a current path when the IGBTs are off and clamp node X to the two supplies. Parasitic collector-gate capacitances C_{CGH} and C_{CGL} are explicitly shown as are supply loop inductance L_S and low-side source inductance L_{SL}. The actual circuit has many other parasitic elements that are not shown.

A passive RC snubber consisting of R_X and C_X damps oscillations and slows transitions on node X. The snubber components are not parasitics. They are components that we add to the half bridge to improve its performance.

The gate drive circuits are shown as black-boxes. These circuits drive the gates of each IGBT with appropriate voltages, drive current, and timing. As we shall see, it is critical that the gate drive voltages be referenced directly to the emitter voltage of the corresponding IGBT with no intervening shared inductance. To reduce the tendency for the opposing IGBT to turn-on when

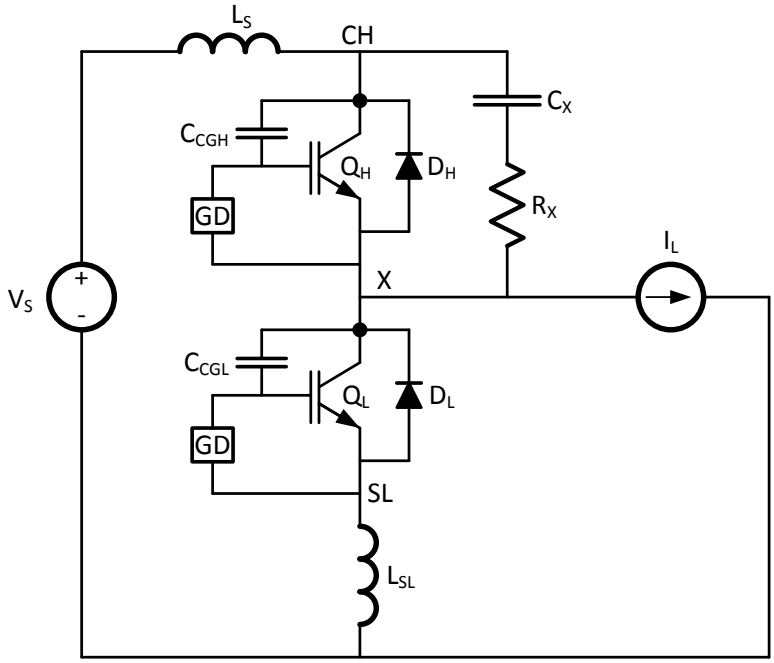

Figure 21.1: An IGBT half-bridge showing key circuit parasitics

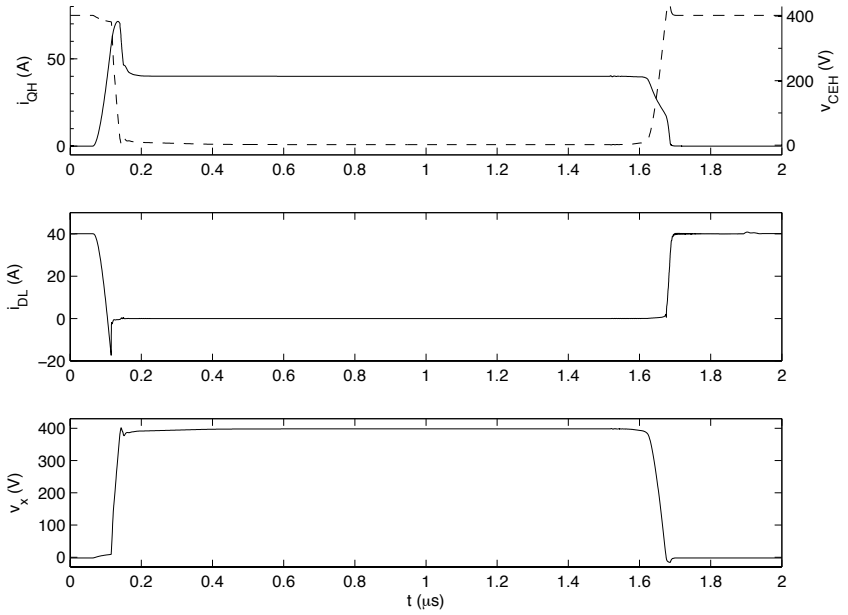

Figure 21.2: Waveforms of key variables during one switching cycle of the IGBT half-bridge.

node X is changing state due to capacitive coupling to the collector, we hold each gate 10V below the emitter voltage with a low impedance when the IGBT is off. We will discuss gate-drive for both high-voltage and low-voltage circuits in more detail in Chapter 22.

In the waveforms shown in this Chapter, we simulate a 16kW half-bridge with $V_S = 400$V and $I_L = 40$A. The IGBTs are Fairchild FGH40N60UF 600V, 40A IGBTs. The diodes are the free-wheeling diodes included in the Fairchild IGBT model. The collector-gate capacitance is 5.8nF and is included in the IGBT model. Inductance L_S =5nH and L_{SL} =5nH. There is additional parasitic inductance, and other elements, in the IGBT model.

The current in the circuit of Figure 21.1 flows primarily in two loops. The *supply loop* charges node X from the supply V_S when Q_H switches on. This loop consists of V_s, L_S, Q_H, Q_L, and L_{SL}. The *output loop* discharges node X through I_L when Q_H turns off and consists of I_L, L_{SL}, and Q_L.

21.2 A Switching Cycle

Figure 21.2 shows key waveforms during one switching cycle of the IGBT half-bridge of Figure 21.1. The high-side IGBT turns on during the time from 0 to 0.2μs, and this device turns off between 1.6 and 1.8 μs.

The top panel shows the current i_{QH} (solid line, left scale) and voltage v_{CEH} (dashed line right scale) of the high-side IGBT. The product of these two waveforms is the power dissipated in the high-side IGBT. The high-side current overshoots to 71A during turn-on due to three factors we discuss in more detail below.

The middle panel shows the current in the low-side diode i_{DL}. This diode carries the 40A load current when the high-side IGBT is off. The low-side IGBT can only conduct current downward, and hence does not contribute to the load current flowing out of the half bridge. The undershoot of this current is due to the diode reverse recovery, one of the three factors contributing to i_{QH} overshoot. The sharp rising edge after the diode recovers (a *snappy* as opposed to *soft* recovery) causes current slope issues as we shall see below.

The bottom panel shows v_X, the voltage at node X. The rise time t_r of v_x is about 25ns and the fall time t_f is about 40ns. This waveform is approximately the inverse of v_{CEH} since $v_{CEH} = v_{CH} - v_X$.

21.3 High-Side Turn-On

Figure 21.3 shows details of the transient associated with the high-side IGBT turning on. The sequence of events starts at t = 15ns (far off the left side of the figure) when the high-side gate driver drives the gate of Q_H to 15V via a 30Ω external resistor and 10Ω of internal gate resistance (in the IGBT model). At t = 70ns the internal gate voltage V_{GEH} (solid line in the bottom panel) crosses the 5V threshold voltage (dotted line) and the IGBT starts to turn on.

From 70 to 115ns, as V_{GEH} continues to slowly rise, i_{QH}, the current in the top IGBT (top panel, solid line) ramps up. As the IGBT current ramps up, low-side diode current i_{DL} ramps down. As the IGBT carries more of the load current there is less current for the diode to carry. At t = 105ns the diode current passes through zero. At this point the IGBT is carrying all of the load current, i_{QH} = 40A.

From 105ns to 115ns, the IGBT current continues to ramp up — above the load current. At the same time the diode current continues to ramp down — below zero. The diode is in its reverse-recovery period t_{rr} = 10ns conducting current backwards while waiting for charge Q_{rr} = 90nC to be swept from the space-charge layer.

At 115ns and a reverse current of 17A the diode recovers and abruptly cuts its current to zero. This *snappy* recovery is a bad thing as it results in a peak current slope of 30GA/s. Via parasitic inductances, this current slope induces a voltage slope of 45GV/s across the low-side IGBT V_{CEL}. This voltage transient is in turn coupled to the gate via C_{CGL} causing the low-side IGBT to briefly

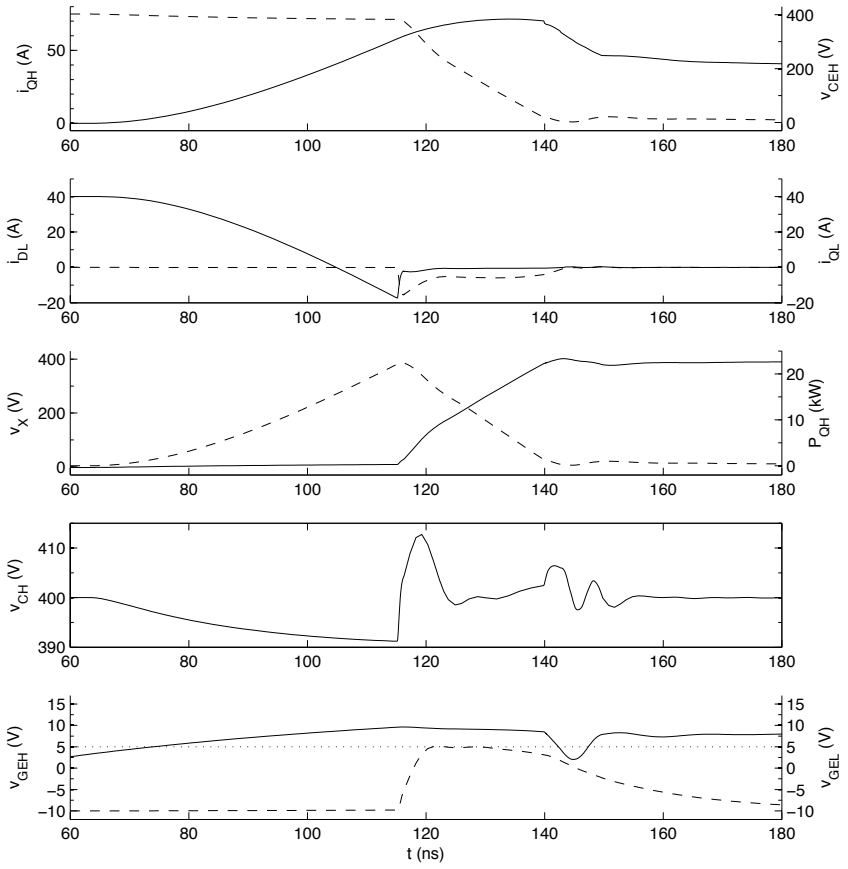

Figure 21.3: Waveforms showing details of the high-side turn-on transient in Figure 21.2.

turn on causing shoot-through current to damp the current and voltage ramps.

The events leading to this shoot-through current can be seen in panels 2 and 5 of Figure 21.3. The bottom panel shows the internal gate voltage of the low-side IGBT (dashed line). Even though the gate driver is holding the external gate voltage 10V below the emitter voltage with low impedance, the capacitive coupling from the collector and the internal 10Ω gate resistance cause the internal gate voltage to rise 15V to turn on the low-side IGBT at 120ns.

The current in the low-side IGBT is shown in panel 2 (dashed line). After an initial spike that is due to capacitive coupling, i_{QL} rises to about 5A at 120ns and then returns to zero at 145ns when the internal gate voltage drops below the threshold.

Using a soft-recovery diode and/or placing a snubber across the low-side IGBT would reduce or eliminate the shoot through caused by large current and voltage ramps due to diode recovery.

After the diode recovers at $t = 115$ns, V_X starts to rise and reaches 400V at $t = 140$ns. There is a slight overshoot due to inductance in the circuit. Current i_{CH} remains at its peak value during this phase as the parasitic capacitance of node X is charged.

Power dissipation in Q_H is shown in panel 3 (dashed line). The power ramps up as i_{CH} increases from $t = 70$ns to 115ns. It reaches a peak value of 22kW at the point where the diode turns off. Power in Q_H then ramps down as V_X rises from t =115ns to 140ns. The total area under the curve is the turn-on switching energy of 773μJ.

The 30A current overshoot in i_{CH} is due to three components. First, the diode reverse-recovery current (panel 2, solid line) reaches about 20A before the diode abruptly reovers. Second, because circuit inductance prevents this current from changing abruptly, Q_L is forced to turn on just enough so that its shoot through current (panel 2, dashed line) almost exactly replaces the diode current before ramping to zero. Finally, additional current is needed to charge the parasitic capacitance of node X as V_X rises.

Panel 4 shows the voltage on node CH. Initially this node droops as the slope of i_{CH} induces a voltage across L_S and L_{SL}. Then, when the diode switches off, the voltage overshoots due to the diode recovery transient. It then ramps up slightly as i_{CH} returns to its steady state value. There is a small amount of ringing before it stabilizes back at 400V.

As discussed in Chapter 22, we set the slope of i_{QH} with our choice of gate driver pull-up resistor, 30Ω in this case. A smaller resistor would give a faster slope, while a larger resistor would give a slower slope. This slope in turn affects the peak reverse recovery current. A faster slope would result in a larger peak reverse recovery current, but a smaller reverse recovery time. A gate-drive pull-up resistor of 10Ω , for example, gives I_{prr}=25A, $t_{rr} = 10$ns, and $Q_{rr} = 120$nC. The gate drive resistor and snubber values are chosen to minimize switching losses and reduce ringing in the circuit.

21.4 High-Side Turn Off

Figure 21.4 shows the high-side turn-off transient. The sequence of events begins with the gate (not shown) falling at 1.52μs, 300ns off the left side of the graph. Over 1μs elapses before the IGBT reacts by reducing i_{QH} (top panel, solid line) at 1.62μs. This large turn-off delay does not itself result in losses. However, because the turn-on delay is much shorter, a large *dead time* is required to avoid shoot-through losses.

As soon as i_{QH} starts falling at 1.62μs, V_{CEH} (top panel, dashed line) starts rising as the difference current $i_L - i_{QH}$ discharges node X. The product $P_{QH} = i_{QH} \times V_{CEH}$ is the power dissipated in Q_H and is shown in panel 3, dashed line. This power peaks at 7.1kW and the area under the curve is a switching energy of 230μJ. For this particular configuration, the turn-off energy is significantly smaller than the turn-on energy. The total switcing energy for the cycle is 1mJ.

As V_{CEH} rises, V_X falls reaching zero about 1.675μs. At this point the diode starts conducting and ramps its current rapidly to 40A in about 15ns. This rapid current transient causes V_X to undershoot by 15V due to voltage induced in L_{SL}. V_{CH} peaks by 15V at the same time due to the current transient in L_S as i_{QH} abruptly falls to zero. This transient is relatively mild due to the tiny 5nH L_S With higher inductances, the overshoot of V_{CEH} during turn-off can be large enough to damage the IGBT and requires countermeasures to protect the device.

If V_X falls too rapidly (V_{CEH} rises too rapidly) the C_{CG} of Q_H can drive the gate high, turning Q_H back on. With the relatively slow fall of V_X in this scenario, coupled with the long turn-off delay and the fact that we are driving the gates to -10V in the *off* state, this transient turn-on is not an issue.

21.5 Exercises

Repeat simulation using MOS FETs

Repeat simulation using SiC FETs

Repeat simulation with 100nH of L_S and L_{SL}.

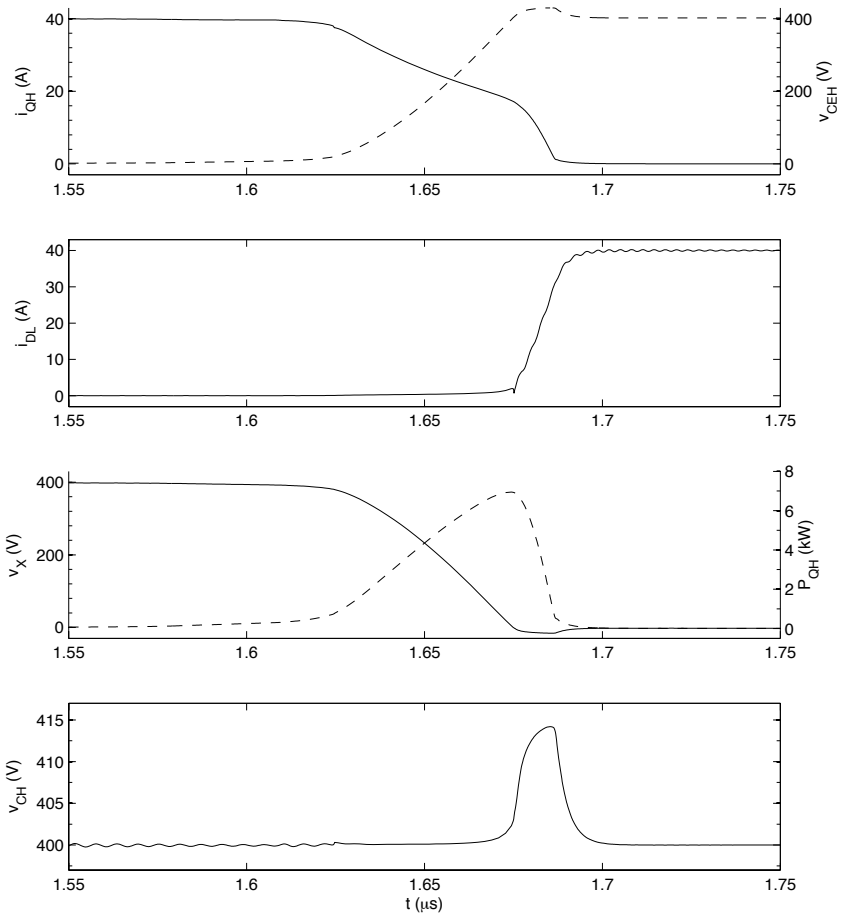

Figure 21.4: Waveforms showing details of the high-side turn-off transient in Figure 21.2.

Chapter 22

Gate Drive

basic
- protection
- dead time
- supply

22.1 Gate Drive Circuit

Figure 22.1 shows a typical gate driver. To first approximation a gate driver is a pair of switches S_H and S_L that connect the gate of MOSFET M_1 to either its source V_S, to turn the FET off, or to a voltage V_{GH} above the source, to turn the FET on. This switching is performed under control of an input signal, *in*. Gate drive resistors R_{GH} and R_{GL} control the speed of the rising and falling transitions respectively. In some gate drivers, the right side of the two switches are tied together and a single resistor R_G controls both rising and falling transition times. In some gate drivers, the low switch drives the gate to a voltage V_{GL} below the source V_S to make the FET more resistant to inadvertent turn-on due to drain-gate feedback. While the gate driver here is shown driving a MOSFET M_1, the same function is used to drive the gate of an IGBT.

High currents are needed to switch large MOSFETs or IGBTs quickly. It is not unusual for high performance gate drivers to have drive currents of 10A or more. The V_{GH} supply must have a bypass capacitor close to the gate driver and inductance in the gate-drive *loop* must be minimized to allow this high drive current to be delivered in a few nanoseconds.

The circuitry contained in the dashed box is typically packaged in a monolithic gate-driver integrated circuit. The box labeled "control and protection" translates the input signal to the levels needed to control the switches. Depending on the gate driver IC, it may also perform additional functions such as:

Level Shifting and Isolation: A high-side switch driver must translate the input signal to a level referenced to the source of the high-side device —

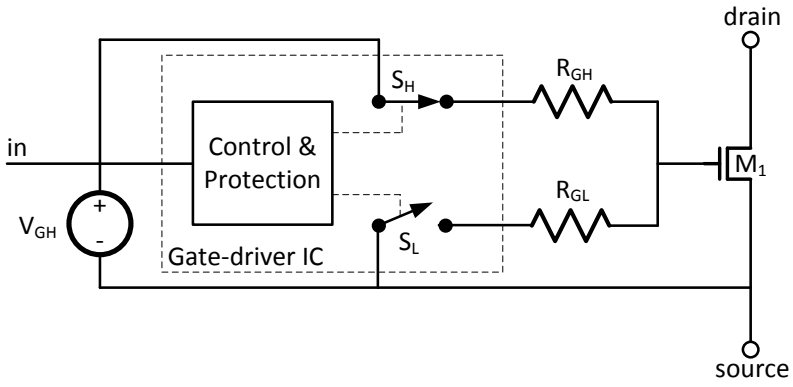

Figure 22.1: A basic gate drive circuit switches the gate voltage of a MOSFET between the source voltage V_S and a positive supply voltage $V_S + V_{GH}$. Gate drive resistors R_{GH} and R_{GL} control the speed of switching. The circuitry contained in the dashed lines is typically packaged in a monolithic gate-drive IC.

that may be 100s of volts higher. To provide noise immunity the gate driver may also provide optical or magnetic isolation of the signal.

Under-Voltage Lockout (UVLO): The gate remains low until the gate drive supply V_{GH} exceeds a threshold $V_{GH\min}$. This prevents FET M_1 from turning on with inadequate gate voltage which could lead to higher than normal on-resistance and ultimately failure due to overheating.

Input Filtering (IF): The input signal is filtered to reject narrow pulses in either direction. This prevents noise from M_1 switching from feeding back and causing the gate to toggle - resulting in excessive switching losses.

Over-Current Protection (OCP): The gate is pulled low if the source current of M_1 exceeds a threshold $I_{S\max}$. This protection requires an external current sensor, typically a current sense resistor in series with the source.

Short Circuit Protection (SCP): The gate is pulled low if the drain of M_1 remains above a threshold $V_{DL\max}$ for longer than an interval $t_{DL\max}$. This protects against the case where the drain is shorted to the positive supply and hence cannot be pulled low.

With either OCP or SCP, once the fault is detected, the gate driver may be disabled until the fault condition is manually reset. This prevents repeated occurances of the condition from stressing the MOSFET.

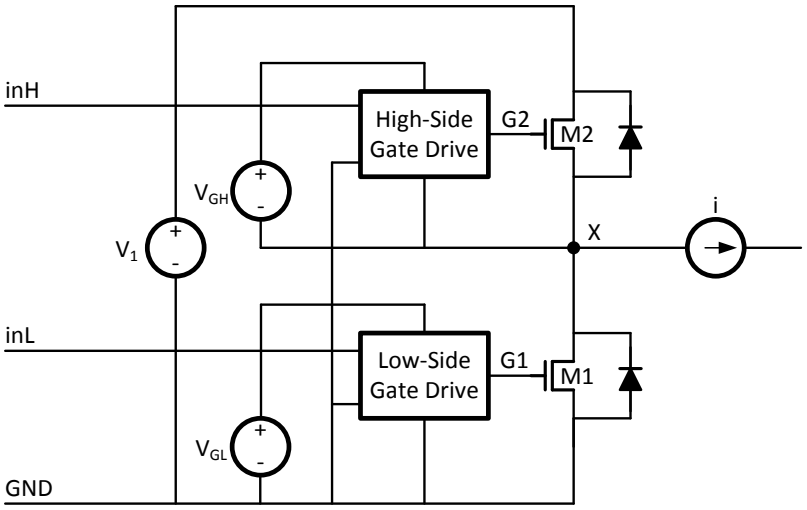

Figure 22.2: A pair of gate drivers are used to drive the two FETs in a half-bridge. *Dead time* must be left between the high periods of the two input signals to avoid *shoot through* current.

The gate-driver supply V_{GH} is a critical component of the gate drive. Both MOSFETs and IGBTs are usually driven with supplies ranging from 12-15V. GaN FETs cannot tolerate such high voltages and must limit their gate drive to about 5V.

22.2 Dead Time

Figure 22.2 shows a typical application in which a pair of gate drivers drive the two switches of a half-bridge. The pair of gate drivers accepts logic-level (typically 0-5V) signals **inH** and **inL** and generates gate-drive (typically 0-15V high-current) signals **G1** and **G2**. The high-side gate drive signal is referenced to the mid-point of the bridge, signal **X**. The low-side gate driver is a buffer and amplifier, translating the 0-5V **inL** to the 0-15V **G1** and increasing current drive. The high-side gate driver must also be a level shifter, translating **inH** referenced to GND, to **G2** referenced to **X**.

The gate-drive supply is referenced to the source of the mosfet V_S. Thus, for a high-side gate driver this is a *flying* supply that rides up and down with the midpoint of a half-bridge. When **X** is high, the source of M_2 is at the positive supply V_1 and the high-side supply is above this point, at $V_1 + V_{GH}$. We discuss

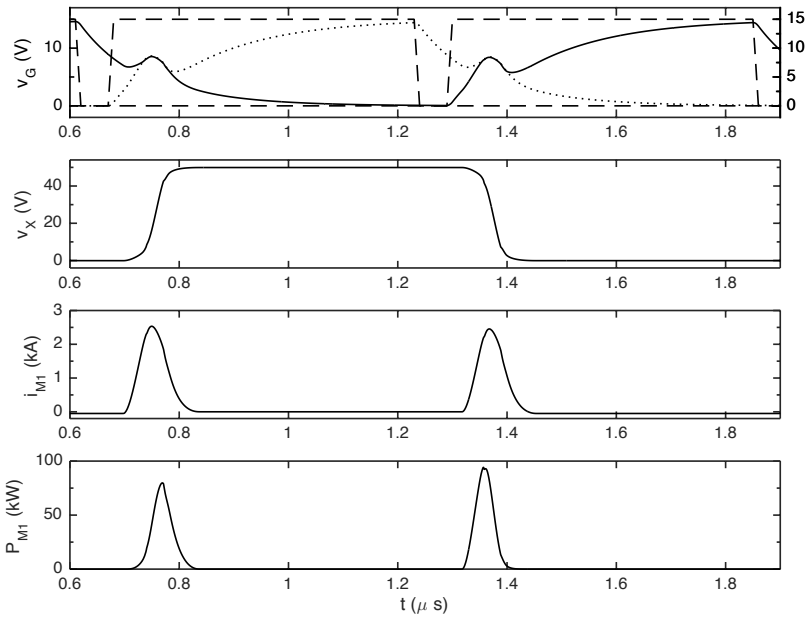

Figure 22.3: Waveforms illustrating a half-bridge driven with too little dead time. From top to bottom: gate drive voltages, mid-point voltage, current in low-side FET M_1, and power in low-side FET M_1. Gate drive voltage at the gate drivers is shown dashed, at the low-side gate dotted, and at the high-side gate solid. With inadequate dead time, overlap current reaches 2,500 A and instantaneous power approaches 100kW.

methods to realize this flying supply below.

The dynamics of the gate-drive cause MOSFETs to turn on faster than they turn off. Thus, to avoid the overlap or *shoot-through* current that occcurs when both FETs in a half-bridge are on simultaneously, sufficient *dead time* must be provided between the two gate drive signals **inH** and **inL** to allow one FET to fully turn off before the other FET turns on. With dead time, it pays to err on the side of caution. There is a very small penalty for having too much dead time. The consequences of having too little dead time are catastrophic.

Figure 22.3 shows waveforms from a SPICE simulation of a half-bridge with too little dead time. The half-bridge is realized with two IRLB3036PBF 60V 2mΩ MOSFETs with a 48V supply and a 50A load. The gate drive at the gate driver has 50ns of dead time between **G1** falling and **G2** rising and vice versa.

The top panel illustrates the dynamics of gate drive that lead to shoot-through.

While the gate drive voltage at the gate drivers (dashed lines) has 50ns of

dead time, the voltage at the gates themselves (solid and dotted lines) do not. This is the voltage measured inside the SPICE model for the FET after both the gate drive resistor and the resistance of the gate itself[1]. The voltage on this internal gate node reflects the actual state of the device.

The asymmetry between the high and low gate drive voltages and the threshold voltage V_T results in an asymmetry between turn-on and turn-off times. With $V_T \approx 5$V, turning off the FET requires discharging the gate node through about 10V (from 15V to 5V), while turning the FET on only requires charging the gate by 5V (from 0V to 5V). Hence the turn-on takes about half the time as turn-off. Even with a 50ns head start, the turn-off of the low-side FET is not complete when the high-side FET turns on at 0.7μs.

The result of both FETs being on simultaneously is catastrophic current draw. The current in the two FETs reaches 2,500A and the power dissipation in just one of the FETs reaches nearly 100kW[2] The width of the current and power puse is about 100ns. Even with much smaller amounts of overlap, overheating and device failure would quickly result.

Figure 22.4 shows waveforms for driving the same half-bridge with adequate dead time. Here the dead time from **G1** falling to **G2** rising is 500ns and the dead time from **G2** falling to **G1** rising is 300ns. The asymmetry here is due to the asymmetric rise and fall times of node **X** due to the 50A current source.

At the start of the figure, the high-side gate drive **G2** (dashed line) falls. The high-side gate (dotted line) follows with an RC decay and drops through $V_T = 5$V about 110ns later. Shortly after that, node **X** falls as the 50A current source discharges its parasitic capacitance. The gate voltage flattens out as **X** falls due to parasitic drain-gate capacitance. Relative to the gate, the drain is rising and feeding charge onto the gate node via C_{GD}.

The low-side gate-drive **G1** rises 300ns after **G2** fell. Even though the high-side gate is still about 3V, this is safe because node **X** has completed its fall. The low-side gate passes through $V_T = 5V$ about 50ns later and there is no overlap current. In fact this switching event is a non-event. Node **X** is already at ground and the body diode of M_1 is already carrying all of the output current. M_1 switching on occurs with no voltage across M_1. That is, it is *zero-voltage switching* or ZVS.

The rising transition of node **X** is a bit more dramatic because of reverse-recovery losses. Here we wait 500ns from **G1** falling before **G2 rises**. At this point the low-side gate (solid line) has completely discharged. When M_2 turns on, it must first source the 50A output load. It then continues to ramp its current to supply the reverse recovery charge of the low-side FET's body diode. The large current spike shown in the bottom trace is reverse-recovery, not overlap current. By the time the reverse recovery process is over a large amount of current is flowing in M_2 Thus node **X** rises very quickly. This rapid rise causes drain-gate feedback that tries to turn-on M_1 and turn off M_2. Because the gate

[1] You cannot probe this gate node on a real device. The node you can probe - after the gate drive resistor but before the gate resistance is intermediate between the two waveforms.

[2] In a real device such high currents would not occur because the device would fail and act as a fuse before the current reached the peak level.

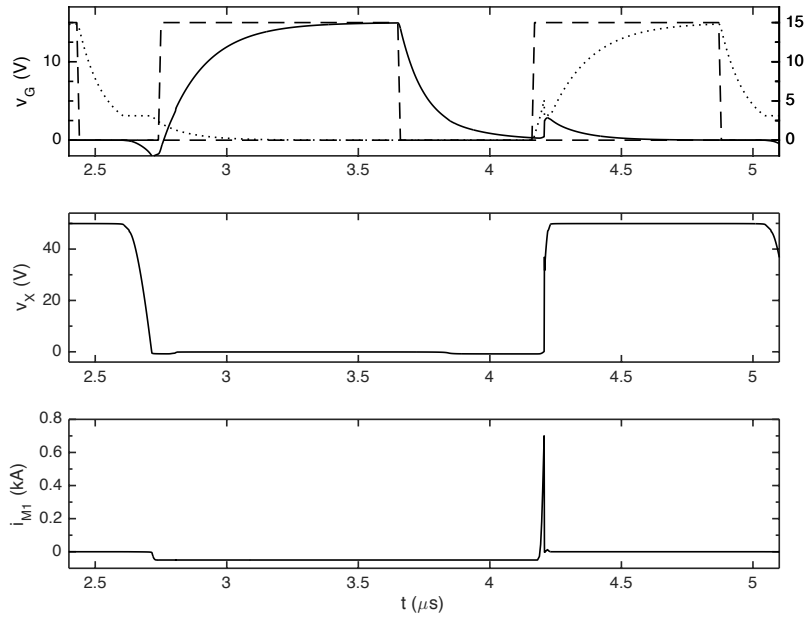

Figure 22.4: Waveforms illustrating a half-bridge driven with adequate dead time. From top to bottom: gate drive voltages, mid-point voltage, current in low-side FET M_1. The spike in I_{M1} here is due to reverse-recovery, not overlap current.

of M_1 started fully discharged, it remains off despite the best efforts of the rising drain to turn it on via C_{GD}. The gate of M_2 dips breifly due to C_{GD} but it quicky recharges.

The reverse recovery current spike reaches 700A and has a width of about 150ns — $Q_{RR} \approx 50\mu$C. This current spike causes about 300μJ of reverse recovery loss in FET M_2.

The downside of requiring such large dead times to avoid overlap or shoot-through current is that it limits the maximum operating frequency of the half-bridge. The half-bridge shown here with 300ns and 500ns dead times — and comparable on times — can't be operated much above 500kHz. Of the 300μJ of reverse recovery losses may limit operating frequencies to something even lower depending on required efficiency and available cooling.

In practice, dead time should be set to a value larger than the minimum shown here to avoid shoot-through current in the presence of component variations. The only downside to increasing deadtime, in addition to a lower maximum operating frequency, is that conduction loss is slightly increased because of the diode drop incurred during the period that the body diode of the *opposing* switch is carrying current before and after the *opposing* FET is on.

22.3 High-Side Supplies

The high-side gate driver in Figure 22.2 requires a high-side supply V_{GH} that is referenced to the switching node, **X**. There are two common ways to realize this high-side supply: using a bootstrap supply, or using an isolated DC-DC converter.

A bootstrap supply, shown in Figure 22.5, uses a flying capacitor charged from the low-side gate drive supply to provide power to the high-side gate driver. When switching node **X** is low, capacitor C_B is charged to $V_{GL} - V_D$, via resistor R_B and diode D_B. V_D is the drop across diode D_B.

When switching node **X** rises, diode D_B turns off, isolating node **V**. Node **V** rises with **X**, giving it a voltage of $V_S + V_{GL} - V_D$. The high-side gate driver is supplied by charge stored on C_B. This charge is restored by recharging C_B the next time node **X** is low.

Capacitor C_B must be sized large enough so that it looses negligible voltage between *charges*. Resistor R_B is sized to limit current during charging and also to limit reverse-recovery current when node **X** rises. Diode D_B must have an extremely fast recovery. For converters with extremely fast rise times on node **X** a SiC diode may be required on the bootstrap supply to reduce reverse recovery losses.

The bootstrap supply is attractive for its simplicity and low component count. However, it has a few limitations. First, it only works if node **X** returns low sufficiently frequently to keep C_B charged. Consider, for example, the case of a battery charger driving a 400V battery. In the idle state (no current in the inductor modeled by the current source) node **X** is at 400V, diode D_B is off, and

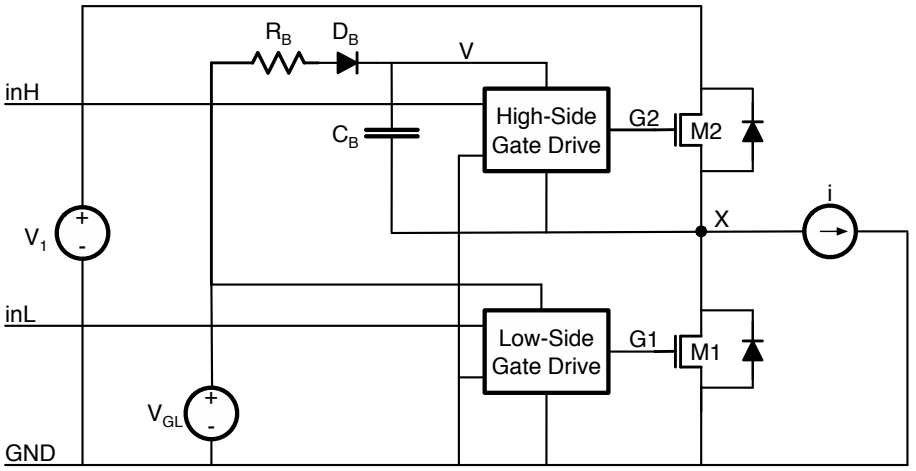

Figure 22.5: A half-bridge gate driver with a *bootstrap* high-side supply. When the switch node **X** is low, capacitor C_B charges from V_{LG} via resistor R_B and diode D_B. When switch node **X** is high, C_B supplies the high-side driver.

C_B is uncharged[3]. In cases like this, or other cases where one cannot gurantee that **X** will be driven low sufficiently often, the bootstrap supply cannot be used.

Another limitation of the bootstrap supply is that it requires that the converter stay under a maximum duty factor. If node **X** is low for too short a time each cycle, C_B will not have adequate time to recharge. If a converter is required to operate with a duty factor approaching 100%, it cannot use a bootstrap supply.

In cases where a bootstrap supply is unsuitable, the high-side supply can be provided by an isolated voltage converter as shown in Figure 22.6. The isolated converter accepts V_{GL} from the low-side supply as its input and provides an isolated output at the same voltage, but referenced to node **X**. With a conversion ratio of 1:1, $V_V = V_X + V_{GL}$.

Any of the converter topologies in Chapter 9 can be used to provide a high-side supply. A flyback converter is often used because of its simplicity. When multiple isolated supplies are needed, a multi-output flyback can be realized by winding multiple secondaries on the same core as illustrated in Figure 22.7. In a three-phase inverter, for example, the three high-side switches need independent isolated supplies. These can be provided by a single flyback converter with three secondary windings[4].

If a bipolar gate drive (driving the gate both above and below the source

[3]With the arrangement shown in Figure 22.5, one could always turn on M_1 to pull down **X**. However, this leads to reverse current in the inductor. Also, in many situations the low side switch will be a diode which is unabled to pull down **X** until current builds in the inductor.

[4]In some inverters separate isolated supplies are provided for the low-side switches as well.

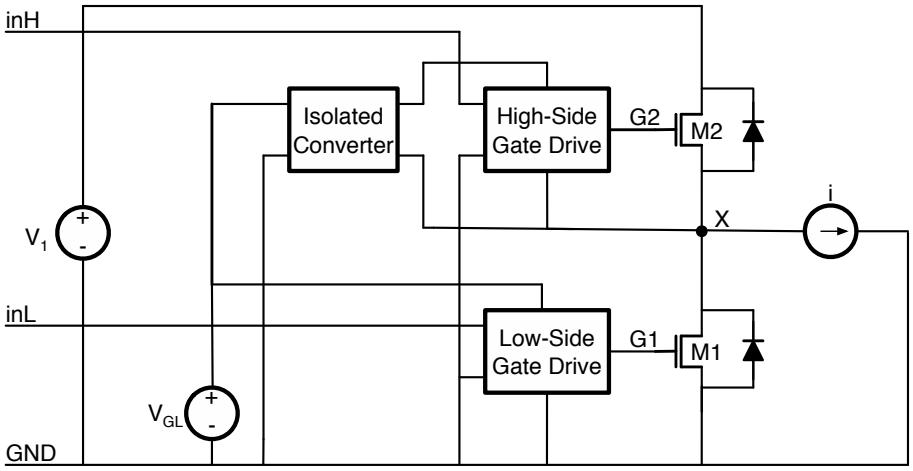

Figure 22.6: An isolated converter (Chapter 9) can also be used to supply power to the high-side gate driver. This approach is required if the switching node cannot be guaranteed to be at ground potential often enough to charge a bootstrap supply.

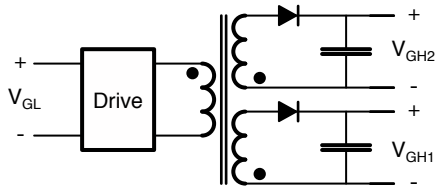

Figure 22.7: Several isolated high-side supplies can be realized with a single multi-winding transformer. A separate secondary winding is provided for each output. A multi-output flyback is shown here, but this approach can be applied to any of the isolated converters of Chapter 9.

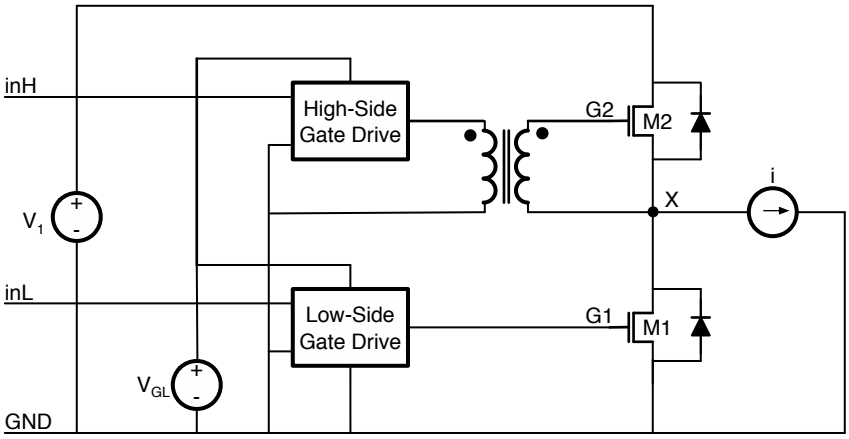

Figure 22.8: The high-side driver can be referenced to the low-side supply and its output level shifted with a transformer. However, it is difficult to maintain volt-seconds balance across the transformer with this approach.

voltage) is needed, the secondary windings of the flyback converter can be center tapped to produce voltages both above and below V_X.

22.4 Transformer Gate Drive

An alternative to using a gate drive with a floating supply is to keep the driver referenced to system ground, but to level shift the output of the driver using a transformer as shown in Figure 22.8.

While conceptually simple, this approach suffers from the inability of transformers to pass DC voltages — and hence the inability to pass any waveform with a DC component. It is very difficult in this configuration to maintain volt-seconds balance across the transformer — which is required to keep the transformer out of saturation. With an equal bipolar gate drive (e.g., +15V and -15V) and a 50% duty factor volt-seconds balance is acheived, but with unidirectional drive or an arbitrary duty factor, balance is not achieved and the transformer will eventually saturate.

22.5 Transient Immunity

A key property of a gate driver is its *transient immunity*. This is the ability of the level shifter part of the gate driver to operate properly in the presence of a fast transient on node **X**. Transient immunity is expressed in units of V/s. For example, if node **X** swings through 500V in 5ns, the gate driver requires a transient immunity of 10^{11}V/s or 100 V/ns. A typical optically-isolated gate

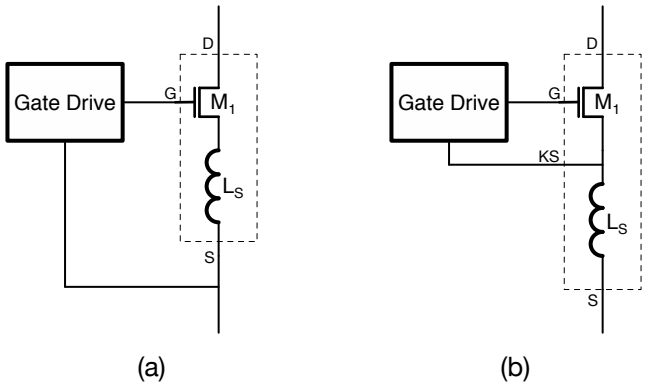

Figure 22.9: (a) A conventional FET has its gate drive affected by voltage across the source inductance. (b) with a *Kelvin* source connection the gate drive is isolated from the source inductance.

driver has transient immunity of 25 V/ns and would not be able to tolerate such a fast tranient.

22.6 Source Reference

Some FETs are packaged with a separate *Kelvin* source terminal for to be used as a reference for the gate driver. This isolates the gate driver reference from voltage induced across the source inductance.

Figure 22.9 shows a gate driver connected to a conventional 3-terminal FET (Figure 22.9(a)) and to a 4-termial FET with a Kelvin gate connection (Figure 22.9(b)). With the configuration in Figure 22.9(a), when the FET switches on, its current ramps up rapidly inducing a voltage across L_S that reduces the gate drive $V_{GS} = V_G - L_S di/dt$ (where V_G is the gate driver voltage. When the FET turns off, the opposite occurs. As current ramps down, a negative voltage is induced across L_S turning the FET back on. With high transient currents this negative feedback through the source inductance can significantly slow turn-on and turn-off.

With the Kelvin source connection shown in Figure 22.9(b), the gate drive is isolated from the source inductance and turn-on and turn-off are not slowed by voltages induced across the inductor.

There is inductance on the Kelvin terminal of the FET. However, the source current of the FET does not pass through this inductance and hence no significant voltage is induced. Only the gate-drive current passes through the Kelvin terminal.

22.7 Protection

Gate drive integrated circuits often include *protection* circuitry to turn the FET off in the event that a parameter is deteted outside of the normal operating *envelope*. Conditions that are checked include:

Over Current: If the FET current is sensed to exceed its maximum rating by a significant margin.

Over Temperature: If the FET temperature is sensed to exceed its maximum rating.

Short Circuit: If V_{DS} remains above a threshold V_{SC} (typically just a few volts) for more than a period t_{SC} after the FET is turned on, the FET output is detected to be shorted.

Over Voltage: When the FET turns off, if V_{DS} exceeds the maximum rating of the FET.

For the first three faults, the fault is managed by turning the FET off. This action may be transient or semi-permanent. For example, when an over-current condition is detected the FET is turned off for the remainder of this PWM cycle. However, it may be allowed to turn-on again on teh next PWM cycle (a transient response affecting just one PWM cycle), or the FET may be disabled until the fault condition is reset by control software (a semi-permanent response).

For the last fault (over voltage), the fault is managed by turning the FET back on - slightly. This discharges the drain node - reducing voltage, and absorbs the inductor current that caused the over-voltage condition. Over voltage protection is typically used for IGBTs rather than FETs. The avalanche body diode in a FET provides over-voltage protection without need for action by the gate driver.

22.8 Exercises

Design a gate driver using discrete FETs
Simulate a gate driver usign LT Spice
With bootstrap supply
Calculate minimum capacitance for a bootstrap supply

Chapter 23

Magnetic Components

Magnetic components, inductors and transformers, are key components of most voltage converters and most green electronic systems. They often account for a sizeable fraction of the system cost and their characteristics often constrain key parameters of the system such as switching frequency. The core losses and copper losses in magnetic components are usually a significant fraction of total losses in these sytems.

In this chapter we review the basics of magnetic circuits and the relationship between magnetic circuits and electric circuits. We build on this material to discuss the design of inductors in Chapter ?? and transformers in Chapter ??.

23.1 Magnetic Circuits

Magnetic components involve *magnetic flux* ϕ flowing through magnetic circuits in response to a *magnetomotive force* (MMF) $\mathcal{F}$. The *reluctance* $\mathcal{R}$ of the circuit opposes the flow of flux so that the resulting flux is given by:

$$\phi = \frac{\mathcal{F}}{\mathcal{R}} \tag{23.1}$$

Equation (23.1) is analogous to Ohm's law, but for magnetic circuits. Here flux ϕ (in Webers) is analogous to current, magnetomotive force $\mathcal{F}$ (in Amperes[1]) is analogous to voltage, and reluctance $\mathcal{R}$ (in Amperes/Weber) is analogous to resistance.

The reluctance of a magnetic path, $\mathcal{R}$, is determined by the length of the path l, its cross-sectional area A, and the permeability of the material making up the path μ.

[1]Sometimes magnetomotive force is referred to in units of Ampere-turns. However, what matters is how many Amperes pass through the core. It doesn't matter how many parallel wires carry this current.

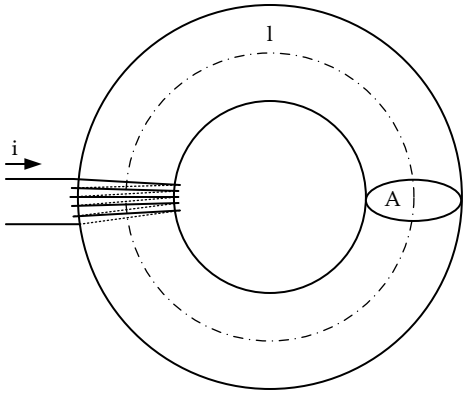

Figure 23.1: An example magnetic circuit. A toroidial core with centerline length l and cross sectional area A made from a ferrite with $\mu_r = 10^3$. A magnetomotive force $\mathcal{F} = Ni$ is applied to the core by applying a current i to N turns of wire around the core.

$$\mathcal{R} = \frac{l}{A\mu} \tag{23.2}$$

If we think of reluctance as being analogous to resistance, doubling the area is the equivalent of putting two of the original paths in parallel — hence halving the reluctance. Doubling the length of the path is like putting two paths in series — doubling the reluctance. Carrying our analogy further, permeability μ is analogous to conductivity. The higher μ, the lower $\mathcal{R}$.

The permeability of free space is $\mu_0 = 4\pi \times 10^-7$. We often refer to the relative permability of a material μ_r where:

$$\mu = \mu_r \mu_0. \tag{23.3}$$

Typical magnetic materials have relative permeabilities μ_r ranging from 10^2 to 10^4.

Consider the example magnetic circuit of Figure 23.2 with $l = 0.1$m and $A = 10^{-4}\text{m}^2$. The reluctance $\mathcal{R}$ of this core is

$$\begin{aligned}\mathcal{R} &= \frac{0.1}{(10^{-4})\,(10^3)\,(4\pi \times 10^{-7})} \\ &= 7.96 \times 10^5.\end{aligned}$$

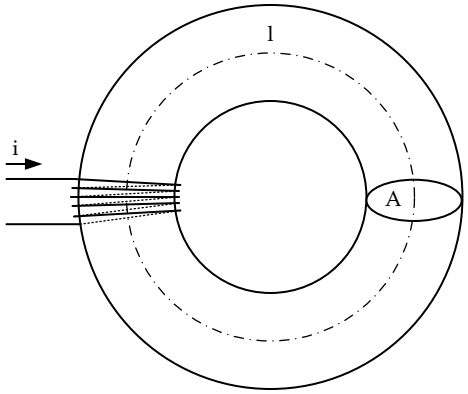

Figure 23.2: An example magnetic circuit. A toroidial core with centerline length l and cross sectional area A made from a ferrite with $\mu_r = 10^3$. A magnetomotive force $\mathcal{F} = Ni$ is applied to the core by applying a current i to N turns of wire around the core.

We can create a magnetomotive force $\mathcal{F}$ by passing current through turns of wire about a core. Consider the *core* of Figure 23.2 with N turns of conductive wire wound through it. The magnetomotive force $\mathcal{F}$ is given by:

$$\mathcal{F} = NI \tag{23.4}$$

The flux produced in response to the current is given by substituting (23.4) into (23.1):

$$\phi = \frac{Ni}{\mathcal{R}} \tag{23.5}$$

$$= \frac{NiA\mu}{l} \tag{23.6}$$

For example, consider our torus of Figure 23.2 with N=5 turns of wire carrying $i = 1$A of current. The magnetomotive force is $\mathcal{F} = 5$A and the flux is $\phi = \frac{5}{7.96\times10^5} = 6.28 \times 10^{-6}$ Webers.

The *flux density* or *B-field* in a magnetic circuit is the flux per unit area:

$$B = \frac{\phi}{A} \tag{23.7}$$
$$= \frac{Ni}{A\mathcal{R}} \tag{23.8}$$
$$= \frac{Ni\mu}{l} \tag{23.9}$$

From (23.9) we see that the flux density B induced by a given current Ni is independent of core area. Flux density B has mks units of *Tesla*.

For our example core of Figure 23.2 with 5A of MMF, the flux density is $B = \frac{\phi}{A} = \frac{6.28\times10^{-6}}{10^{-4}} = 6.28 \times 10^{-2}$Tesla.

The *magnetic field* or *H-field* in a magnetic circuit is given by

$$H = \frac{B}{\mu} \tag{23.10}$$
$$= \frac{Ni}{l} \tag{23.11}$$

Where we derive (**??**) by substituting (**??**) into (**??**). Manetic field H has units of A/m.

To carry our electrical analogy futher note that H and $\mathcal{F}$ are related by

$$\mathcal{F} = Ni = Hl \tag{23.12}$$

Thus magnetomotive force $\mathcal{F}$ is the line integral of the magnetic field H just as voltage V is the line integral of the electric field E.

For our example magnetic circuit of Figure 23.2 with 5A of MMF applied to a core with a length of 0.1m, the magnetic field is 50A/m.

23.2 Equivalent Magnetic Circuits

To solve for ϕ with more complex core geometries we often invoke our electrical analogy and draw an *equivalent electrical circuit* for our magnetic circuit. An equivalent circuit for the simple magnetic circuit of Figure 23.2 is shown in Figure 23.3.

More complex geometries result in multiple resistors that represent different segments of the magnetic circuit. Consider, for example the *gapped* core shown in Figure 23.4(a). The toroidal core of magnetic material ($\mu_r = 1000$) with path length l_c is interrutped with an air gap of length l_g. Gapped cores are frequently used for inductors, as we shall see in Chapter 26, because storing energy as flux in the air gap avoids saturating the magnetic material.

We model this magnetic circuit as the series combination of the core reluctances $\mathcal{R}_c$ and the gap reluctance $\mathcal{R}_g$. Applying (23.2) we can write:

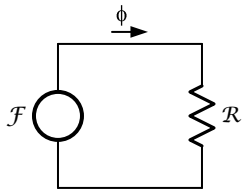

Figure 23.3: We can model our example magnetic circuit with an equivalent electrical circuit. The magnetomotive force is modeled as a voltage source with magnitude $\mathcal{F}$. The core is modeled as a resistor with value $\mathcal{R}$. The current that flows in the electrical circuit represents the flux in the magnetic circuit.

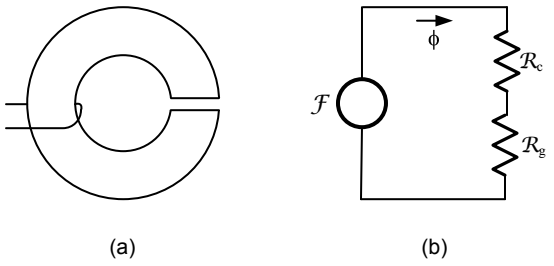

Figure 23.4: Equivalent electrical circuit for a gapped core. The reluctance of the core and the reluctance of the gap are modeled as series resistors.

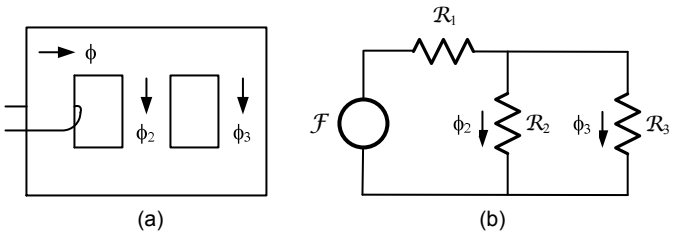

Figure 23.5: Equivalent electrical circuit for a core with three branches. Each branch is modeled as a separate reluctance.

$$\begin{aligned}\mathcal{R} &= \mathcal{R}_c + \mathcal{R}_g \\ &= \frac{l_c}{\mu_c A_c} + \frac{l_g}{\mu_0 A_g}\end{aligned} \tag{23.13}$$

Because the permeability of the core is usually much greater (often 1000 times greater) than that of free space, we often ignore R_c and approximate the reluctance as

$$\mathcal{R} \approx \mathcal{R}_g = \frac{l_g}{\mu_0 A_g} \tag{23.14}$$

We typically assume A_g is the same as A_c near the gap. However, due to fringing fields $\mathcal{R}_g$ is usually a slightly lower than would be calculated this way. While fringing flux, and hence gap reluctance, can be approximated analytically, it is more common to use a field solver to compute the fringing flux and gap reluctance numerically.

Figure 23.5 shows a core with three branches (often called an E-core). A coil is wound around one branch - generating MMF $\mathcal{F}$. We model each branch of the core with a separate reluctance. All flux from the source flows through the first branch $\mathcal{R}_1$. This branch includes the entire C-shaped core part left of the middle column. The flux then splits and flows partly through each of $\mathcal{R}_2$ and $\mathcal{R}_3$. The middle column is modeled by $\mathcal{R}_2$ while the reverse C-shaped section to the right of the middle column is modeled by $\mathcal{R}_3$. The total reluctance seen by the coil is:

$$\mathcal{R} = \mathcal{R}_1 + \mathcal{R}_2 || \mathcal{R}_3 \tag{23.15}$$

$$= \mathcal{R}_1 + \frac{\mathcal{R}_2 \mathcal{R}_3}{\mathcal{R}_2 + \mathcal{R}_3} \tag{23.16}$$

This results in a total flux of:

$$\phi = \frac{\mathcal{F}}{\mathcal{R}} \tag{23.17}$$

$$= \frac{\mathcal{F}(\mathcal{R}_2 + \mathcal{R}_3)}{\mathcal{R}_1\mathcal{R}_2 + \mathcal{R}_1\mathcal{R}_3 + \mathcal{R}_2\mathcal{R}_3} \tag{23.18}$$

This flux is split by the *flux divider* formed by the reluctance of the right two branches so that

$$\begin{aligned}\phi_2 &= \phi\frac{\mathcal{R}_3}{\mathcal{R}_2 + \mathcal{R}_3} \\ &= \frac{\mathcal{F}\mathcal{R}_3}{\mathcal{R}_1\mathcal{R}_2 + \mathcal{R}_1\mathcal{R}_3 + \mathcal{R}_2\mathcal{R}_3}\end{aligned} \tag{23.19}$$

$$\begin{aligned}\phi_3 &= \phi\frac{\mathcal{R}_2}{\mathcal{R}_2 + \mathcal{R}_3} \\ &= \frac{\mathcal{F}\mathcal{R}_2}{\mathcal{R}_1\mathcal{R}_2 + \mathcal{R}_1\mathcal{R}_3 + \mathcal{R}_2\mathcal{R}_3}.\end{aligned} \tag{23.20}$$

23.3 Inductance

Changes in i cause changes in ϕ which in turn induce voltage across our coil. The voltage across each turn of our coil is given by Faraday's law.

$$V_{\text{Turn}} = \frac{d\phi}{dt} \tag{23.21}$$

So the voltage across the coil of N turns is:

$$V = N\frac{d\phi}{dt} \tag{23.22}$$

Substituting (23.5) gives

$$V = \frac{N^2}{\mathcal{R}}\frac{di}{dt} \tag{23.23}$$

We refer to the constant $N^2/\mathcal{R}$ in (23.23) as *inductance* or L. So we have

$$L = \frac{N^2}{\mathcal{R}} \tag{23.24}$$

$$L = \frac{N^2 A\mu}{l} \tag{23.25}$$

Note that inductance is proportional the cross-sectional area of the magnetic circuit A and inversely proportional to the length of the magnetic path (not the length of the coil). Adding turns to an inductor increases inductance quadratically.

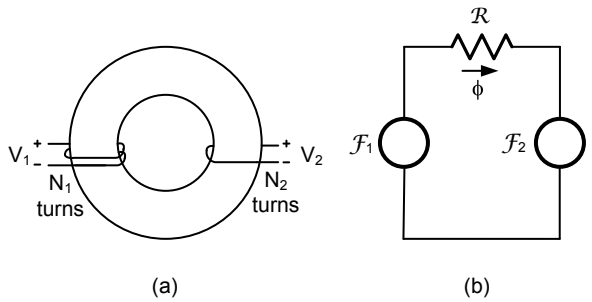

Figure 23.6: (a) A transformer is created by placing two windings on the same core so that the flux in the cores passes through both windings. (b) The equivalent circuit has a source of MMF for each winding and a reluctance $\mathcal{R}$ for the core.

23.4 Transformers

A transformer (introduced in Section 9.1) is created by arranging two or more windings so that some or all of the flux through one winding also passes through the other winding. Consider the arrangement of two windings, one with N_1 turns and one with N_2 turnn, on a toroidal core shown in Figure 23.6(a). Assume for now that all flux remains in the core and hence passes through both windings. By Faraday's Law we know that the change in flux induces a voltage in each turn around the core of $V_{\text{turn}} = \frac{d\phi}{dt}$. Thus we have

$$V_1 = N_1 \frac{d\phi}{dt} \tag{23.26}$$

$$V_2 = N_2 \frac{d\phi}{dt} \tag{23.27}$$

$$V_2 = \frac{N_2}{N_1} V_1 \tag{23.28}$$

The equivalent magnetic circuit for the transformer is shown in Figure 23.6(b). Summing MMF around this loop we can write:

$$N_1 i_1 - N_2 i_2 = \phi \mathcal{R} \tag{23.29}$$

If the reluctance were zero we would have the current relationship of an ideal transformer

$$i_2 = \frac{N_1}{N_2} i_2. \tag{23.30}$$

With non-zero reluctance, however additional current i_M flows that is proportional to ϕ. Suppose the secondary were open circuited — i.e., $i_2 = 0$. Then we have

$$i_M = \frac{\mathcal{R}}{N_1}\phi \tag{23.31}$$

$$\frac{di_M}{dt} = \frac{\mathcal{R}}{N_1}\frac{d\phi}{dt} \tag{23.32}$$

$$\frac{di_M}{dt} = \frac{\mathcal{R}}{N_1}\frac{1}{N_1}V_1 \tag{23.33}$$

$$\tag{23.34}$$

This current is that due to the *magnetizing inductance* L_M of the transformer which, referenced to the primary side, has a value of:

$$L_M = \frac{N_1^2}{\mathcal{R}} \tag{23.35}$$

The electric circuit realized by this component is an ideal transformer with a magnitizing inductor with magnitude L_M in parallel with the primary, like Figure 9.3 without the leakage inductance.

Now suppose the transformer the geometry of Figure 23.7(a) where some of the flux that passes through the primary coil is diverted through the middle branch of the circuit and does not link the secondary. The equivalent circuit is shown in Figure 23.7(b). The two windings are represented by two MMF sources $\mathcal{F}_1$ and $\mathcal{F}_2$. The three branches of the core are represented by three reluctances.

Performing a $Y - \Delta$ transformation on the three reluctances gives the circuit of Figure 23.7(c). The three reluctances have values:

$$\mathcal{R}_A = \frac{\mathcal{R}_1\mathcal{R}_2 + \mathcal{R}_1\mathcal{R}_3 + \mathcal{R}_2\mathcal{R}_3}{\mathcal{R}_\in} \tag{23.36}$$

$$\mathcal{R}_B = \frac{\mathcal{R}_1\mathcal{R}_2 + \mathcal{R}_1\mathcal{R}_3 + \mathcal{R}_2\mathcal{R}_3}{\mathcal{R}_\infty} \tag{23.37}$$

$$\mathcal{R}_C = \frac{\mathcal{R}_1\mathcal{R}_2 + \mathcal{R}_1\mathcal{R}_3 + \mathcal{R}_2\mathcal{R}_3}{\mathcal{R}_\ni} \tag{23.38}$$

Reluctance $\mathcal{R}_C$ here represents the magnetizing inductance of the transformer. Reluctances $\mathcal{R}_A$ and $\mathcal{R}_B$ represent the primary- and secondary-side leakage inductances respectively.

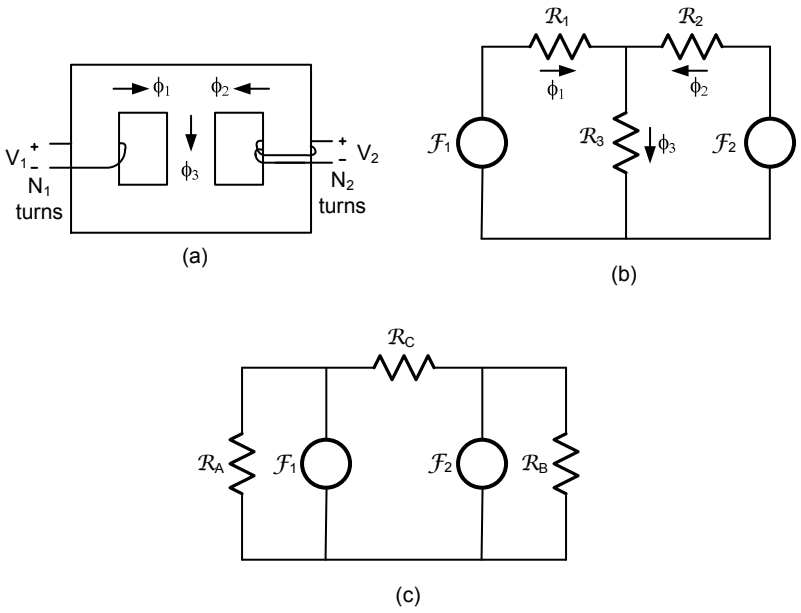

Figure 23.7: (a) A transformer with a *leakage* path for flux that does not link both windings. (b) The equivalent circuit has a source of MMF for each winding and three reluctances, one for each branch of the core. (c) A $Y - \Delta$ transformation of the reluctances separates the reluctance that determines the magnetizing inductance $\mathcal{R}_C$ from the reluctances that determine the leakage inductances.

The resulting circuit is an ideal transformer with a magnetizing inductance in parallel with the primary and leakage inductances in series with both the primary and secondary — like Figure 9.3 but with an additional leakage inductor in series with the secondary.

23.5 Exercises

Solve for flux given core geometry and number of turns

Solve for inductance with and without gap

Sove for flux with complex geometry - multiple branches with and without gaps

Solve for transformer parameters

Solve for mag inductance

Solve for leakage inductance

Chapter 24

Magnetic Materials

24.1 Magnetic Materials

In an ideal material there is a linear relationship between B and H, $B = \mu H$, as given by (23.10). Real magnetic materials exhibit *saturation* and *hysteresis* that lead to a non-linear $B - H$ characteristic.

24.1.1 Saturation

Magnetic materials with high relative permeability μ_r can only sustain this high permeability when the B field is below a maximum level B_{sat}. Figure 24.1 shows a piecewise linear $B - H$ function that exhibits saturation. Below saturation $B = \mu_r \mu_0 H$ giving a steep slope $\mu_r \mu_0$. Above B_{sat}, the slope is reduced to μ_0. Silicon steel has B_{sat} ≈1.5T. Ferrites have B_{sat} ≈0.5T.

Magnetic saturation limits the maximum current that can be handled by an inductor.

$$I_{\max} = B_{\text{sat}} \frac{l}{N\mu} \tag{24.1}$$

If current i exceeds $I_{\max}$ the incremental permeablility of the material falls by a factor of μ_r. This causes the inductance to fall by a factor of μ_r. In an inductor with an applied voltage, this causes the current to increase μ_r times faster — often with catastrophic results.

The maximum through current of a transformer is not limited by core saturation. The magnetizing current of a transformer (viz. the current flowing in the magnetizing inductance) is limited by (24.1).

Figure 24.1 shows an abrupt transition in incremental permeability (the slope of the B-H curve) at B_{sat}. While some magnetic materials have such a hard transition, in other materials the transition is more gradual. For example, Figure 24.2 shows the B-H curve for a Ferroxcube ferrite this material has an initial permeability of $\mu_r = 3000$ and saturates gradually at 0.4T at 100°C. In

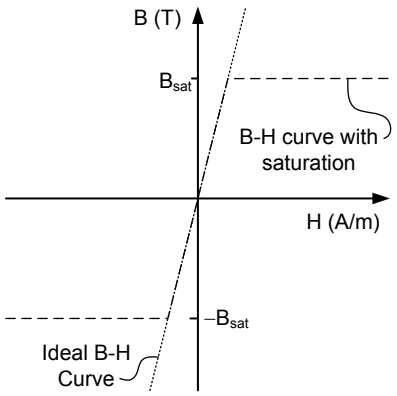

Figure 24.1: B-H curves with and without saturation. Above the saturation flux density B_{sat} the permeability, and hence the slope of the curve, is reduced from $\mu_r\mu_0$ to μ_0.

practice one would not operate cores build from this material with a B field above 0.3T.

The B-H curve in Figure 24.2 also has hysteresis which we shall discuss in the next section.

24.1.2 Hysteresis

A ferromagnetic material has *memory*[1]. When a strong B field is applied to the material, it becomes *magnetized*. That is, it behaves like a permanent magnet. When a material is magnetized, the B field is given by:

$$B = \mu(H + M) \tag{24.2}$$

The magnetization M provides a MMF in the magnetic circuit, just like a winding carrying current.

Consider a cycle of a magnetic component. Initially a high negative H field is applied to the component this causes magnetic domains to aligh so that $M = -H_c$. As H is increased the B field follows the right side of the curve of Figure 24.3. When $H = H_c$, $B = \mu(H + M) = 0$. That is, we have to apply a magnetic field of H_c, called the *coercivity* of the material, to zero the flux density. For silicon steel $H_c \approx 40$ A/m. For ferrites $H_c \approx 10$ A/m.

[1]At one time, almost all computers used *magnetic-core* memories based on this principle.

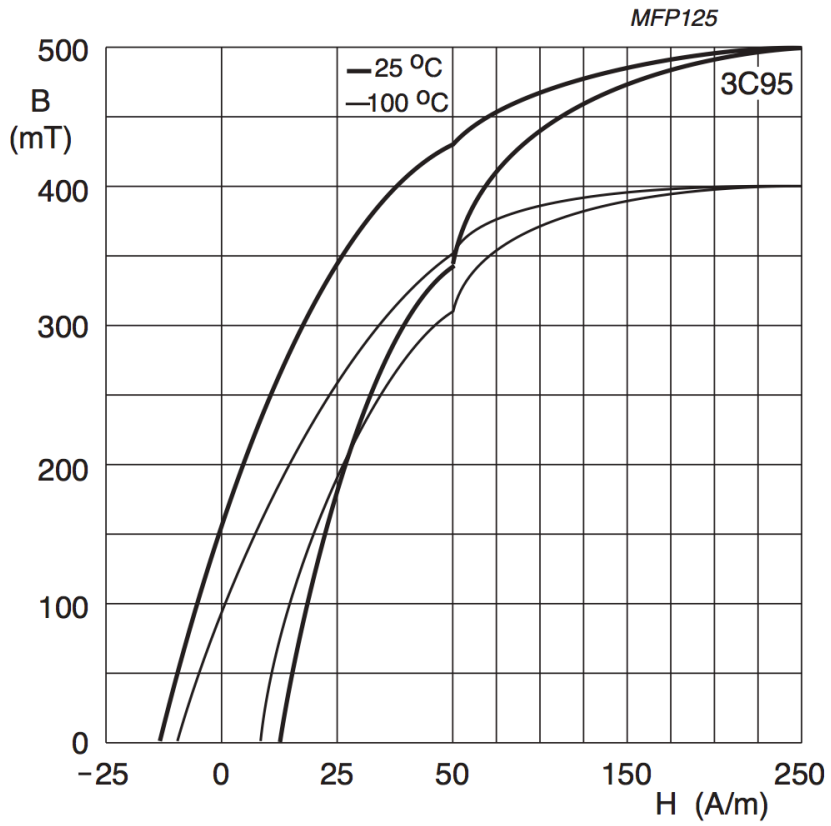

Figure 24.2: B-H curves for Ferroxcube 3C95 ferrite material at 25°C and 100°C. These curves exhibit both saturation and hysteresis and have a soft saturation characteristic.

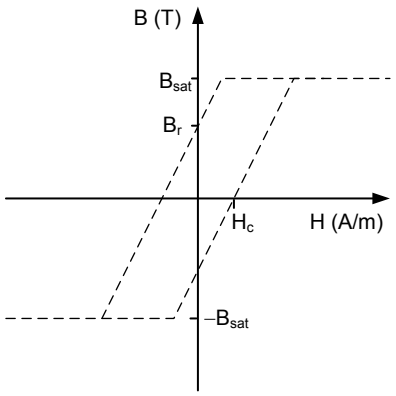

Figure 24.3: A B-H curve with hysteresis. Initially the core is magnitized with $M = -H_c$ so that an H-field of H_c is required to bring the flux density to 0. After driving the core to saturation the magnetiztion is reversed to $M = H_c$ so that when the H-field is removed a residual flux density of B_r remains.

As we increase H above H_c, the B field goes positive and we continue to move up the right side of the curve in Figure 24.3. When $B = B_{\text{sat}}$ the core saturates. At this high field, the magnetic domains in the material reverse and $M = H_c$.

With positive magnetization, as we bring the H field back to zero we now follow the left side of the curve in Figure 24.3 as we reduce H. When $H = 0$, a residual flux density $B_r = \mu H_c$ remains.

As we continue to reduce H below zero we follow the left side of the curve of Figure 24.3 until we reach $B = -B_{\text{sat}}$. At this high negative field the magnetization reverses so that $M = -H_c$ and we start the cycle over again.

Hysteresis in the B-H curve causes core losses. Each cycle around the B-H loop dissipates energy proportional to the area of the loop and the volume of the core. To see this note that:

$$v = \frac{d\phi}{dt} = A\frac{dB}{dt} \tag{24.3}$$

$$i = lH \tag{24.4}$$

$$P = vi = AlH\frac{dB}{dt} \tag{24.5}$$

$$E = \int Pdt = Al\int HdB \tag{24.6}$$

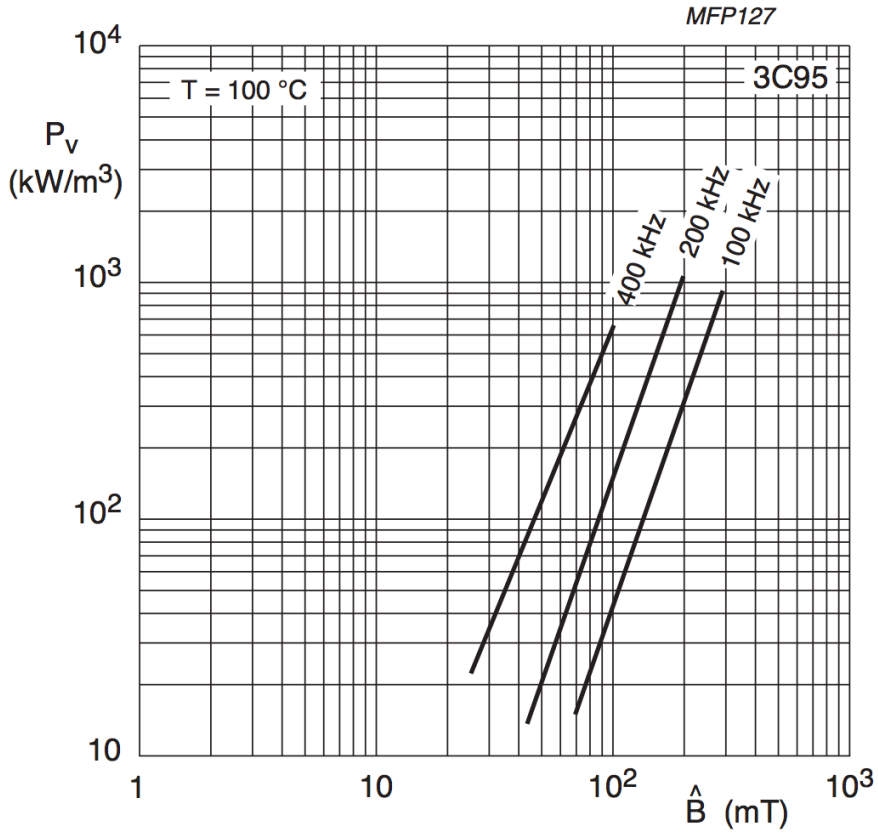

Figure 24.4: Loss per unit volume, P_V, for Ferroxcube 3C95 material at 100kHz, 200kHz, and 400kHz.

Here Al is the volume of the core — the product of its area and length. The integral $\int HdB$ over one cycle is the area of the loop.

The actual B-H curve, and hence the core loss, varies with frequency f and ΔB_{max}. Figure 24.4 shows the core loss per unit volume (kW/m^3) for the Ferroxcube 3C95 material at three frequencies. This relationship between loss, frequency, and $B_{\max}$ can also be described as a power law:

$$P_V = K f^{\alpha} \Delta B_{\max}^{\beta}. \tag{24.7}$$

For Magnetics Inc. R material, for example, K =36, $\alpha = 1.64$, and β =2.68 — with P_V in units of W/m^3.

24.2 Exercises

simulate inductor saturating
    calculate loss in one cycle of B-H curve
    calculate core loss for typical situation

Chapter 25

Wire for Magnetics

Our magnetic components, inductors and transformers, are realized by winding conductive wire on a core of magnetic material. Most often the winding is made using copper *magnet wire* — bare copper wire with a very thin layer of insulating material. However copper tape, bar, and tube and conductors made from aluminum may be used in certain situations. The conductivity of the wire material and wire cross-section area determine the DC resistance of the winding.

For many green electronics applications, however, the AC performance of the magnetic component is important. Typical operating frequencies for switching power converters run from 10s of kHz to low MHz. At these frequencies the *skin effect* and *proximity effect* cause the bulk of the AC current to flow in a thin layer on the surface of the conductor. This causes the AC conductivity of a wire to increase only linearly with radius. At high frequencies thick (low gauge) wires waste copper and magnetic core *window area*.

A solution to the problem of AC conductivity is to realize a winding with many parallel, fine-gauge *strands* of wire. The diameter of each strand is chosen so that the bulk of the wire area carries current. In some cases *Litz wire* which consists of many parallel, fine-grain strands braided together may be employed.

25.1 Skin Effect

AC current in a wire flows primarily in a narrow layer near the surface of the wire with the current density falling off exponentially with depth according to:

$$J = J_0 exp\left(\frac{d}{\delta}\right) \tag{25.1}$$

The *skin depth* δ is the the depth at which the current has fallen off to $\exp -1$ of its starting magnitude and is given by:

$$\delta = \sqrt{\frac{\rho}{\pi f \mu}} \tag{25.2}$$

Where ρ is the conductivity of the wire material and f is the frequency of the AC current. From (25.2) we see that δ is proportional to $f^{-1/2}$. For copper at 100kHz the skin depth is about 0.2mm.

Above a *skin-depth frequency* f_s, a wire of radius r can be considered to be carrying current uniformly over its entire cross section.

$$f_S = \frac{\rho}{\pi \mu r^2} \tag{25.3}$$

Below f_s the AC and DC resistance of the wire are identical. Above f_s the AC resistance becomes appreciably larger than the DC resistance of the wire and is given by:

$$R_{\mathrm{AC}} = R_{\mathrm{DC}} \left(\max 1, \sqrt{\frac{f}{f_S}} \right) \tag{25.4}$$

AC current in a wire generates a time varying magnetic field that penetrates the wire. This field, in turn, opposes the flow of current in the core of the wire. This results in the distribution of current given by (25.1) (which can be derived fairly easily from Maxwell's equations).

To avoid having high AC resistance, or wasting copper, or both, we choose the radius of the wire to satisfy (25.3) so that:

$$r < r_s = sqrt \frac{\rho}{\pi \mu f} \tag{25.5}$$

25.2 Proximity Effect

$$l = \frac{2nNr}{w} \tag{25.6}$$

$$R_{\mathrm{AC}} = R_{\mathrm{DC}} \left(1 + \frac{(\pi n N)^2 r^6}{12 \delta^4 w^2} \right) \quad = R_{\mathrm{DC}} \left(1 + \frac{(\pi l)^2}{12} \left(\frac{r}{\delta} \right)^4 \right) = R_{\mathrm{DC}} \left(1 + \frac{(\pi l)^2}{12} \left(\frac{f}{f_S} \right)^2 \right) \tag{25.7}$$

25.3 Litz Wire

25.4 Winding Geometry

inter-winding capacitance

interleaving primary and secondary

25.5 Exercises

Chapter 26

Designing Magnetic Components

Now that we understand the basics of magnetic circuits and the properties of magnetic materials we are prepared to discuss how to design magnetic components: inductors and transformers. We consider inductors in this chapter and deal with transformers in Chapter 27.

The design of an inductor is at its simplest specified by three parameters: a core, the number of turns of wire on that core N, and the length of the air gap in the core l_g. Given the specification, the parameters of a core, and $B_{\max}$ (a key design variable), we derive formulae for calculating N and l_g.

Thus, the design of an inductor is a search for the best core and the optimal value of $B_{\max}$. To conduct this search we consider a range of cores and use a numerical solver (such as the "solver" in Microsoft Excel) to solve for the optimal value of $B_{\max}$, the value that gives the lowest power. We then pick the smallest core that meets our specification.

26.1 Inductor Design

The process of inductor design starts with a specification and a set of optimization constraints. The specification typically includes the inductance L, the operating frequency f, the maximum current $I_{\max}$, the current ripple ΔI, and the RMS current $I_{\rm RMS}$.

The optimization constraints deal with the tradeoff between cost and power loss. In some cases we want to design the lowest cost inductor that has losses less than a specified maximum P_{Lmax}. In other cases we want the inductor with the lowest possible loss below a given maximum cost.

In some cases, some of our specification variables may be part of a larger optimization. For example, in the design of a switching converter, one may vary switching frequency f to trade off switching losses against core losses. For now, however, we will consider f to be fixed by the specification.

Given a specification and an optimization constraint inductor design is an iterative process involving the following steps — using a spreadsheet, Matlab, or computer program to perform the optimization.

1. Pick a core.

2. Calculate the gap length l_g and number of turns N required to achieve inductance L on this core as a function of $B_{\max}$. Note that $B_{\max}$ is not a part of the specification, but rather an optimizaiton variable that we will adjust in step 4 to optimize our design.

3. Select a wire gauge and compute the number of parallel *strands* that will fit on the core — also as a function of $B_{\max}$.

4. Perform an optimization to find the $B_{\max}$ that gives the lowest loss on this core. As $B_{\max}$ is swept one trades off core loss against wire loss. Minimum loss typically happens at the point where the two are roughly balanced.

5. Iterate steps 1-4 over other cores to optimize the solution.

The three parts of the current specification characterize the current waveform. $I_{\max}$ is the peak of the waveform and determines $B_{\max}$ that must be kept within limits to avoid saturating the core. $\Delta I_{\max}$ is the current ripple which determines ΔB and hence the area of the path traced in operation on the $B-H$ curve and thus the core loss. $I_{\rm rms}$, along with the resistance of the winding, determines the wire loss of the inductor.

To compute the number of turns required to realize an inductance of L with a maximum current of $I_{\max}$ and flux density of $B_{\max}$ we substitute (23.24) into (23.8) at this operating point and derive:

$$B_{\max} = \frac{NI_{\max}}{A\mathcal{R}} \tag{26.1}$$

$$B_{\max}A\mathcal{R} = NI_{\max} \tag{26.2}$$

$$B_{\max}A\frac{N^2}{L} = NI_{\max} \tag{26.3}$$

$$N = \frac{LI_{\max}}{B_{\max}A} \tag{26.4}$$

We see that the number of turns required decreases with $B_{\max}$ and increases with $I_{\max}$.

To determine l_g, the length of the gap needed we assume that all of $\mathcal{R}$ is across the gap. Substituting (23.24) into (23.14) we write:

$$\mathcal{R} = \frac{l_g}{\mu_0 A_g} \tag{26.5}$$

$$\frac{N^2}{L} = \frac{l_g}{\mu_0 A_g} \tag{26.6}$$

$$l_g = \frac{N^2 \mu_0 A_g}{L} \tag{26.7}$$

In step 3 we need to select a wire gauge that is fine enough so that the skin effect and proximity effect will be negligible, since we will be ignoring them for the rest of our calculations.

26.2 Inductor Example

Suppose we need a 22 μH inductor that will handle a sawtooth current waveform that varies between 5 A and 10 A with a frequency of 200 kHz. Further suppose we desire the lowest cost inductor that has total losses (core + wire) less than 1 W. From the current waveform we calculate $I_{\max} = 10$ A, $\Delta I_{\max} = 5$ A, and $I_{\rm rms} = 7.64$ A.

We start by tentatively choosing a Ferroxcube PQ32/30 core. (We will revisit this choice as we iterate our design later.) From the manufacturers data sheet we get (or calculate) the core area A, the core volume V, the copper *window* area A_W, and the length of the average turn of wire:

$$\begin{aligned} A &= 1.67 \times 10^{-4}\text{m}^2 \\ V &= 1.25 \times 10^{-5}\text{m}^3 \\ A_W &= 1.49 \times 10^{-4}\text{m}^2 \\ l_{\rm turn} &= 0.064\text{m} \end{aligned}$$

We initially assume a value of $B_{\max} = 0.2$T. Later we will optimize this value to minimize losses. However, for now, we can calculate:

$$\begin{aligned} N &= \frac{L I_{\max}}{B_{\max} A} \\ &= \frac{(22 \times 10^{-6})(10)}{(0.2)(1.67 \times 10^{-4})} \\ &= 6.59 \approx 7 \text{ turns} \end{aligned}$$

We round the calculated number of turns, 6.59, to the nearest integer, 7. We use the integral value in the calculation of the required gap.

$$\begin{aligned} l_g &= \frac{N^2 \mu_0 A_g}{L} \\ &= \frac{(7^2)(1.26 \times 10^{-6})(1.67 \times 10^{-4})}{22 \times 10^{-6}} \\ &= 0.042\text{mm} \end{aligned}$$

At this point we have the initial design for our inductor. The next step is to calculate the core loss and winding loss.

We calculate the core loss from (24.7) using $\alpha = 1.65$ and $\beta = 2.5$ and choosing K to match a loss of 80 kW/m^3 at 0.1 T and 100 kHz.

$$\begin{aligned} P_C &= V_{\text{core}} K f^{\alpha} \Delta B_{\text{max}}^{\beta}, \\ &= (1.25 \times 10^{-5})(80 \times 10^3) \left(\frac{2 \times 10^5}{10^5}\right)^{1.65} \left(\frac{0.1}{0.1}\right)^{2.5} \\ &= 3.14\text{W}. \end{aligned}$$

[Redo this with 30AWG or finer]

We calculate the wire loss from the resistance of the winding and I_{rms}. Assuming 24AWG magnet wire with a cross sectional area of 0.2 mm^2, each $N = 7$ turn *strand* of the winding has an area of $A_S = 1.4 \times 10^{-6}$ m^2. If we assume a *fill factor*, the fraction of the window filled by wire, of $K_f = 0.7$, we can fit $S = \lfloor \frac{A_W}{A_S} \rfloor = 74$ parallel strands of $N = 7$ turns each on our core. With a resistance of 84mΩ /m, the resistance of the winding is

$$\begin{aligned} R_W &= \frac{\rho l_{\text{turn}}}{S} \\ &= \frac{(0.084)(0.064)}{74} \\ &= 5.12 \times 10^{-4}\Omega. \end{aligned}$$

Which gives

$$\begin{aligned} P_W &= I_{\text{rms}}^2 R_W \\ &= (7.64^2)(5.12 \times 10^{-4}) \\ &= 0.03\text{W}. \end{aligned}$$

So our total losses are $P_T = P_C + P_W = 3.17W$. This is more than the 1 W of our specification, but before we start looking for a larger core, we observe that the loss is hugely unbalanced — with P_C being 100 times larger than P_W.

To reduce P_C at the expense of P_W we reduce our optimization variable B_{max}. Using the *solver* feature of Microsoft Excel, we search for the *optimal* value of B_{max}, i.e., the value that gives the lowest total loss. We find the optimum values are:

Core	P18/11	PQ20/16	PQ20/20	PQ26/20	PQ26/25	PQ32/30	
A	4.33	6.19	6.26	12.1	12.0	16.7	$m^2 \times 10^{-5}$
V	1.12	2.33	2.85	5.47	5.82	12.5	$m^3 \times 10^{-6}$
A_W	2.09	4.74	6.44	5.75	8.45	14.9	$m^2 \times 10^{-5}$
l_{turn}	33	42	42	53	54	64	mm
$B_{\max}$	0.27	0.17	0.16	0.11	0.10	0.06	T
N	18	20	22	16	18	22	turns
L_g	0.80	1.41	1.73	1.77	2.22	4.62	mm
P_C	0.62	0.41	0.39	0.31	0.25	0.15	W
P_W	0.73	0.52	0.45	0.35	0.30	0.30	W
P_T	1.35	**0.92**	0.84	0.66	0.55	0.46	W

Table 26.1: Optimization of the 22uH inductor. For each core we run an optimizer to find the $B_{\max}$ that gives the lowest total power P_T. The smallest core that meets the specified P_T is selected.

$$\begin{aligned} B_{\max} &= 0.06 \text{ T} \\ N &= 22 \text{ turns} \\ l_g &= 4.62 \text{ mm} \\ P_C &= 0.15 \text{ W} \\ P_W &= 0.30 \text{ W} \\ P_T &= 0.45 \text{ W} \end{aligned}$$

Reducing $B_{\max}$ directly reduces core loss from (24.7). However, from (26.4) it also requires more turns for the same inductance — increasing wire loss. Thus reducing $B_{\max}$ trades core loss for wire loss until the optimal point is found.

We now have a design that meets our specification. However, since our total power is less than half the required value, it is worth seeing if we can economize by using a smaller core.

The process of core selection is illustrated in Table 26.2. This table shows relevant parts of an Excel spreadsheet used for the optimization. The cores considered are listed in the first row — from smallest on the left to largest on the right. The core parameters for each core are listed in the next four rows. Row six gives the optimal value of $B_{\max}$ found by the solver. The design parameters N and l_g for the inductor with this value of $B_{\max}$ are given in the next two rows. The final three rows show the calculated power.

The smallest core that meets our specification of $P_T < 1$ W is the PQ20/16 with $N = 20$ turns and a gap of $l_g = 1.41$ mm.

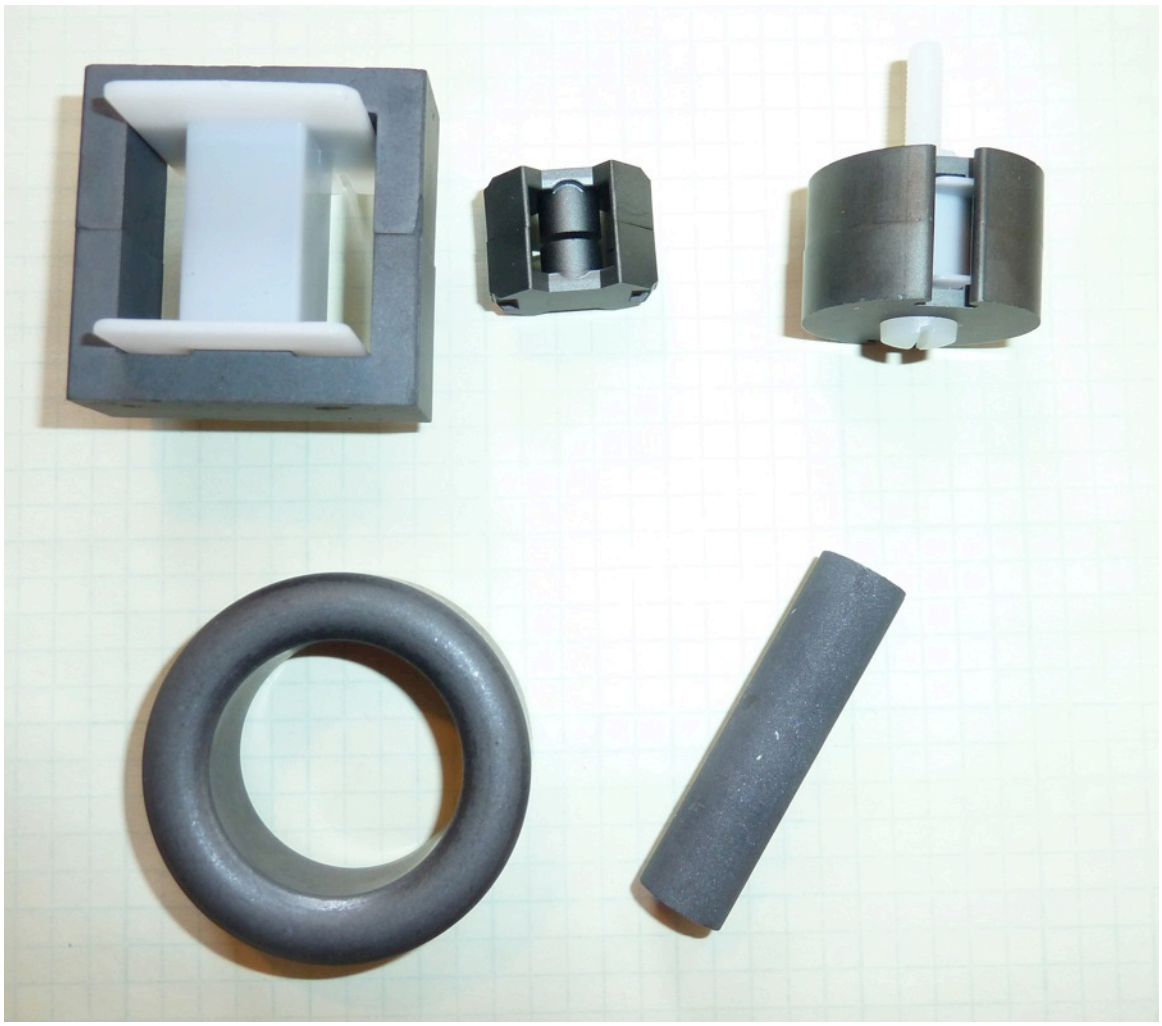

Figure 26.1: Examples of ferrite cores: an E-core, a PQ-core, a pot core, a toroid, and a ferrite rod.

26.3 Core Selection

Magnetic cores come in many shapes and sizes as shown in Figure 26.1. Common shapes include toroids, E-cores, and P-cores. As with our example, core selection is an iterative process: we pick a core, optimize a design using that core, and then adjust our core size up or down depending on the result.

We can however, inform our initial choice of a core by understanding how two key parameters of the core impact properties of inductors using that core. The core cross-sectional area A determines the amount of flux the core can carry at a specified $B_{\max}$, and hence the amount of energy, $E = 1/2I^2L$, that can be stored in the core. The core window area A_W determines the amount of effective current $NI_{\max}$ the core can carry at a given maximum current density $J_{\max}$.

To get a rough idea of sizing, consider an inductor with a single turn, $N = 1$, a gap that is the square root of the core area, $l_g = \sqrt{A}$, and for which $B_{\max} = 0.2T$. We can calculate

$$\begin{aligned} L &= \frac{1}{\mathcal{R}} \\ &= \mu_0\sqrt{A} \\ I_{\max} &= B_{\max}A\mathcal{R} \\ &= \frac{B_{\max}\sqrt{A}}{\mu_0} \\ E &= \frac{1}{2}I_{\max}^2 L \\ &= \frac{B^2A^{3/2}}{2\mu_0} \qquad (26.8) \\ A &= \left(\frac{2E\mu_0}{B_{\max}^2}\right)^{2/3} \qquad (26.9) \end{aligned}$$

From (26.8) we see that the energy that can be handled by a core is proportional to its volume $A^{3/2}$. If an optimized inductor for a certain energy E requires a core with an area A, an inductor that can store $2E$ will require a core with area $1.6A$.

The window area A_W for standard cores scales proportionally with A. Thus the relative magnitudes of core loss, which is proportional to volume, and wire loss, which is determined by window area remains roughly constant.

26.4 Details

The example above glosses over three important details in the design process: selection of the magnetic material, selection of the wire, and how the wire is wound on the core.

The core material is typically selected up front based on the operating frequency of the inductor. A low-frequency (10kHz) inductor might select a silicon

steel core material while a medium frequency (200kHz) inductor would use a ferrite core. If there is a question about which core material is better for a given application, the design can be carried out for both materials and the better one selected. For example, we could carry out the design of our 22μH inductor for both FerroxCube 3C95 and 3C96 materials.

The wire gauge is determined by frequency of operation and the current to be handled (Chapter 25). At low frequencies a single large-gauge wire can be used. At higher frequencies multiple fine-gauge strands are preferred to reduce the increase in AC resistance due to the skin effect. Litz wire may be used to mitigate the proximity effect. The selection is also influenced by current ripple — since the DC component of current sees the DC resistance of the wire and the AC component of the current sees the higher AC resistance. The triangle waveform typical of switching power converters has frequency components at odd harmonics of the fundamental switching frequency. Each harmonic sees the AC resistance associated with its frequency. An inductor that has a very small ripple can get by with a higher wire gauge than an inductor that must handle a high ripple.

Finally, how the wire is wound on the core (or the bobbin that slides on the core) greatly affects the parasitic inter-winding capacitance of the inductor. Ideally, an inductor should be wound from one end of the core to the other with no "doubling back". The individual strands of each winding should be stacked on top of one another. If the winding "doubles back", two turns that are electrically far apart (and hence at very different voltages) wind up in close proximity. Capacitive coupling between these turns adds to the parasitic capacitance of the inductor and increases losses.

26.5 Exercises

repeat optimization for Ferroxcube E-cores -
repeat optimization for 10uH
repeat optimization for 2x current
optimization of frequency as well with switching energy

Chapter 27

Transformer Design

We design a transformer by designing an inductor that realizes the magnetizing inductance of the transformer using half of the window area for the primary winding. The magnetizing inductance required is determined by the volt-seconds the transformer must handle each half-cycle and the maximum flux density $B_{\max}$ that is allowed. As with inductor design, $B_{\max}$ is a design parameter that is optimized to give the lowest total loss. Because a transformer does not need to store energy, it is typically designed with no gap in the core.

27.1 Design Procedure

Transformer design, like inductor design, is an iterative process. We pick a core, solve for the value of $B_{\max}$ that gives minimum power dissipation and then repeat the process for other cores. There are two key differences. The first is that transformers, since they transmit rather than store energy, do not have gapped cores[1].

The second difference is that we compute the number of turns so that the magnetizing current reaches the value that gives $B_{\max}$ at the end of each half cycle. If we apply a voltage V_P to the primary for time t_P, the number of turns required so that the flux swings from $-B_{\max}$ to $B_{\max}$ is:

$$N_P = \frac{V_P t_P}{2 B_{\max} A}. \tag{27.1}$$

Typically $t_P = t_{cy}/2$ so we can rewrite (27.2) as:

$$N_P = \frac{V_P t_{cy}}{4 B_{\max} A}. \tag{27.2}$$

[1]Transformers for flyback converters do store significant amounts of energy and hence do have gapped cores. Flyback transformers can be thought of as inductors with a secondary winding and our inductor design procedure can be applied to them

Once we have computed the number of turns, we calculate core loss using (24.7) and $\Delta B_{\max} = 2B_{\max}$.

To compute wire loss, we divide the core window area A_W into seperate areas for the primary and secondar (and any additional windings) and compute the loss for each winding separately. The lowest loss typically occurs when area is split evenly between the current-carrying windings.

27.2 Transformer Design Example

Suppose we require a 1.7kW 2:1 step-down transformer for a full-bridge converter that takes a 170 V, 10 A primary and steps it down to an 85 V, 20 A secondary. We require that the total loss of the transformer be less than 20 W. Our switching frequency is $f = 100$ kHz which gives $t_{cy} = 10$ μs and at 100% duty factor $t_P = 5$ μs. We assume a fill factor of $k_W = 0.7$ and allot half of this for each of the two windings.

We start by choosing a Ferroxcube E41/17/12 core with parameters:

$$\begin{aligned} A &= 1.49 \times 10^{-4}\text{m}^2 \\ V &= 1.15 \times 10^{-5}\text{m}^3 \\ A_W &= 1.68 \times 10^{-4}\text{m}^2 \\ l_{\text{turn}} &= 8.21 \times 10^{-2}\text{m} \end{aligned}$$

Running the "solver" we find minimum power is achieved with $B_{\max} = 0.12$ T and $\Delta B_{\max} = 0.24$ T. At this optimal point we have:

$$\begin{aligned} N_P &= 46 \text{ turns} \\ N_S &= 23 \text{ turns} \\ P_P &= 5.30 \text{ W} \\ P_W &= 10.60 \text{ W} \\ P_C &= 8.87 \text{ W} \\ P_T &= 19.46 \text{ W} \end{aligned}$$

The primary wire loss P_P is half the total wire loss P_W. The primary is 46 turns of six strands/turn while the secondary has 23 turns with 12 strands per turn. The secondary is one quarter the resistance of the primary (half its length and twice the strands per turn), but with twice the current, its power is the same.

This design meets our specification, but we decide to iterate with other cores to see if we can do better.

Table 27.2 shows the results of optimizing the transformer design for two additional cores. The E36/21/12 core gives slighty smaller losses. This core is not actually *smaller* than the E41/17/12, but rather has a different shape. It has a smaller A than the E41 but a larger A_W. In this case the larger window

Core	E35/18/12	E36/21/12	E 41/17/12	
A	1.00	1.26	1.49	m^2 $\times 10^{-4}$
V	8.07	12.2	11.5	m^3 $\times 10^{-6}$
A_W	1.81	2.25	1.68	m^2 $\times 10^{-4}$
l_{turn}	6.90	6.94	8.21	m $\times 10^{-2}$
$B_{\max}$	0.14	0.12	0.12	T
N_P	62	56	46	turns
P_C	8.20	8.76	8.87	W
P_W	14.26	9.33	10.60	W
P_T	22.46	18.09	19.46	W

Table 27.1: Optimization of the 1.7kW 2:1 transformer.

area gives lower wire loss with roughly the same core loss. If we were designing a custom core for this transformer, we might design one with an even larger A_W. The E35/18/12 core is a strictly smaller core, but its total power exceeds our target, so we stick with the E36/21/12 for our design.

27.3 Flyback Transformers

Unlike the transformers in other isolated converters, the transformer in a flyback converter stores energy. During the first half of each cycle, the primary winding stores energy in the magnetizing inductance. This energy is released during the second half of the cycle.

Because its primary function is to store energy, a flyback transformer is designed like an inductor, not a transformer. The design procedure is exactly that described for inductors in Chapter 26. The one difference is that only half the copper area is available for the primary winding. Thus, compared to a straight inductor, the resistance — and hence the wire loss — is doubled. During each half cycle current only flows in half the windings (either the primary or the secondary). At any instant in time, half the copper area is unused.

We leave the design of a flyback transformer as an exercise.

27.4 Exercises

transformer as above but 4:1 step down

340V to 48V 5A

flyback transformer design - 170V, 0.25 A input, 100kHz, 3012V output.

transformer with gap - repeat design example but with 1mm gap in each core.

Chapter 28

Photovoltaic Cells and Modules

28.1 Photovoltaic Cells

28.2 Photovoltaic Modules

28.3 Single-Phase PV Systems

28.4 Three-Phase PV Systems

28.5 Module Imbalance

28.6 PV Economics

28.7 Exercises

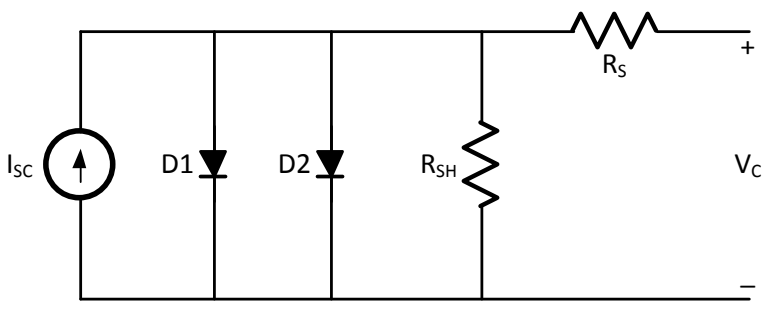

Figure 28.1: A photovoltaic cell can be modeled as a current source, a pair of diodes, and series and shunt resistances.

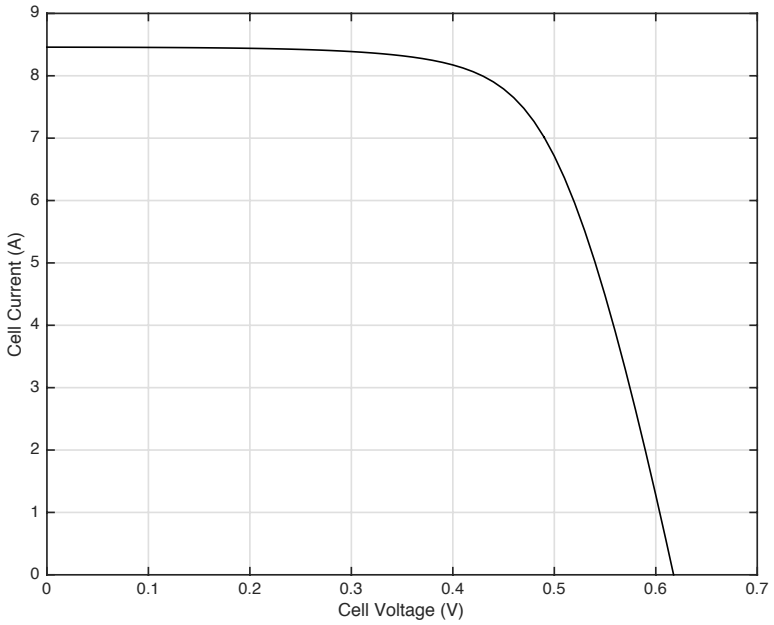

Figure 28.2: The IV curve of a typical PV cell. The y-intercept is the short-circuit current I_{SC} — determined by the current source in Figure 28.1. The curve follows a diode IV curve (with current inverted) reaching the x-axis at the an open-circuit voltage V_{OC} when all of the current flows through the diodes.

```
* pv cell model
.param isc = 8.46
.param rsh = 300
.param rs = 0.01
.model pvdiode D(Is=5e-4 N=3)
.model pvdiode2 D(Is=3e-10 N=1)

.subckt pv_cell anode cathode
I1 cathode a {isc}
D a cathode pvdiode n=1
D2 a cathode pvdiode2 n=1
RSH cathode a {rsh}
RS a anode {rs}
.ends pv_cell
```

Figure 28.3: SPICE model for a typical photovoltaic cell.

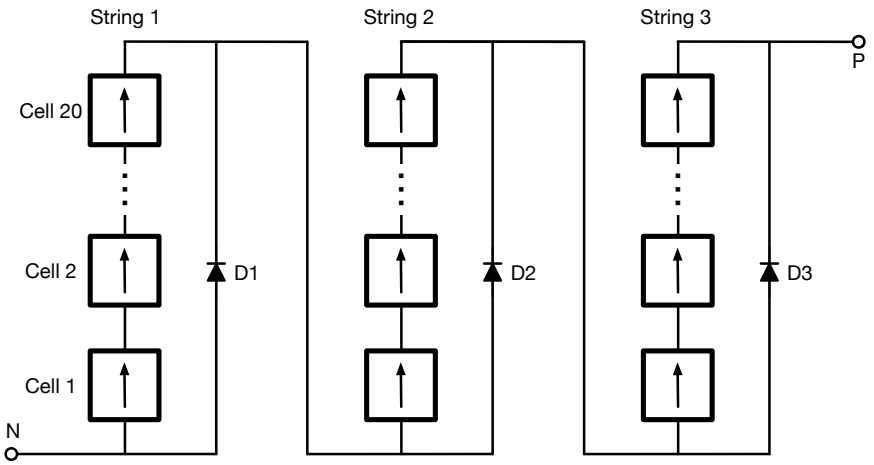

Figure 28.4: A typical photovoltaic module consists of 60 cells arranged in three substrings of 20 cells each. Bypass diodes are placed across each substring to bypass the substring if the voltage across the substring goes negative.

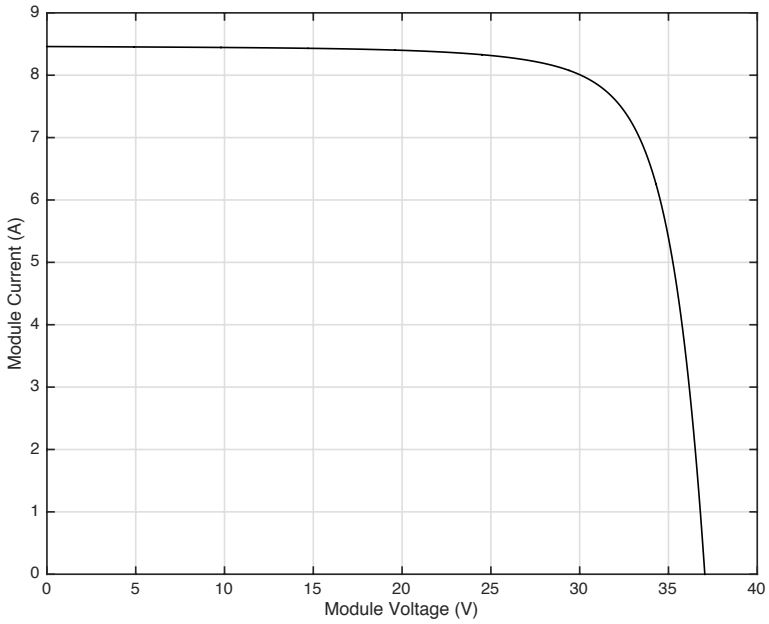

Figure 28.5: VI-characteristics of a typical photovoltaic module with balanced substrings. I_{SC} is the same as for the individual cells. V_{SC} is 60× the value for an individual cell.

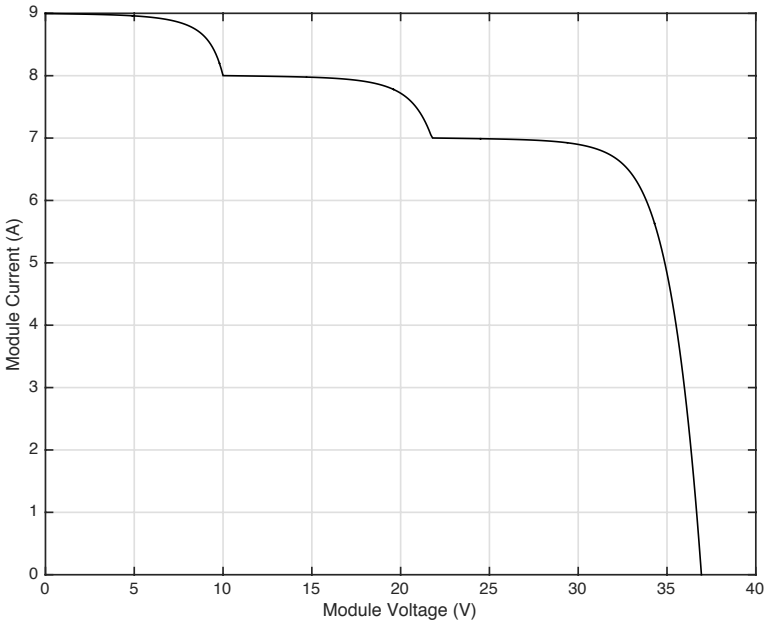

Figure 28.6: VI-characteristics of a module with mismatched substrings with I_{SC} values of 7, 8, and 9A respectively. The current of the weakest substring dominates until the voltage across that substring goes negative — forward biasing its bypass diode.

Part III. Motor Control

Brushed PM motor
Motor Control
Brushless PM motor
AC Induction motor

Chapter 29

Electric Motors

Many green electronic systems involve the conversion between mechanical energy and electrical energy. A wind turbine converts mechanical energy (from the wind) to electrical energy to drive the grid. An electric car converts mechanical energy (from the car's battery) to mechanical energy (to propel the car) and then back again (during regenerative braking). The device most commonly used to convert from electrical to mechanical energy is a *motor*. We refer to this device as a *generator* when it is used in the opposite direction — to convert from mechanical to electrical.

In this chapter we introduce the basic concepts of electric motors and generators by examining the operation and control of a brushed permanent magnet motor.

29.1 Lorentz Force and Faraday Induction

The force generated by a motor is due to the *Lorentz force* that is generated by a moving charge in a magnetic field.

$$F = q\,(v \times B) \tag{29.1}$$

Where F (the force), B (the magnetic field), and v (the charge velocity) are 3-vectors, and q (the charge) is a scalar. If the charges are confined to a wire we can rewrite this as

$$F = i\,(L \times B) \tag{29.2}$$

Where L is a 3-vector representing the length and direction of the wire and i is the scalar current.

Figure 29.1 shows (29.2) graphically. Here the B field is to the left and the current is flowing away from the viewer. The resulting force is upward — perpendicular to both B and L following the *right-hand* rule.

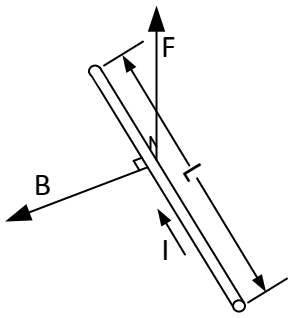

Figure 29.1: Current I flowing in a wire of length L in a magnetic field B induces a *Lorentz force* $F = I(L \times B)$.

If the wire in Figure 29.1 is moving in the direction of the force, *Faraday induction* induces a voltage across the wire that pushes back on the current. As the wire moves through the B field, it cuts field lines changing the amount of flux on one side of the wire. According to Faraday's Law, the magnitude of the voltage induced by the changing flux is given by:

$$V_E = \frac{d\phi}{dt}$$
$$V_E = |L|\,(B \times v) \tag{29.3}$$

Where V is the voltage and v is the velocity of the wire.

Equations (29.2) and (29.3) together describe the two most important properties of an electric motor. From (29.2) we see that the force, and hence torque for a rotating machine, is proportional to current, $F \propto i$. From (29.3) we see that open circuit voltage across the terminals of a motor (or generator) is proportional to velocity $V \propto v$, angular velocity $Vv \propto \omega$ for a rotating machine.

In summary, the electrical quantities motor current i_M and induced voltage V_E are related to the mechanical quantities torque τ and angular velocity ω as:

$$i_M = \frac{\tau}{K_M} \tag{29.4}$$
$$V_E = K_M \omega \tag{29.5}$$

Where K_M is the *motor constant*, a parameter that is related to the motor geometry, magnet properties, and winding parameters. The voltage V_E generated by faraday induction in a motor is often called the *electromotive force* or EMF

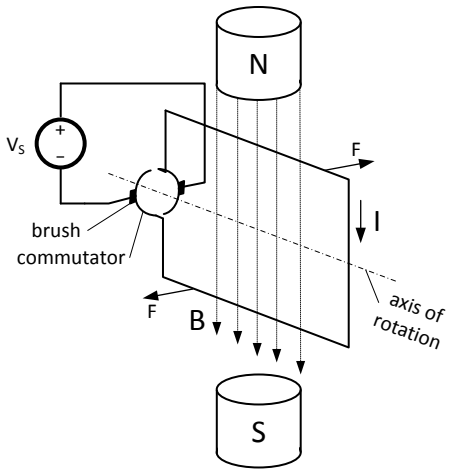

Figure 29.2: A brushed permanent-magnet motor. Permanent magnets generate a B field (dotted arrows) that cross a wire loop *rotor*. Current I flowing in the rotor induces forces F on the upper and lower loop segments that together constitute a torque τ. A commutator and a pair of brushes reverse the current flow each half cycle to keep the torque in the same direction.

of the motor. In some contexts, V_E is referred to as *back EMF* since it acts in a direction that *pushes back* on current i_M.

Note that with these equations mechanical power $P_M = \tau\omega$ is equal to electrical power $P_E = i_M V_E$.

29.2 The Brushed Permanent Magnet Motor

Figure **??** shows a very simple permanent magnet motor. The *stator* (stationary part) of this motor is a pair of permanent magnets that generates a B field oriented downward. The *rotor* (rotating part) of the motor consists of a rectangular wire loop that passes twice through this field, once above the axis of rotation, and once below.

A *commutator* ring switches the connection between a voltage source and the rotor loop so that the positive terminal of the voltage source is always connected to the upper loop segment. When the motor rotates 90-degrees from the position shown the connection to the voltage source is reversed. After rotating an additional 180-degrees (a total of 270 degrees) the connection is switched back to its original position.

If we pass a current i through the wire loop a Lorentz force is induced

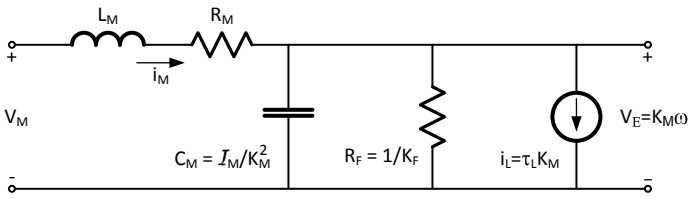

Figure 29.3: Electrical model of a brushed permanent magnet motor. V_M is the voltage applied at the motor terminals. Elements L_M and R_M represent the inductance and resistance of the rotor winding. Capacitor C_M represents the inertia of the rotor. Resistor R_F represents the friction of the motor. Current source i_L represents the load on the motor. The voltage across C_M, R_F, and i_L is $V_E = K_M\omega$, the EMF of the motor.

according to (29.2) in the upper and lower segments of the rotor. The force in the upper segment is into the page and the force in the lower segment is out of the page so that collectively they generate a torque about the axis of rotation. The magnitude of this torque is proportional to the current as per (29.4).

If we disconnect the voltage source and rotate the motor, Faraday induction generates a voltage across our rotor according to (29.3). The commutator ensures that the positive induced voltage is always on the upper terminal of the disconnected source. As per (29.5) the magnitude of this voltage is proportional to the angular velocity of the motor ω.

An actual motor is slightly more complex than what we have described here. Rather than a single loop in the rotor, it has several loops spaced by a small number of degrees so that the torque generated by the motor is approximately constant with angle. With a single loop the torque is a rectified sine wave. That is, the torqe at angle θ is proportional to $|\sin(\theta)|$. With multiple loops the torque is approximately constant, independent of angle.

The rotor of an actual motor also typically has a laminated steel core to provide a low reluctance path for the magnetic field to give a larger B field across a relatively short *gap* between the magnets in the stator and the rotor core. The fields and forces across the gap are complex, vary slightly with angle, and are usually modeled using field solvers. To a very good approximation, however, the motor V-I characteristics are still governed by (29.5) and (29.4).

29.3 Motor Model

For the purposes of motor control, we model a motor as shown in Figure 29.3. The two horizontal elements, R_M and L_M represent the resistance and inductance of the rotor windings. When a voltage V_M is applied to the motor ter-

minals, a current i_M flows giving a drop of $i_M R_M + L_M \frac{di_M}{dt}$ across the rotor windings. The voltage left after these winding drops is V_E, the EMF of the motor caused by faraday induction.

$$V_E = V_M - i_M R_M - L_M \frac{di_M}{dt} = K_M \omega \tag{29.6}$$

The inertia of the motor is modeled by capacitor C_M. A torque $\tau_I = \mathcal{I}_M \frac{d\omega}{dt}$ is required to accelerate the motor. Substituting (29.4) this torque requires a current

$$\begin{aligned} i_I &= \frac{\mathcal{I}_M}{K_M} \frac{d\omega}{dt} \\ &= \frac{\mathcal{I}_M}{K_M^2} \frac{d\vee_E}{dt} \end{aligned} \tag{29.7}$$

So the capacitance required to model the moment of inertia $\mathcal{I}_M$ is

$$C_M = \frac{\mathcal{I}_\mathcal{M}}{K_M^2} \tag{29.8}$$

We model the internal friction of the motor as resistor $R_F = \frac{1}{K_F}$. This friction draws a current $i_F = \frac{V_E}{R_F} = V_E K_F$ and hence has a torque of $\tau_F = V_E K_F K_M$.

Finally, the load on the motor is modeled as a current source, $i_L = \tau_L K_M$. In some cases we may use a more sophisticated model for the load. For example, in an electric vehicle we may use a capacitor to model the momentum of the vehicle, a resistor to model the friction of the vehicle, and a current source to represent changes in elevation.
And the kinetic energy stored in the inertia of the rotor is modeled by the energy stored in the capacitor.

$$E_I = \frac{1}{2} \mathcal{I}_\mathcal{M} \omega^2 = \frac{1}{2} C_M V_E^2 \tag{29.9}$$

29.4 The Torque Curve

We can characterize the steady-state response of a motor with a fixed appled voltage V_M as a *torque-speed curve*. As shown in Figure 29.4 this curve shows the torque the motor is delivering to the load as a function of motor speed.

At one extreme, if there is no load on the motor (zero torque), and if we ignore motor friction, $V_E = V_M$ and

$$\omega_{\max} = \frac{V_M}{K_M} \tag{29.10}$$

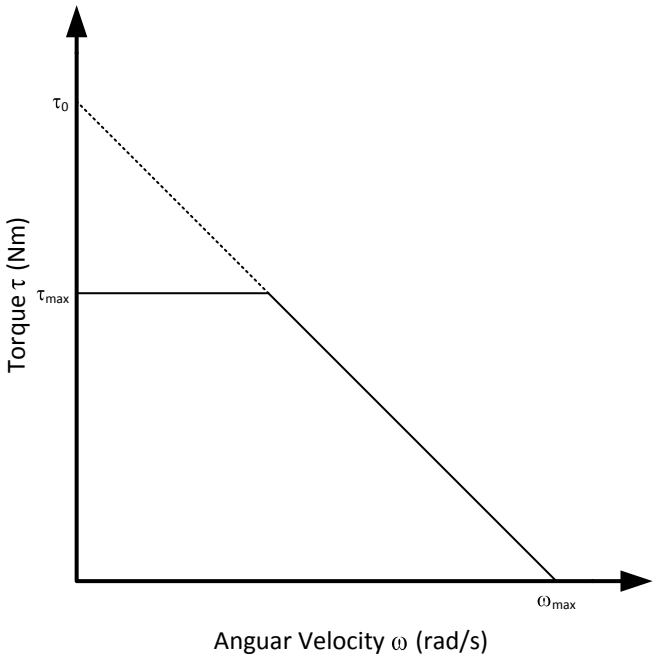

Figure 29.4: Torque vs speed curve for an electric motor. For a given applied voltage, the motor is characterized by its *locked rotor torque* τ_0, and its unloaded speed $\omega_{\max}$. Thermal limits often impose a maximum torque $\tau_{\max} < \tau_0$.

L_M	2	mH
R_M	0.22	Ω
K_M	0.66	$\frac{Vs}{rad}$
$\mathcal{I}_M$	0.083	

Table 29.1: Parameters for a typical brushed, permanent-magnet DC motor.

If we take friction, modeled by R_F into account, we get the expression

$$\omega_{\max} = \frac{R_F V_M}{(R_F + R_M) K_M} \tag{29.11}$$

Thus, if we ignore friction (or if we know both R_F and R_M) we can derive the motor constant K_M from the unloaded speed $\omega_{\max}$.

At the other extreme, if the rotor is stationary, $V_E = K_M\omega = 0$ and so the current drawn by the motor is $I_M = \frac{V_M}{R_M}$. Hence the *locked-rotor torque* is

$$\tau_0 = K_M I_M = \frac{K_M V_M}{R_M} \tag{29.12}$$

Thus, if we know R_M we can also derive K_M from a measurement of τ_0, or alternately if we know K_M we can determine R_M from τ_0. Note that the effective resistance here is the entire resistance around the loop that connects the source to the motor, not just the winding resistance of the motor. The effective resistance includes, for example, the internal resistance of the source and the resistance of the wires connecting the source to the motor.

For many motors, passing a current $i_0 = \frac{V_M}{R_M}$ through the rotor windings for more than a very short period of time will result in destructive temperatures. To avoid burning out the motor, we often cap the torque curve at a maximum continuous torque $\tau_{\max}$ that corresponds to a maximum continuous current $i_{\max}$. It is possible to exceed this value for short periods of time as long as the motor components are kept within their safe temperature zones.

29.5 Step Response

The step response of an unloaded motor is determined by its electrical and mechanical time constants, and observing the response can help determine the motor parameters. In this section we develop a Matlab model of a typical motor and using this model simulate the motor's step response. The parameters of our motor are listed in Table 29.5.

We simulate one time step dt of motor operation using the code shown in Figure 29.5. The two state variables are inductor current `i_m` and motor velocity `omega`. This code computes the back EMF of the motor V_E and from this the drop across the inductor V_L. The inductor voltage is integratd to give the next

```
v_e = omega * k_m ;          % motor emf
v_l = v_m-v_e-i_m*r_m ;      % inductor voltage
i_m = i_m + v_l*dt/l_m ;     % motor current
tau = i_m*k_m ;              % torque
omega = omega + tau*dt/mi ; % angular velocity
```

Figure 29.5: Matlab code to simulate one timestep of a motor

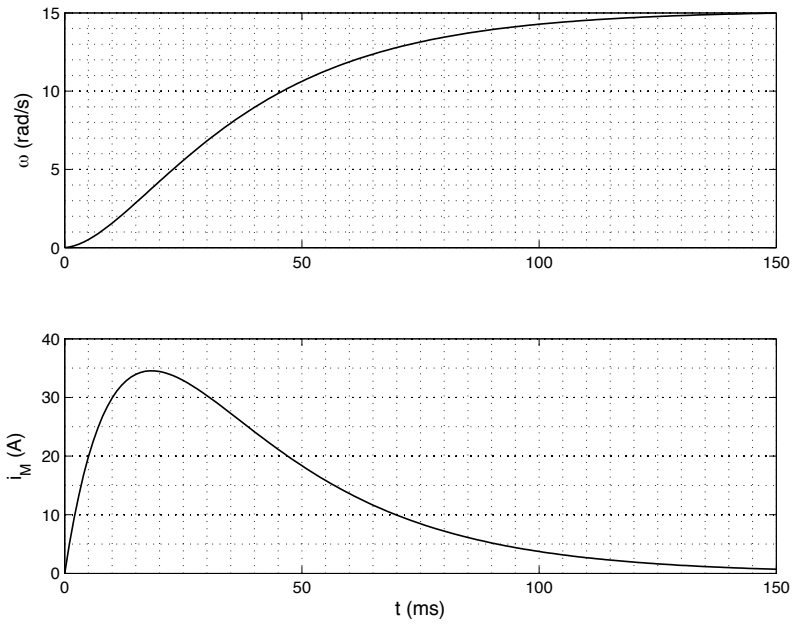

Figure 29.6: Response of a motor with the parameters of Table 29.5 to a 10V step on V_M.

value of inductor current I_M. Torque τ which is proportional to I_M is then integrated to give the next value of ω.

The response of our simulated motor to a 10V step in input voltage is shown in Figure 29.6. The figure shows ω on the top panel and I_M on the bottom panel. Motor inductance L_M can be determined from the initial slope of the current curve. At this point $\omega = 0$ and $I_M = 0$ so $V_E = 0$ and $V_R = R_M I_M = 0$ so the entire V_M drops across L_M thus we have:

$$L_M = \frac{V_M}{di/dt} \tag{29.13}$$

In this case $V_M = 10$V and the initial slope $di/dt = 5$A/ms so we have $L_M = 2$mH which is consistent with Table 29.5.

Ignoring friction we can determine K_M from the steady state value of ω. In this case we have $\omega = 15$rad/s at $V_E = 10$V so we have $K_M = V_E/\omega = 0.66$.

Once K_M is known, the moment of inertia can be determined from the ratio of τ and $d\omega/dt$ at any point on the curve. At 19ms, for example the current reaches a peak value of 34.5A corresponding to a torque of 22.8Nm and $d\omega/dt = 275\text{rad/s}^2$. This gives $\mathcal{I}_M = .083\text{Nms}^2/\text{rad}$.

Finally, the time constant on the tail of the curve, about 40ms, is equal to $\tau_M = R_M C_M$ so we have:

$$R_M = \frac{\tau_M}{C_M} \tag{29.14}$$

$$= \frac{\tau_M K_M^2}{\mathcal{I}_\mathcal{M}} \tag{29.15}$$

$$\tag{29.16}$$

Plugging in our estimated values of $\tau_M = 40$ms, K_M=0.66, and $\mathcal{I}_M$=0.083. We get $R_M = 0.21\Omega$ which is off slightly because we have mis-estimated τ_M. Its actual value is 42ms.

29.6 Motor Drive

A motor *drive* is a circuit used to apply current and voltage to the terminals of a motor. A simple motor drive, shown in Figure 29.7, is just a voltage source V_1 and a switch. The switch is used as a *half-bridge* to connect one motor terminal to either the supply voltage V_1 or to ground.

This configuration is really a (Chapter 2) where the inductor and capacitor are provided by L_M and C_M of the motor itself. If the switch is operated width a duty factor of D in position **a**, then the effective voltage seen by the motor terminals is

$$V_M = DV_1 \tag{29.17}$$

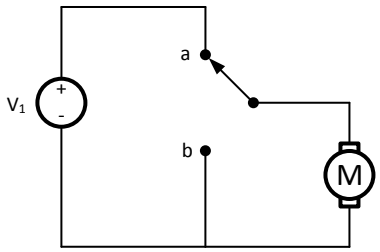

Figure 29.7: A motor drive is a buck converter where the inductor is internal to the motor and the motor's inertia acts as the filter capacitor. The switch is pulse-width modulated with duty factor D to apply an effective voltage $V_M = DV_1$ to the motor.

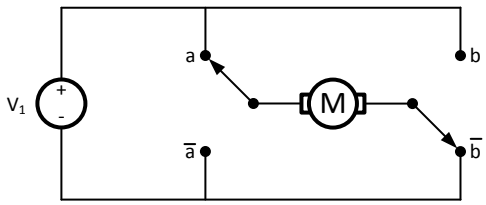

Figure 29.8: To drive the motor in either direction a full-bridge drive circuit is used with a half-bridge connecting each motor terminal to either GND or the supply voltage V_1.

The PWM switching frequency of the motor drive must be high enough to avoid excessive ripple current. Motors are typically wound with fairly heavy gauge magnet wire and have high resistivity at high frequencies. For example, a motor with L_M=1mH and a switching frequency of f=20kHz operating from V_1=100V with D=0.5 will have a ripple current of $I_r = \frac{V_1 D}{L_M f} = 2.5A$.

If the motor needs to be driven in both directions, a full-bridge drive, as show in Figure 29.8, can be used. To apply a positive voltage to the motor, switch **b** is left in the lower position and switch **a** is modulated at the PWM frequency. To apply a negative voltage to the motor the roles are reversed: switch **a** is left in the lower position and switch **b** is modulated at the PWM frequency.

The motor drives of Figure 29.7 and Figure 29.8 often include current sensing circuits so that the motor controller can monitor motor current. Motor velocity and temperature are also often monitored.

In some cases rather than directly measuring a parameter we estimate the parameter from other measurements. Motor velocity, ω, for example, can be esimated by measuring the voltage across the motor with no current applied. Even if current is applied, we can estimate V_E and hence ω if we know the current and subtract the drops across L_M and R_M from the applied V_M.

29.7 Exercises

model an actual motor

equivalent capacitance

Chapter 30

Motor Control

In addition to a motor drive (Section 29.6) a controller is needed to drive a motor to a desired state. Depending on the application we may drive the motor to a desired position, a desired velocity, or to apply a desired torque. For example, a motor controlling an air vent may be commanded to a particular position Θ corresponding to the vent being open or closed. When a motor is geared, this position may involve a number of revolutions, not just an angle between 0 and 2π. When a cruise control is engaged, a motor may be controlled to a constant velocity ω. Finally, the controller in an electric vehicle may drive the motor to produce a torque that is a function of the position of the accelerator pedal.

In performing its control functions, the controller may make use of sensors that measure the motor position, velocity, and/or torque. Current sensors are often used to approximately measure motor torque, and measures or estimates of back EMF V_E may be used to approximate velocity.

In this chapter we develop a simple controller for motor velocity. We then modify this controller to operate within the constraints of a current limit.

30.1 A Naive Controller

Suppose we have a motor with the parameters shown in Table 29.5 and we wish to command it to operate at 30rad/s for 0.3s 20rad/s for 0.3s followed by 40 rad/s for 0.3s. Further assume that our motor has a strict current limit of $I_{\max}$ =20A. That is $|I_M| < I_{\max}$ at all times.

The simplest approach is to use an open-loop controller that applys the control law:

$$V_M = K_M \omega_d \tag{30.1}$$

Where ω_d is the desired velocity.

Our controller, of course sets D, not V_M by applying (29.17) to give $D = V_M/V_1$. Because it is simpler, we will deal with V_M for the remainder of this

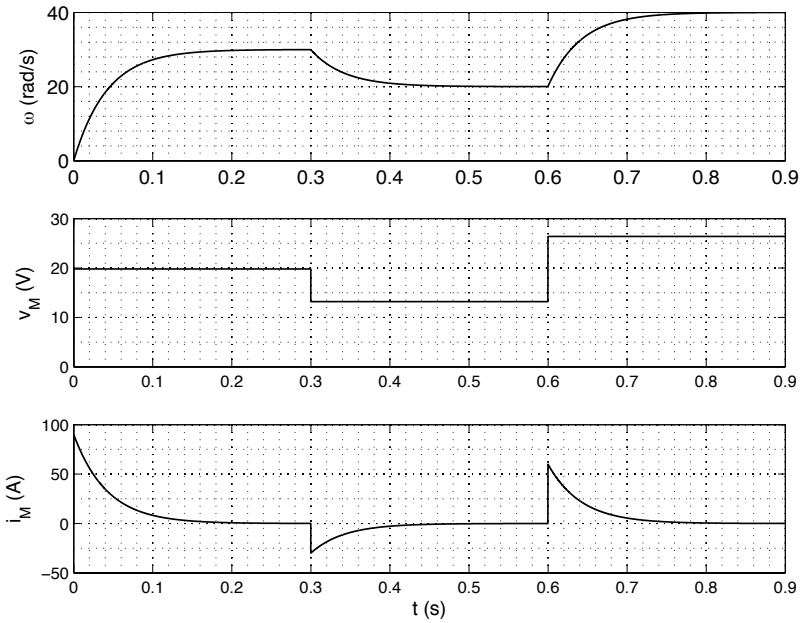

Figure 30.1: Response of our example motor to the naive control law of (30.1).

discussion of control. It is understood that what is actually being set by the drive is D.

The response of simulating our example motor with the naive control law of (30.1) is shown in Figure 30.1. At time zero we step V_M to 19.8V and the motor slowly responds by winding up to 30rad/s. At 0.3s we apply 13.2V and the motor responds by slowing to 20rad/s. At 0.6s we apply 26.4V and the motor accelerates to 40rad/s.

We refer to the process that starts at 0.3s as *regenerative braking*. At 0.3s we drop V_M to 13.2V and the motor slows to 20rad/s. By applying $V_M < V_E$, negative current flows and torque is applied in a direction opposite the rotation of the motor — braking the motor. During this period the motor is acting as a *generator* with power flowing from the motor to supply V_1. We could brake the motor harder — apply a stronger reverse torque — by dropping V_M even lower, or by applying a negative V_M using the drive of Figure 29.8.

While we think of the motor drive as a buck converter — with the inductor and capacitor provided by the equivalent circuit of the motor, during regenerative braking the drive is acting as a boost converter. During this period it is taking current at a potential of V_E and boosting it to V_1 to return energy to the supply. The motor's inductance L_M serves as the inductor of the boost converter.

```
err = omega_d - omega ;
ierr = ierr + err*dt ;
v_m = err*p + ierr*q ;
```

Figure 30.2: Matlab code for PI control law

Looking just at the ω panel of the figure our naive approach to motor control doesn't look too bad. The motor goes to each commanded speed with a delay determined by the motor time constant $\tau_M = R_M C_M$. In fact this plot is just three copies of the motor step response Figure 29.6 appropriately translated.

Unfortunately this naive control scheme has two problems. First it exceeds our current limit by a large margin in both directions. Current I_M reaches extreme values of +90A and -30A. Second, ignoring the current for the moment, the response is not as fast as it could be with a more sophisticated controller.

30.2 PI Motor Control

Since our motor drive is really just a buck converter we can apply the PI control law we developed in Chapter 5 to control this drive. Regulating the velocity ω of the motor is exactly analgous to regulating the output voltage V_C of the buck regulator. The Matlab code for a PI control law for our motor is shown in Figure 30.2. We leave the derivation of the gain constants `p` and `q` as an exercise.

The simulated response of our example motor to the PI control law is shown in Figure 30.3. The good news is that with the PI controller the motor response is much faster. However, achieving this fast response has made our over-current problem even worse. Current I_M now has extremes of 245A and -80A.

30.3 Current-Limited Control

To achieve a rapid motor response while staying within our current limit requires that we regulate motor current. One approach to this is shown in the Matlab code of Figure 30.4. Here we assume we have exact knowledge of ω, R_M, and K_M and hence can compute bounds on V_M that we must stay within to prevent excessive current. The control code computes V_M using a PI control law, just as in Figure 30.2 and then clamps V_M to the calculated limits.

The last three lines of code are needed to prevent a control pathology known as *integral wind up*. If the control output of a controller with an integral error term is clamped (as we are doing here), the integral error will continue to grow without bound. When the clamp is eventually released the large integral error will cause an overshoot, often quite large, in the response. To avoid integral windup, whenever our clamp overrides the value of V_M, we also clamp `ierr` to the value that would produce our clamped V_M if there were no proportional

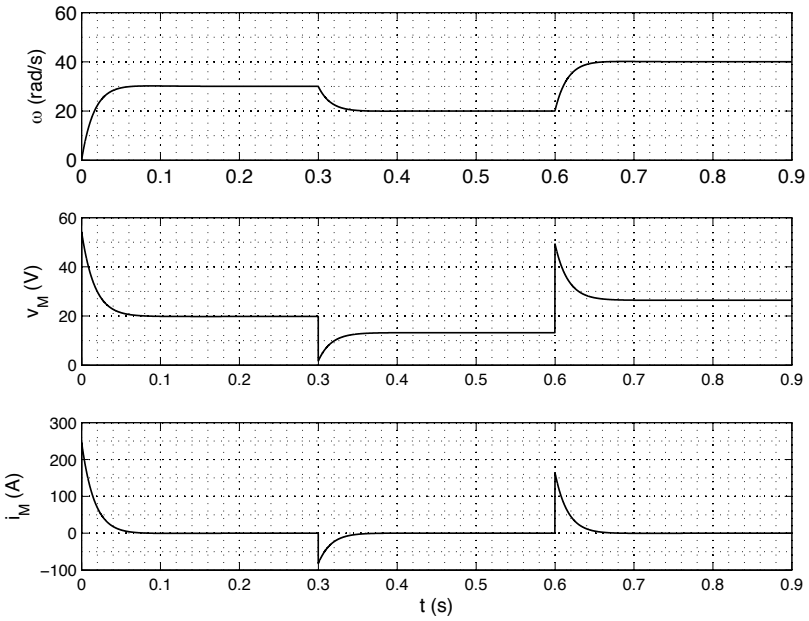

Figure 30.3: Response of our example motor to the PI control law of Figure 30.2. Response is faster than the naive control law but the excess current problem is even larger.

```
% original PI controller
err = omega_d - omega ;
ierr = ierr + err*dt ;
v_m_1 = err*p + ierr*q ;

% now apply current limit
v_m_max = i_max*r_m + v_e ;
v_m_min = -i_max*r_m + v_e ;
v_m = min(v_m_max, max(v_m_min, v_m_1)) ;
if(v_m ~= v_m_1)
  ierr = v_m/q ;
end
```

Figure 30.4: Matlab code for PI control law with current limiting. The override of `ierr` is used to prevent *integral wind up* when the control is limited.

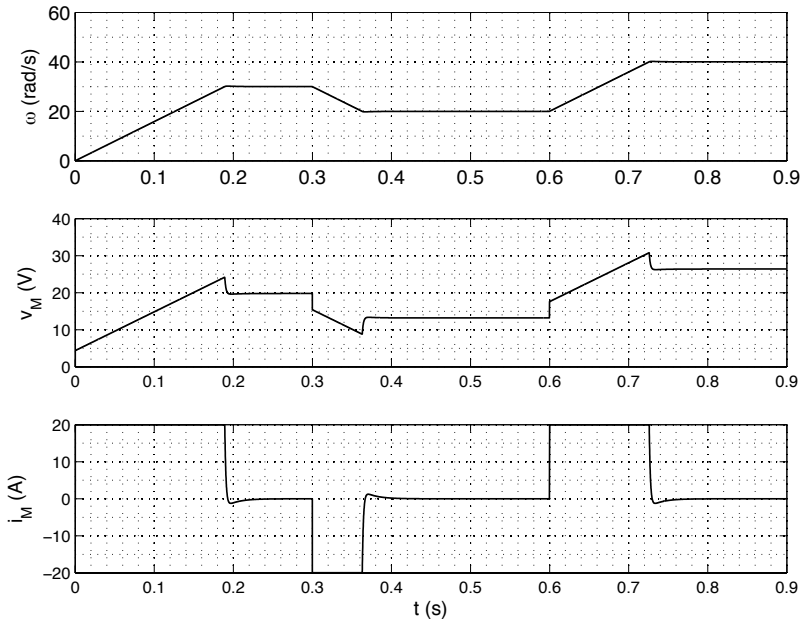

Figure 30.5: Response of our example motor to the current-limited PI control law of Figure 30.4. The motor velocity ω slews to each set point as fast as possible subject to the current constraint.

error (V_M/q). With this approach the integral error has a value that gives continuity in V_M when the clamp is released.

The simulated response of our example motor to the current-limited PI control law of Figure 30.4 is shown in Figure 30.5. With this control law the motor velocity ω slews as fast as possible with the current constraint to each set point. During these transitions the current is clamped to its limits of 20A or -20A giving the largest possible torque and hence the fastest possible response time. At the end of each transition the clamp is removed and the PI control law resumes control - smoothly settling on the set point with no overshoot.

At first glance, one might think that you could get the same response using a *bang-bang* controller. A bang-bang controller would set the current to 20A whenever ω is below the set point and to -20A whenever ω is above the set point. The problem with a bang-bang controller is that you would either have a dead zone around the set point — where errors are not corrected — or the current would oscillate constantly between -20A and +20A in response to even very small disturbances. The resulting extreme torques would result in objectionable vibrations.

Our current-limited PI controller is in effect a bang-bang controller for large

errors (where the current limit takes effect) but transitions smoothly to a linear control law for small errors. Hence it suffers from neither a dead zone nor from oscillating current — and torque — extremes.

Figure 30.4 clamps current by computing a set of voltage limits that correspond to the current limits and then clamping V_M to these limits. An alternate approach, that can be used when motor parameters are not exactly known, is to build a current regulator that runs in parallel with the velocity regulator. Whenever I_M is within limits the velocity regulator sets V_M. The current regulator takes over when I_M exceeds the current limit. We leave implementing such a current regulator as an exercise.

30.4 Exercises

derive the coefficients

run the current limit controller without integral wind up prevention

build a current regulator

develop a controller for motor position simulate in Matlab

develop a controller for motor torque simulate in Matlab

Chapter 31

Brushless Permanent Magnet Motors

While the brushed permanent magnet motor is simple to understand and to control it has the disadvantage of having a commutator and brushes. These parts add to the expense of the motor and the brushes wear out with operation requiring service.

The brushless permanent magnet motor eliminates the commutator and brushes and replaces them by electronic swithches that connect the windings (which are moved to the stator) to the supply. The behavior of the brushless permanent magnet motor is exactly the same as the brushed motor with one exception: the applied voltage and current must be rotated so they are perpendicular to the current position of the rotor. This is done by modulating the drive to multiple, usually three, phases. As long as this orientation is maintained, the behavior is exactly as described in Chapter 29: torque is proportional to current and voltage is proportional to angular velocity.

In this chapter we introduce the brushless permanent magnet motor. We show how it can be driven either using hard commutation or by synthesizing a multi-phase sine wave using an inverter Chapter 15. We adapt our motor drive Section 29.6 and our motor control methods Chapter **??** to brushless motors.

31.1 Motor Geometry

Figure 31.1 shows a brushless permanent magnet motor in schematic form. The roles of the stator and rotor are reversed compared to the brushed motor of Figure 29.2. With the brushless motor the rotor contains the permanent magnet and the stator contains the windings. The windings are divided into multiple phases — three in this case — to control the direction of the applied field.

To analyze the brushless motor we can consider separately the field and torque produced by the phase currents and the back EMF voltage induced by the rotating magnet. Let us first consider the back EMF.

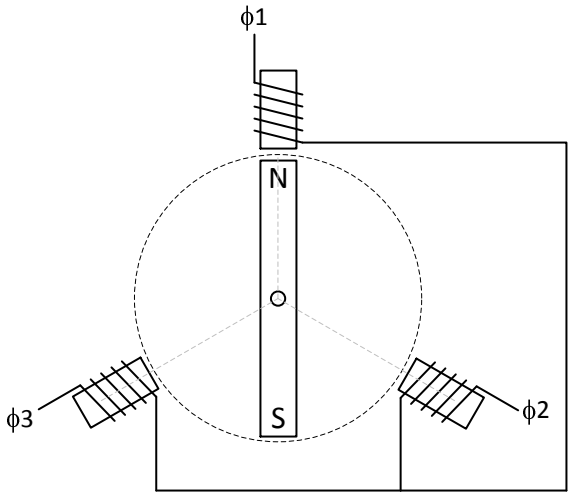

Figure 31.1: A three-phase brushless permanent-magnet motor has a stator with three windings that generate a rotating magnetic field. The rotor contains a permanent magnet that generates torque in response to the stator's magnetic field. As the permanent magnet rotates, it also induces voltage in the stator windings.

31.1.1 Back EMF

Suppose the rotor is currently at angle θ, rotating with angular velocity $\omega = \frac{d\theta}{dt}$, and generates a magnetic field across the *gap* to the stator with flux density B_r. The position shown in Figure 31.1 with the north pole of the rotor straight up corresponds to $\theta = 0$. The field across the gap as a function of angle ψ is $B(\psi) = B_r \cos(\theta - \psi)$, and if the effective area captured by each phase winding is A the flux is $\phi(\psi) = B_r A \cos(\theta - \psi)$. Then for phase winding i at angle ψ_i, the voltage induced is

$$\begin{aligned} V_i &= \frac{d\phi_i}{dt} \\ &= -\omega B_r A \sin(\theta - \psi_i) \\ &= -\omega K_M \sin(\theta - \psi_i). \end{aligned} \tag{31.1}$$

From (31.1.1) we see that the voltage induced across a phase varies sinusoidally with θ and has a maximum positive value when the rotor is 90° behind the phase position, i.e., when $\theta - \psi = -\pi/2$.

If we consider the voltage at angle α to be a weighted sum of the N phase voltages at angles ψ_i:

$$\begin{aligned} V(\alpha) &= \sum_{i=1}^{N} V_i \cos(\alpha - \psi_i) \\ &= -\omega K_M \sum_{i=1}^{N} \sin(\theta - \psi_i) \cos(\alpha - \psi_i). \end{aligned} \tag{31.2}$$

Then the voltage at angle $\theta + \pi/2$ for rotor position θ is

$$V(\theta + \pi/2) = \omega K_M \sum_{i=1}^{N} \sin^2(\theta - \psi_i) \tag{31.3}$$

And for N=3 and angles ψ of 0, $\pi/3$, and $-\pi/3$ we have

$$V(\theta + \pi/2) = 1.5 K_M \omega. \tag{31.4}$$

Thus, we see that rotated into a frame of reference 90° ahead of the rotor, the back EMF of the motor remains proportional to ω.

31.1.2 Torque

It is easier to explain torque if we redraw the brushless permanent magnet motor with the three phases energizing coils that surround the rotor as shown in

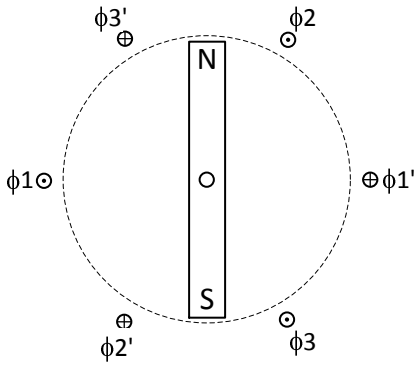

Figure 31.2: Alternate view of the brushless permanent magnet motor with the stator drawn with coils that surround the rotor rather than the *salient poles* of Figure 31.1.

Figure 31.2. Each phase is shown as two wires carrying current in opposite directions. While this appears different than the phases energizing *salient poles* as in Figure 31.1, the two are functionally identical. In both figures, a current i_1 in ϕ_1 induces a vertical magnetic field and the induced voltage in ϕ_1 is proportional to ω and peaks when the north pole of the rotor is pointing to the left.

Recall that the field at angle ψ is $B(\psi) = B_r \cos(\theta - \psi)$. Phase ϕ_i oriented at ψ_i — zero radians for ϕ_1 — has two wires oriented at $\psi_i - \pi/2$ and $\psi_i + \pi/2$. The field seen by the first wire is $B(\psi_i - \pi/2) = B_r \sin(\theta - \psi_i)$ and the field seen by the second wire is $B(\psi_i + \pi/2) = -B_r \sin(\theta - \psi_i)$. Thus, if current i_i is flowing in ϕ_i, then (**??**) gives a force on each wire of ϕ_i of $F = i_i L B_r \sin(\theta - \psi_i)$. Converting this to torque gives

$$\begin{aligned} \tau &= -i_i 2 L B_r r \sin(\theta - \psi_i) \\ &= -i_i K_M \sin(\theta - \psi_i) \end{aligned} \tag{31.5}$$

We can apply current i at an angle α by applying current $i_i = \frac{2}{3} i \cos(\alpha - \psi_i)$ to phase ϕ_i. If we apply the current at $\alpha = \theta + \pi/2$, $i_i = \frac{2}{3} i \sin(\theta - \psi_i)$. we have

$$\tau = i K_M \tag{31.6}$$

Thus, we see that once we rotate both the voltage and current to the frame of reference of the rotor (at angle $\theta + \pi/2$) the behavior of the brushless permanent magnet motor is idential to that of the brushed motor.

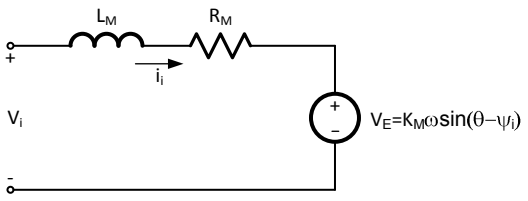

Figure 31.3: Equivalent circuit for one phase of a brushless permanent-magnet motor.

31.2 Equivalent Circuit

To our motor controller, each phase of a brushless permanent magnet motor has an I-V characteristic identical to that of the equivalent circuit shown in Figure **??**. The voltage source V_E represents the back-EMF of the motor which varies sinusoidially with rotor position θ. The phase terminals are connected to this voltage source via the inductance L_M and resistance R_M of the phase windings. In many motors with *Y-connected* phases, only one terminal of each phase is brought out. The other terminal of each phase is tied to a common point.

The power $P_i = i_i V_E$ into or out of the voltage source V_E represents the conversion between electrical and mechanical power by phase i. If P_i is positive phase i is converting electrical power to mechanical power — operating as a motor. When P_i is negative, the motor is operating as a generator — converting mechanical power to electrical power.

31.3 Brushless Motor Drive

Conceptually we control a brushless permanent magnet motor as show in Figure 31.4, by generating a motor voltage V_M as it were a brushed motor and then generating a phase voltage V_i for each phase ϕ_i to rotate V_M into the frame of reference of the rotor.

$$V_i = \frac{2}{3} V_M sin(\psi_i - \theta) \tag{31.7}$$

This is just conceptual of course. We don't actually generate any of these voltages. Instead we control the motor using three *half-bridge* switches as shown in Figure **??**. Each switch can connect its corresponding phase to ground or to supply voltage, V_S. We apply a pulse-width modulated signal with duty factor D_i to phase i to generate an average voltage that corresponds to V_i. With the conceptual phase voltages V_i in the range $[-V_S, V_S]$ we have:

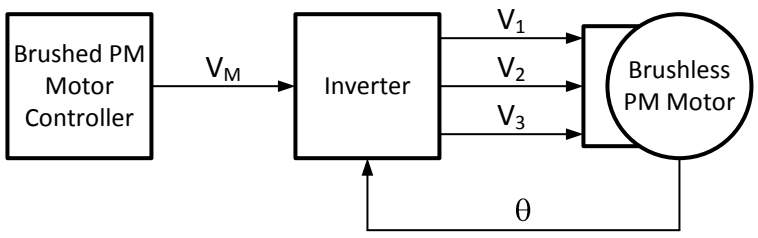

Figure 31.4: We conceptually divide a brushless motor controller into two blocks. The first block computes the motor drive voltage V_M that would be required to drive an equivalent brushless motor. The second block generates phase voltages V_1, V_2, and V_3 to rotate V_M to $\theta + \pi/2$ for motor position θ.

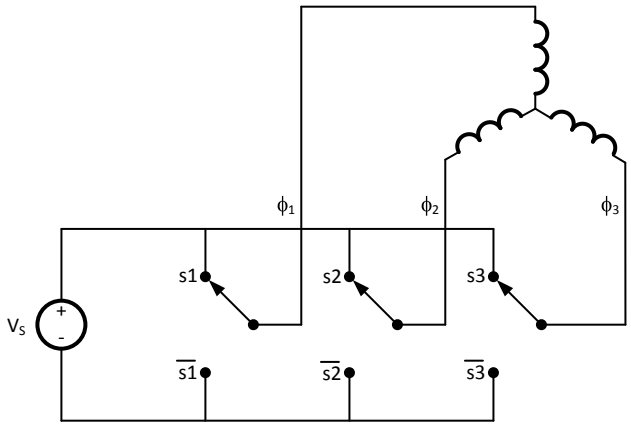

Figure 31.5: A three-phase motor driver consists of three *half bridges* that generate the three phase voltages. For each phase ϕ_i switch s_i generates a pulse-width modulated signal with duty factor D_i and mean value $D_i V_1$.

$$D_i = 0.5(1 + \frac{V_i}{V_S}) \tag{31.8}$$

$$= 0.5(1 + \frac{V_M sin(\psi_i - \theta)}{V_S}) \tag{31.9}$$

$$\tag{31.10}$$

This converts each conceptual voltages V_i to an actual average phase voltage of $\hat{V_i}$ according to

$$\hat{V_i} = 0.5(V_S + V_i) \tag{31.11}$$

That is $V_i = 0$ maps to $\hat{V_i} = 0.5V_S$, $V_i = V_S$ maps to $\hat{V_i} = V_S$, and $V_i = -V_S$ maps to $\hat{V_i} = 0$.

Because the three phases are distributed evenly around the unit circle the average of the three phase voltages is always $V_S/2$ and this is the voltage at the common point at the center of the three windings.

In some cases it is desirable to minimize the phase voltages. We can accomplish this by subtracting the minimum duty factor from all of the duty factors. This effectively lowers the common mode voltage so that one phase is always at ground potential. The modified duty factors D'_i are calculated as

$$D_{\min} = \min_{i=1}^{n}(D_i)$$

$$D'_i = D_i - D_{\min} \tag{31.12}$$

This transformation does result in an AC component on the common node of the motor. If this is objectinable the minimum duty factor of any phase over an entire cycle can be subtracted from each phases duty factor instead of the instantaneous minimum.

31.4 Hard Electronic Commutation

[to be written]

31.5 Excercises

Chapter 32

AC Induction Motors

AC induction motors are among the most common motors in use for higher-output applications. They have attractive power for a given weight and size and do not require rare-earth magnets. They tend to be slightly less efficient than a comparable brushless permanent magnet motor because they introduce rotor condution losses.

32.1 Motor Geometry

As shown in Figure 32.1 an AC induction motor has a stator identical to a brushless permanent magnet motor but substitutes a *sqirrel cage* rotor of conductive bars that are shorted together at both ends for the permanent magnet rotor. The motor can be constructed either with coils wound on *salient poles* as shown in Figure 32.1 or with a stator consisting of coils surrounding the entire rotor as showni n Figure 32.2. As in the brushless motor the stator is driven with a three-phase sinusoidal signal to produce a rotating magnetic field with frequency ω_s. If the rotor is turning with frequency ω_r, each rotor bar sees a magnetic field at the difference freqency $\omega_d = \omega_s - \omega_r$. This changing magnetic field induces a voltage across the rotor bars and this voltage in turn causes current to flow in the bars. This rotor current interacts with the magnetic field to generate torque.

In effect the coupling between the stator and the rotor in an induction motor acts as a transformer (Section 9.1) generating voltage and current in the rotor in response to voltage applied to the stator. The rotor current then generates a motor torque and, if the rotor is moving, a back EMF.

32.2 Induction Motor Operation

Operation of the induction motor can be understood in terms of rotating voltage, current, and fluxes as shown in Figure 32.3. Here we represent AC quantities in the two-dimensional plane as an amplitude and a phase. For example, if we

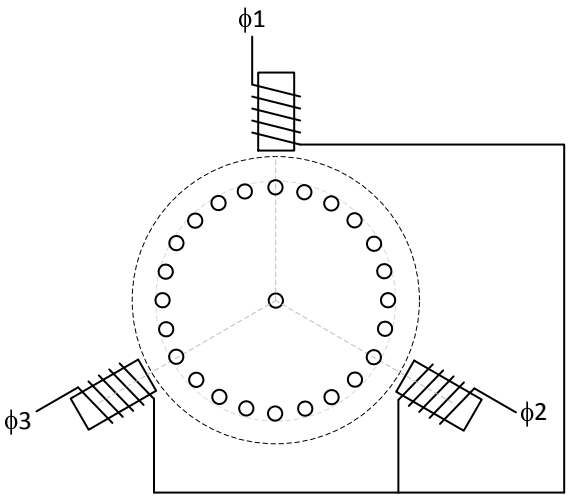

Figure 32.1: A three-phase AC induction motor has a stator identical to the brushless permanent magnet motor of Figure 31.1 but substitutes a *squirrel cage* rotor of shorted conductive bars for the permanent magnet rotor.

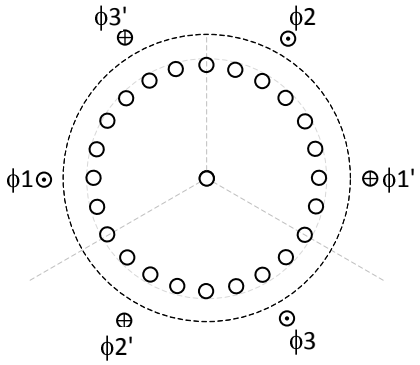

Figure 32.2: A three-phase AC induction motor drawn with the stator consisting of three coils (shown in cross section). Operation is identical to the salient-pole rotor shown in Figure 32.1

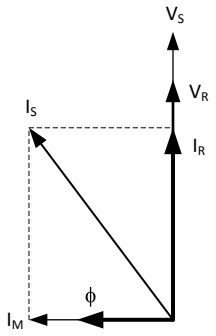

Figure 32.3: Operation of the AC induction motor can be understood by representing the key quantities as *phasors* — i.e. in polar coordinates as magnitude and phase (angle) relative to the sinusoidal stator voltage V_s at frequency ω_s.

apply a three-phase sinusoidal voltage to the stator so $V_1 = V_a \sin(\omega_s t + \theta)$, $V_2 = V_a \sin(\omega_s t + \theta - 2\pi/3)$, and $V_3 = V_a \sin(\omega_s t + \theta + 2\pi/3)$, we simply denote this as $V_s = V_a \angle(\theta)$.

If we ignore the stator resistance and inductance for the moment, applying a sinusoidal three-phase voltage $V_s \angle 0$ with frequency ω_s results in a magnetic flux across the gap of

$$\begin{aligned} \phi &= \int V_s dt \\ &= \frac{V_s}{\omega_s} \angle -\pi/2 \end{aligned} \tag{32.1}$$

A magnetizing current flows in the stator proportional to this flux.

$$I_m = \frac{V_s}{L_m \omega_s} \angle -\pi/2 \tag{32.2}$$

Where L_m is the magnetizing inductance of the motor. Ignoring the rotor, our induction motor is just an inductor and thus the current I_m lags the voltage V_s by $\pi/2$.

The rotor sees the flux varying sinusoidally with frequency ω_d. Assuming the *turns ratio* between the stator and rotor is 1:1, this induces a rotor voltage (in the stator frame of reference) across each rotor bar of

$$\begin{aligned} V_r &= \frac{d\phi}{dt} \\ &= \frac{V_s \omega_d}{\omega_s} \angle 0 \\ &= \sigma V_s \angle 0. \end{aligned} \tag{32.3}$$

As expected for a transformer, our rotor voltage is in-phase with the stator voltage. Unlike a stationary transformer our rotor voltage is scaled by $\sigma = \frac{\omega_d}{\omega_s}$ a quantity referred to as the *slip* of the motor[1].

If the rotor bars have an aggregate resistance of R_r the rotor voltage V_r causes a rotor current with value

$$\begin{aligned} I_r &= \frac{V_r}{R_r} \angle 0 \\ &= \frac{\sigma V_s}{R_r} \angle 0 \end{aligned} \tag{32.4}$$

In practice rotor inductance will cause I_r to slightly lag V_s. We ignore that effect for now.

Because of the scaling of rotor voltage by σ, the rotor resistance reflected back to the stator — i.e. the rotor resistance seen by V_s is

$$\begin{aligned} R'_r &= \frac{R_r}{\sigma} \\ &= \frac{\omega_s R_r}{\omega_d} \end{aligned} \tag{32.5}$$

When there is no slip, $\sigma = 0$, R'_r is infinite and no rotor current flows. When the rotor is locked $\sigma = 1$ and $R'_r = R_r$. For other cases of positive slip $R'_r > R_r$. If $\omega_r > \omega_s$, we have a negative resistance, $R'_r < -R_r$. In this case the motor is operating as a generator.

The overall stator current is the vector sum of the magnetizing current and the rotor current reflected to the stator.

$$I_s = I_m + I_r \tag{32.6}$$

As shown in Figure 32.3, This stator current is at a negative angle intermediate between 0 and $-\pi/2$.

Current I_r interacts with the magnetic field to produce a torque

[1] Many texts use the variable s to represent slip. We use σ here to avoid confusion with laplace transforms.

$$\begin{aligned}\tau &= I_r\phi \\ &= \frac{\sigma V_s^2}{\omega_s R_r} \qquad (32.7) \\ &= \frac{V_s^2\omega_d}{\omega_s^2 R_r} \qquad (32.8)\end{aligned}$$

Torque increases quadratically with V_s since both the ϕ and I_r depend linearly on V_s. Torque decreases with increased stator frequency ω_s because ϕ decreases with frequency. Note that τ is a DC value, not an AC signal, so it has no phase.

The output power delivered by the motor to its shaft (or received from the shaft if negative) is

$$\begin{aligned}P_O &= \tau\omega_r \\ &= I_r\phi\omega_r \\ &= I_r V_s\frac{\omega_r}{\omega_s} \\ &= I_r^2 R_r\frac{\omega_r}{\omega_d} \\ &= I_r^2 R_r\frac{1-\sigma}{\sigma} \qquad (32.9)\end{aligned}$$

We rewrite P_O in the form I^2R to facilitate development of our equivalent circuit in the next section.

32.3 Equivalent Circuit

The behavior of an AC induction motor can be analyzed using the equivalent circuit of Figure 32.4. The circuit represents one of the three phases of the motor.

Voltage V_s is applied to one stator phase. Stator current I_s passes through the stator winding resistance R_s and leakage inductance L_s before splitting. Magnetizing current (32.2) passes through magnetizing inductance L_m. The reflected rotor current passes through the leakage inductance of the rotor L_r and the scaled rotor bar resistance R_R/σ which we have split into loss and output components. The dotted line represents the air-gap between the stator and the rotor.

From (32.4) we see that the effective rotor resistance seen by V_s is $R_r' = R_r/\sigma$. We split this scaled rotor resistance into two components to separate the component of the resistance that represents conduction losses in the rotor R_{rl} from the component that represents the output power produced by the motor R_{ro}. From (32.9) we see that the motor power is equivalent to the power dissipated by passing the rotor current through a resistor with value:

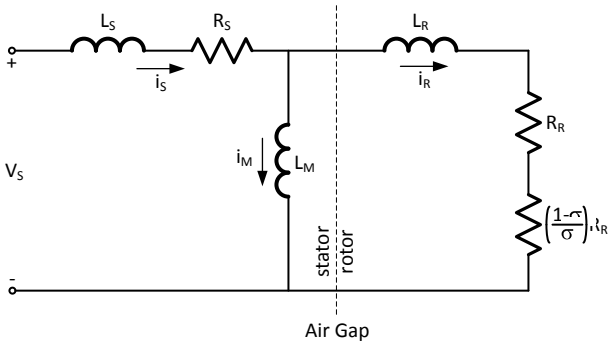

Figure 32.4: Equivalent circuit for one phase of an AC induction motor. The stator voltage induces a magnetizing current I_m (32.2) across magnetizing inductance L_m and a rotor current (reflected to the stator) I_r (32.4) across a scaled rotor resistance R_r. We add stator resistance R_s and stator and rotor leakage inductances L_s and L_r to complete our model.

$$\begin{aligned} R_{ro} &= \left(\frac{\omega_r}{\omega_d}\right) R_r \\ &= \left(\frac{1-\sigma}{\sigma}\right) R_r \end{aligned} \tag{32.10}$$

The remaining portion of R'_r is:

$$\begin{aligned} R_{rl} &= R'_r - R_{ro} \\ &= \left(\frac{1}{\sigma} - \frac{1-\sigma}{\sigma}\right) R_r \\ &= R_r \end{aligned} \tag{32.11}$$

This is as expected since the power dissipated by I_r flowing through the actual rotor resistance R_r is in fact dissipated in the resistance of the rotor.

32.4 Torque Curve

In operating a motor we must trade off torque against efficiency, $\nu = P_o/P_i$. Ignoring stator losses, we can calculate the efficiency of our motor from the resistive voltage divider of Figure 32.4:

$$\begin{aligned}\nu &= \frac{R_{ro}}{R_{ro}+R_{rl}} \\ &= \frac{\frac{\omega_r}{\omega_d}}{\frac{\omega_r}{\omega_d}+1} \\ &= \frac{\omega_r}{\omega_r+\omega_d} \\ &= \frac{\omega_r}{\omega_s} \\ &= 1-\sigma \end{aligned} \tag{32.12}$$

We see that efficiency is determined entirely by slip. When $\sigma = 1$ (locked rotor) there is no output power and $\nu = 0$. When $\sigma = 0$ (synchronous operation) there is no rotor current and hence no rotor losses — unfortunately there is also no torque. Slip is required to generate torque. As slip increases we get more torque for a given V_s but we also operate less efficiently.

From (32.7) we see that torque is directly proportional to slip. Thus to generate torque, we must sacrifice efficiency and, for fixed V_s and ω_r, as torque increases, efficiency drops.

Figure 32.5 shows torque τ and efficiency ν as a function of ω_s and Figure 32.6 shows the same two curves plotted as a function of σ. Efficiency is maximum $\nu = 1$ at $\omega_s = \omega_r$ — i.e., at $\sigma = 0$. As σ increases efficiency falls while torque increases. Torque reaches a maximum at $\sigma = 0.5$ where $\nu = 0.5$. Typically a motor is operated at much lower slip, and hence much higher efficiency.

Of course the efficiency ν here is just the rotor efficiency — accounting only for the resistive loss in the rotor windings. The actual motor efficiency will be considerably lower due to friction and windage, stator winding losses, and core losses.

32.5 Induction Motor Control

The drive used to control a three-phase (or N-phase) induction motor is identical to the drive used for a brushless permanent magnet: a three-phase inverter. As shown in Figure 31.5, a half-bridge driven with a PWM signal generates an appropriate average voltage, $\sin(\omega_s t - \psi_i)$, for each phase ϕ_i.

While the drive itself is identical, the algorithm used to control the drive differs significantly between the two motors. For a brushless motor, we determine the rotor position θ and apply a current at $\theta - \pi/2$ to generate torque. In contrast, we don't care about the position of an induction motor. Instead we determine rotor velocity ω_r and determine the optimum combination of voltage V_s and stator frequeny ω_s to apply the desired torque τ per equation (32.8). Of the many possible solutions, one is chosen to optimize a figure of merit subject to constraints.

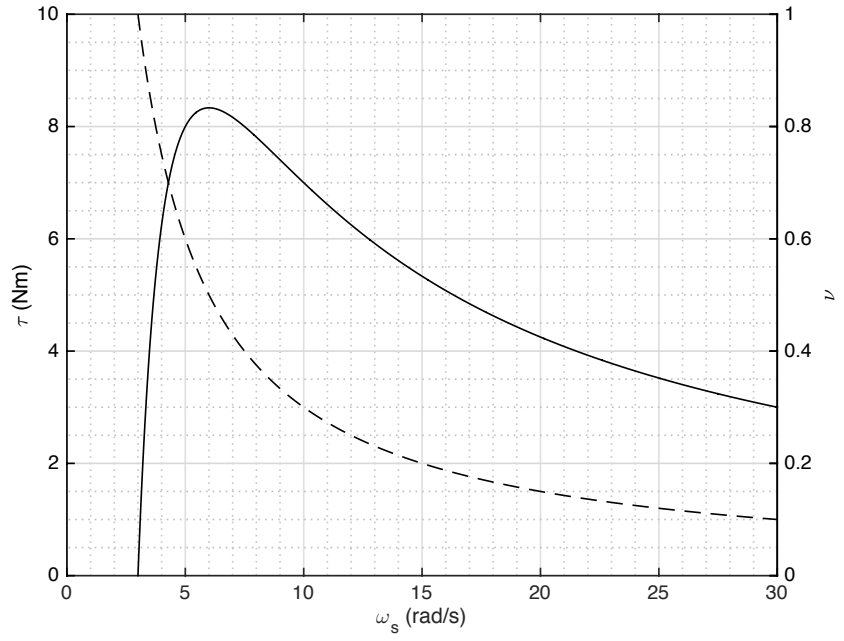

Figure 32.5: Torque τ and efficiency ν as a function of ω_s for $\omega_r = 3$. Maximum efficiency, $\nu = 1$, occurs at $\sigma = 0$ where $\tau = 0$. As σ increases ν drops and τ increases, reaching a maximum at $\sigma = 0.5$ and $\nu = 0.5$.

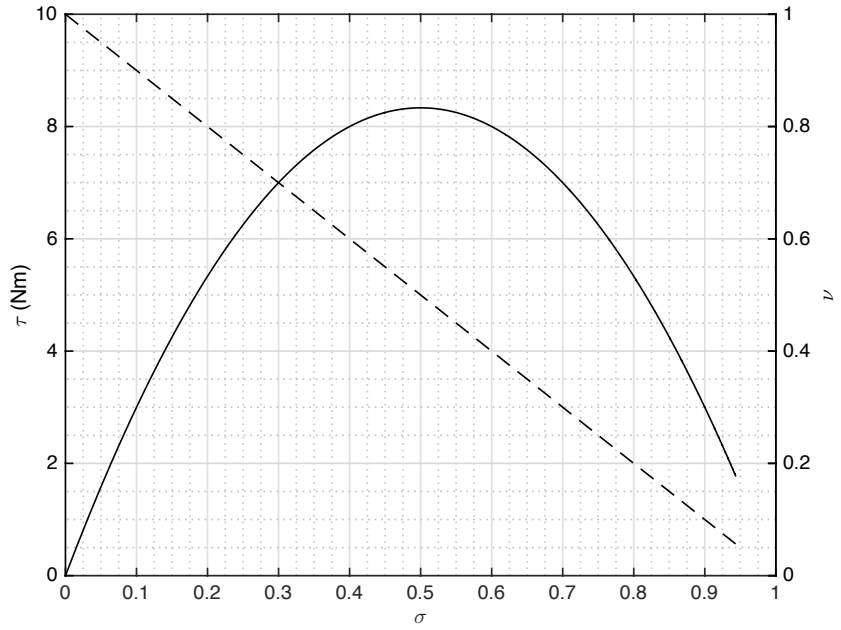

Figure 32.6: Torque τ and efficiency ν as a function of ω_s for $\omega_r = 3$. Maximum efficiency, $\nu = 1$, occurs at $\sigma = 0$ where $\tau = 0$. As σ increases ν drops and τ increases, reaching a maximum at $\sigma = 0.5$ and $\nu = 0.5$.

A typical motor controller will be subject to the constraints:

$$V_s \leq V_{\max} \tag{32.13}$$
$$I_s \leq I_{\max} \tag{32.14}$$

We can think of the constraints as establishing a feasible region in the ω_s, V_s plane (Figure 32.7). The voltage constraint (32.14) specifies that all possible solutions must be below the $V_{\max}$ line in this plane. The current constraint (32.14) applied to the magnetizing current (32.2) specifies that all solutions be to the right of an $I_{\max}$ diagonal:

$$\omega_s \geq \frac{V_s}{L_m I_{\max}} \tag{32.15}$$

These two constraints are shown as thick dotted lines in the figure. The point where the two constraints intersect occurs at

$$\omega_x = \frac{V_{\max}}{L_m I_{\max}} \tag{32.16}$$

For an AC induction motor operating as a motor ($\sigma > 0$) the lowest σ, and thus highest ν is achieved at the lowest ω_s that stays in the feasible region. This optimal point will occur at a boundary of the feasible region. For $\omega_s \leq \omega_x$ the controller is constrained by the diagonal $I_{\max}$ line. For $\omega_s \geq \omega_x$ the controller is constrained by the horizontal $V_{\max}$ line.

Figure 32.7 shows a motor with $\omega_x = 1.67$ operating with $\omega_r = 1$. The voltage and current constraints are shown as dotted lines. Loci of constant positive torque are shown as solid curves in the figure. Loci of constant negative torque are shown as dashed curves in the figure.

The third solid curve from the bottom corresponds to $\tau = 1$ any V_s, ω_s combination along this curve will produce $\tau = 1$. However, the intersection of this curve with the $I_{\max}$ constraint line gives the most efficient operating point for $\tau = 1$. We can calculate this current-constrained operating point, V_{sc}, ω_{sc} as:

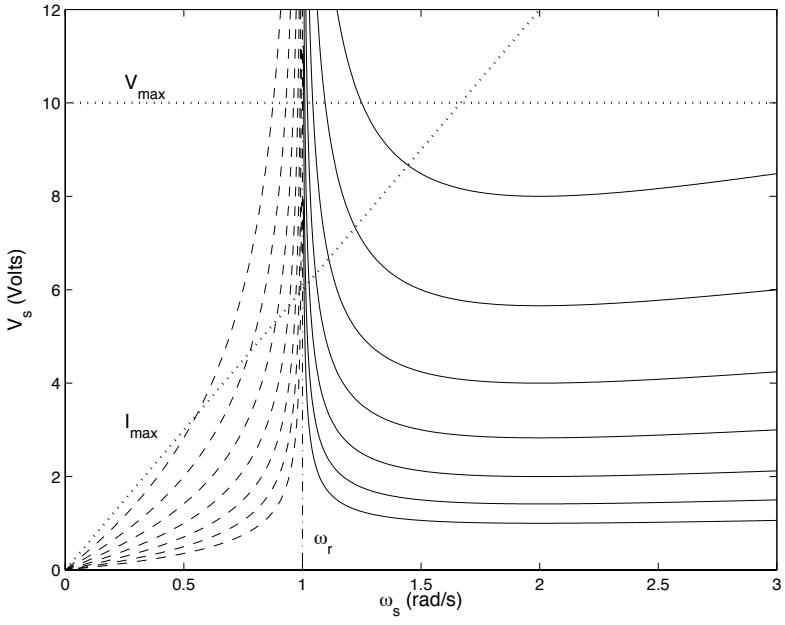

Figure 32.7: Lines of constant torque in the ω_s,V_s plane for $\omega_r = 1$. The thick dotted lines represent the $V_{\max}$ and $I_{\max}$ constraints. The solid lines are positive lines of constant torque with values (starting at the bottom) of 0.25, 0.5, 1, 2, 4, 8, and 16. The dashed lines are negative lines of constant torque with the same absolute values. At this low ω_r the controller is constrained by $I_{\max}$.

$$V_{sc} = L_m I_{\max} \omega_{sc} \tag{32.17}$$

$$\begin{aligned} \tau &= \frac{\omega_{sc}^2 L_m^2 I_{\max}^2 \omega_{dc}}{\omega_{sc}^2 R_r} \\ &= \frac{L_m^2 I_{\max}^2 \omega_{dc}}{R_r} \\ \omega_{dc} &= \frac{\tau R_r}{L_m^2 I_{\max}^2} \end{aligned}$$

$$\omega_{sc} = \frac{\tau R_r}{L_m^2 I_{\max}^2} + \omega_r \tag{32.18}$$

$$V_{sc} = \frac{\tau R_r}{L_m I_{\max}} + \omega_r L_m I_{\max} \tag{32.19}$$

At this operating point we have:

$$\sigma = \frac{\tau R_r}{\tau R_r + \omega_r I_{\max}^2 L_m^2} \tag{32.20}$$

$$\nu = \frac{\omega_r I_{\max}^2 L_m^2}{\tau R_r + \omega_r I_{\max}^2 L_m^2} \tag{32.21}$$

Figure 32.8 shows the case where $\omega_s > \omega_x$ and the controller is constrained by voltage rather than current. In this case $\omega_r = 3$ and the constant-torque curves are all constrained by the $V_{\max}$ line. To solve for the point of intersection we substitute (32.13) into (32.8) to give a quadratic equaiton that can be solved for the voltage-constrained operating frequency, ω_{sv}.

$$\begin{aligned} \tau &= \frac{V_{\max}^2 \omega_{dv}}{\omega_{sv}^2 R_r} \\ 0 &= \tau R_r \omega_{sv}^2 - V_{\max}^2 \omega sv + V_{\max}^2 \omega_r \\ \omega_{sv} &= \frac{V_{\max}^2 - \sqrt{V_{\max}^4 - 4\tau R_r V_{\max}^2 \omega_r}}{2\tau R_r} \end{aligned} \tag{32.22}$$

Solving (32.22) for $\tau = 1$ for example gives $\omega_{sv} = 3.1$. We pick the smaller of the two solutions to the quadratic equation because it results in less slip. The other solution represents the point where the constant-torque curve rises above $V_{\max}$ at very high frequency. This point is almost visible only for the $\tau = 8$ curve at $\omega_s = 7.5$.

In practice the controller finds both voltage-constrained and current-constrained operating points and and selects the point with the lower V_s and the higher ω_s as being the most constrained.

Figures 32.7 and 32.8 also show curves of constant negative torque — as dashed lines. When braking (producing negative torque) the controller computes operating points in the same manner as for positive torque. Current-constrained and voltage constrained operating points are calculated using (**??**)

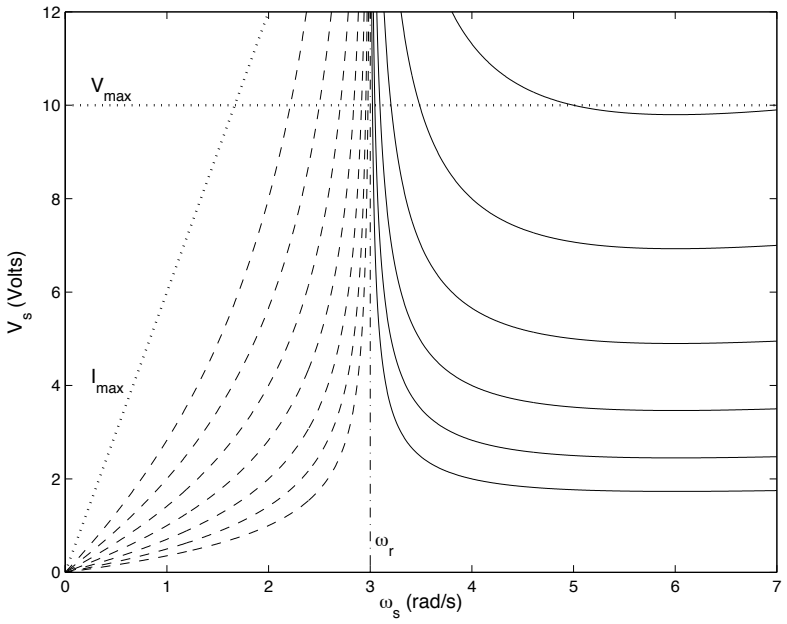

Figure 32.8: Lines of constant torque in the ω_s,V_s plane for $\omega_r = 3$. At this higher speed the controller is entirely constrained by $V_{\max}$ and the torque value of +16 is not achievable.

and (32.22) with negative values of torque. For negative torque the constrained operating point with the lower ω_s is chosen as being the most constrained.

32.6 Exercises

Part Part IV - Control

Revisiting Buck control
PD controller
Current-mode control
Energy-based control
Charge based control
A hybrid controller
Bode plots
Phase margin and stability
Integral windup
Predictive control
Model-based control
Sensing and estimation

Part V - Practical Design

An example buck converter
Power factor correction
An inverter

Part VI - Efficient Design

Losses
Conduction losses
Switching losses
Gate losses
Leakage inductance losses (in flyback)
Soft switching
Quasi-square wave
Active clamp
Isolated sensing

Chapter 33

Soft Switching

The switching losses of a converter can be largely eliminated by changing the state of the switching elements (typically FETs or IGBTs) only when they have either zero voltage across them or zero current through them. This is called zero voltage switching (ZVS) or zero current switching (ZCS). Collectively the two techniques are called *soft switching*. In contrast the conventional converters described in Chapters 2 to **??** are *hard-switching* converters because they change the state of switches with both non-zero current and voltage — leading to switching losses.

In this chapter we introduce soft switching by examining a quasi-square-wave (QSW) buck converter, one of the simplest ZVS converters. This converter is just a conventional buck converter with one capacitor added to slow the voltage transient. The control of the QSW converter, however, differs substantially from the control of a conventional buck and requires a gate drive with variable frequency.

In subsequent chapters we will explore quasi-resonant converters and converters employing an active clamp.

33.1 The Quasi-Square-Wave Converter

Figure 33.1 shows a schematic diagram of a QSW buck converter. The circuit is identical to the buck converter of Figure 2.3 except for the addition of capacitor C_X to slow the voltage transitions on node X, the drain of M_2.

The main difference between the QSW converter and the hard-switched buck is not the topology but rather how it is controlled and the very high inductor ripple currents that result. To turn M_1 and M_2 on ZVS, the QSW converter uses current in the inductor to charge and discharge C_X so there is no voltage across each FET when it is switched on. To charge C_X to V_S before turning on M_1 requires the current in L to be negative. To keep the average inductor current equal to the load current the peak current must be more than twice the load current.

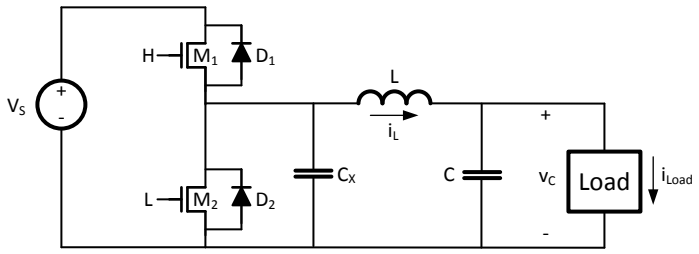

Figure 33.1: A quasi-square-wave (QSW) buck converter adds capacitor C_X to a conventional buck converter to slow the voltage swing on the drain of M_2 to facilitate ZVS of both switches.

Parameter	Value	Units	Description
C	100	μF	filter capacitor
L	10	μH	inductor
C_X	47	nF	drain capacitor
V_S	100	V	supply voltage
V_C	50	V	output voltage
I_L	5	A	load current

Table 33.1: Parameters and component values for the QSW buck simulation of Figure 33.2.

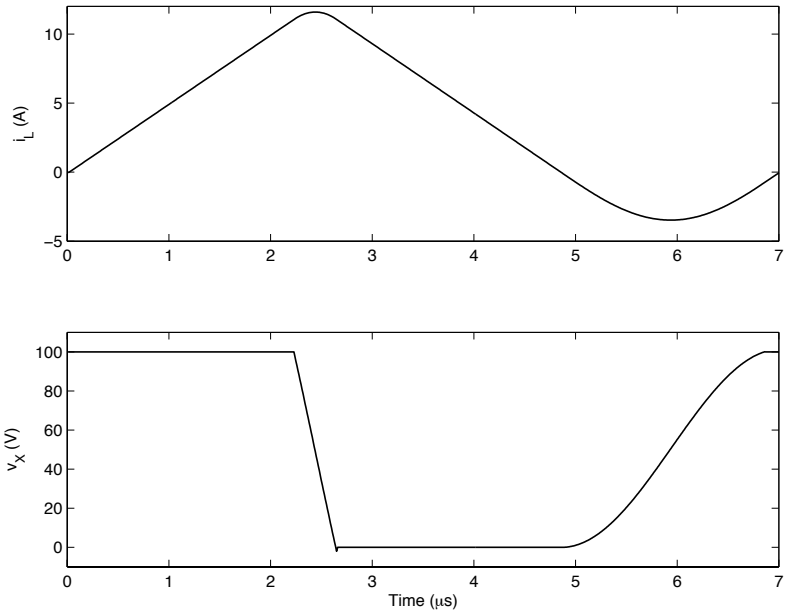

Figure 33.2: Waveforms showing operation of the QSW buck converter of Figure 33.1 with the parameters of Table 33.1. The inductor current (top panel) goes negative to store energy to charge C_X. The voltage on C_X (bottom panel) consists of segments of sine waves during transients.

Figure 33.2 shows inductor current and voltage across capacitor C_X for one cycle of operation for the QSW converter of Figure 33.1 with the parameters shown in Table 33.1. We divide the cycle into four periods that are delimited by switching either M_1 or M_2 on or off.

Initially M_1 is on and the inductor current ramps from 0 to a threshold value of 11A. When this threshold i_{high} is reached, M_1 turns off ZVS at 2.23μs. Capacitor C_X slows the fall of node X so that very little voltage develops across M_1 before its current has gone to zero. The fall time of node X is 420ns. From Figure 17.2 we saw a typical FET can be switched off in 10-20ns. Thus, at most 5V will develop across M_1 while its current ramps down.

During the 420ns fall time of node X both switches are off. Capacitor C_X is discharged by the inductor current which peaks at 11.5A giving a slew rate of 24V/μs. This slew rate is determined by the value of C_X which is chosen large enough to make the loss of turning M_1 off negligible. During the fall of node X both inductor current and node voltage vary sinusoidally[1] as the tank

[1]The node voltage looks like a striaght line but it is in fact part of a sine wave.

circuit formed by L and C_X oscillates. The voltage passes through the peak of the sine wave while the current passes through the portion of the sine wave with the steepest descent.

When node X falls through 0 the body diode of M_2 turns on clamping V_X a diode drop below ground. At this point M_2 is turned on ZVS. With M_2 on and node X at 0 V current in the inductor ramps down from 11A at 2.65μs to a threshold value of $i_{\text{low}} = 0$A at 4.86μs. When i_{low} is reached M_2 is turned off ZVS. As with the turn-off of M_1, capacitor C_X keeps the voltage across M_2 near zero while its current ramps down. In this case, because the inductor current is so low, the voltage change during hte turn-off of M_2 is negligible.

With both switches off again, the tank circuit formed by L and C_X resumes oscillation — but this time with the inductor feeding current into node X causing it to rise. During this period both waveforms are again sinusoidal with the current following the bottom of the sine wave and the voltage the rising segment. With a much lower starting current 1A vs 11A the rise time of node X is much slower than the fall time — taking 1.7μs to slew from 0 to V_S. When v_X reachs v_S the body diode of M_1 clamps it at this value. At this point M_2 switches on ZVS and the cycle has completed.

During the fall time from $2.23 - 2.65\mu$s and the rise time from $4.86 - 6.86\mu$s both switches are off and the tank circuit formed by C_X and L oscillates. The frequency of this oscillation is $f_X = \frac{1}{2\pi\sqrt{LC_X}} = 232kHz$. This frequency is unrelated to frequency of the switching cycle, which in this case is 148kHz. To make the figures easier to understand we have set f_X unusually low. In a typical converter it would be many times this value.

For a fixed value of C_X, this rise time could be made faster by setting a lower (more negative) value of i_{low}. However, we typically want to keep i_{low} to the highest value that provides sufficient current to drive node X to V_S. This value is given by:

$$i_{\text{low}} = -\left(V_S - 2V_C\right)\sqrt{\frac{C_X}{L}} \tag{33.1}$$

Setting i_{low} below this point increases current ripple which in turn leads to increased conduction and core losses.

If i_{low} is a little too high the inductor current will cross zero before v_X reaches v_S. This case is illustrated in Figure 33.3. Here we have set v_C to 45V, so starting from zero inductor current the LC tank rings v_X up to 90V (twice the initial voltage). The FET switches on when i_L passes through zero and pulls v_X up the remaining 10V. While this is not exactly ZVS operation, switching with 10V across the terminals of M_1 requires much less energy than switching with 100V across the terminals. We can avoid this issue altogether by setting i_{low} sufficiently low.

The frequency of a QSW converter varies with average current. The lower the current the higher the switching frequency. Figure 33.4 illustrates converter operation with $\hat{i} = 0.7$A. The cycle time here is just 4.7μs compared to 6.9μs in Figure 33.2.

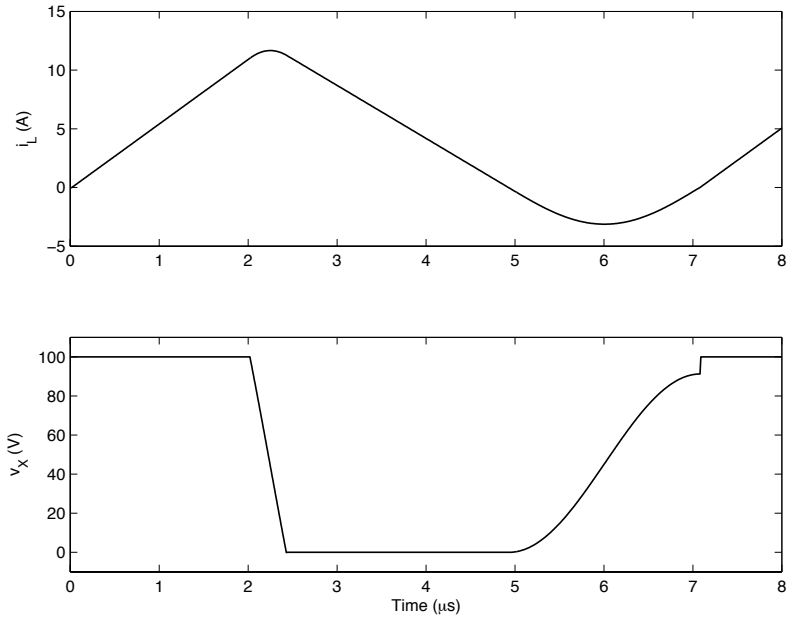

Figure 33.3: Waveforms showing operation of the QSW buck converter of Figure 33.1 with the parameters of Table 33.1 but with $v_C = 45$V. In this case i_L crosses zero before v_X reaches v_S. We turn on M_1 at this point and it pulls v_X up the last 10V.

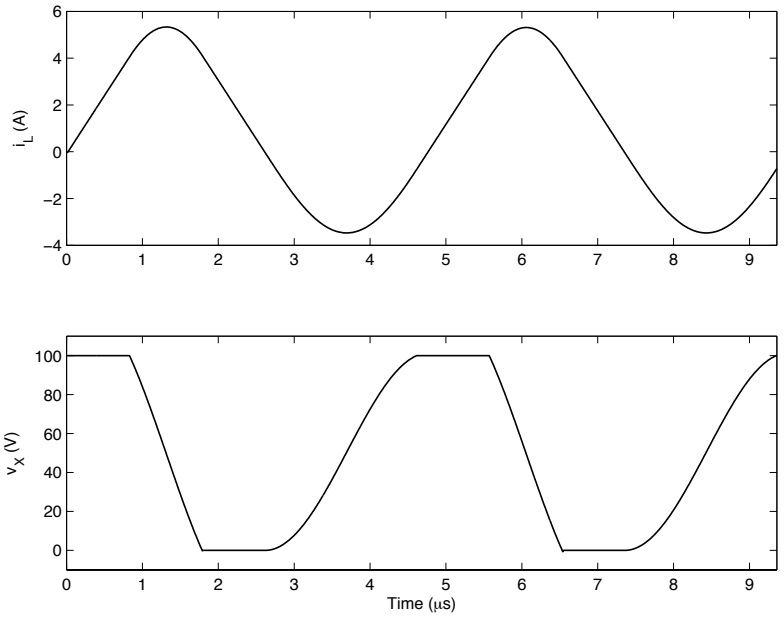

Figure 33.4: Waveforms showing operation of the QSW buck converter of Figure 33.1 with the parameters of Table 33.1 but with $I_L = 0.7$A. At this low operating current the cycle time is reduced to 4.7μs.

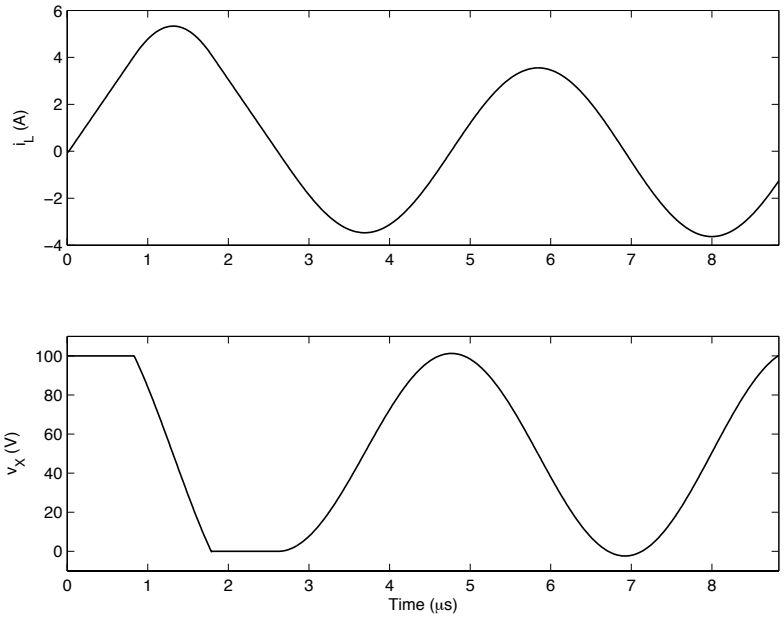

Figure 33.5: Waveforms showing operation of the QSW buck converter of Figure 33.1 with the same parameters as in Figure 33.4. To increase the cycle time the controller skips an entire cycle of the LC tank before turning on M_1 stretching the cycle time to 9.1μs.

To avoid excessive switching at low currents the QSW converter can skip one or more cycles of the LC tank before turning on M_1. Skipping one cycle is illustrated in Figure 33.5. When M_2 turns off at 2.6μs the LC tank is allowed to oscillate not for a quarter cycle (as in Figure 33.2-Figure 33.4) but for 1.25 cycles. The oscillation is halted when M_1 turns on the second time v_X reaches v_S.

33.2 Control of the Quasi-Square-Wave Converter

As described above, the four events in the switching cycle are triggered by i_l and v_X reaching thresholds. The *off* events are triggered by current thresholds and the *on* events are triggered by voltage thresholds.

M_1 **Off:** When $i_L > i_{\text{high}}$.

M_2 **On:** When $v_X < 0$.

M_2 **Off:** When $i_L < i_{\text{low}}$.

M_1 **On:** When $v_X > V_S$ or $i_L > 0$.

We add the *or* condition to the M_1 on trigger to handle cases where the inductor runs out of current before v_X reaches v_S.

From these trigger conditions we see that the control variables are the two current thresholds, and if keep i_{low} near its maximum value to minimize losses, we have a single control variable i_{high}. Thus we control a QSW converter by setting threshold i_{high}. This in turn determines the average current over the cycle $\hat{i}$ which integrates on C to determine the output voltage.

The transfer function of the system is:

$$\frac{v_c(s)}{i_{\text{high}}(s)} = \frac{K}{Cs} \tag{33.2}$$

where K is the gain constant relating i_{high} to $\hat{i}$.

Applying a PI controller with transfer function $P + Q/s$ gives open- and closed-loop transfer functions:

$$G(s) = \frac{KPs + KQ}{Cs^2} \tag{33.3}$$

$$H(s) = \frac{\frac{KP}{C}s + \frac{KQ}{C}}{s^2 + \frac{KP}{C}s + \frac{KQ}{C}} \tag{33.4}$$

So we have

$$\omega = \sqrt{\frac{KQ}{C}} \tag{33.5}$$

$$\zeta = \frac{KP}{2\omega C} \tag{33.6}$$

$$Q = \frac{\omega^2 C}{K} \tag{33.7}$$

$$P = \frac{2C\zeta\omega}{K} \tag{33.8}$$

Figure 33.6 shows a simulation of a QSW converter using a PI controller with ω chosen to be 6.28×10^3. The waveforms show the response of the controller to two shifts in target voltage. The controller goes non-linear during both transients with i_{high} (middle panel) being clamped to its low limit of 0 A during the downward transient and to 20A during the upward transient. There is some integral *wind-up* in the controller during these periods where the output is saturated.

The waveforms show how cycle time varies with current. As i_{high} varies from 0 to 20A, t_{cy} varies by more than 2:1 from a minimum of 4.3μs when the current is 0A to 10.5μs when i_{high} is set to 20A.

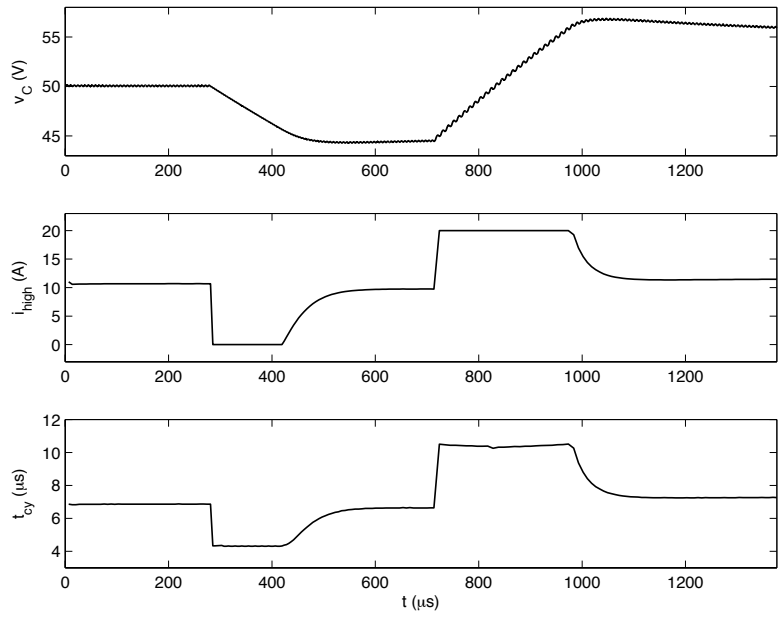

Figure 33.6: Simulated waveforms of a QSW buck converter with a PI controller. The target voltage is changed to 45V at 274μs and to 55V at 707μs.

33.3 Exercises

evaluate switching loss
- play with values of C_X
- design and simulate a QSW boost converter
- explore effects on comonent values and ratings