INTRODUCTION TO **MECHATRONIC DESIGN**

CARRYER
OHLINE
KENNY

Samsung et al. v. Resonant Sys.
IPR2023-00993 - Exhibit 1051

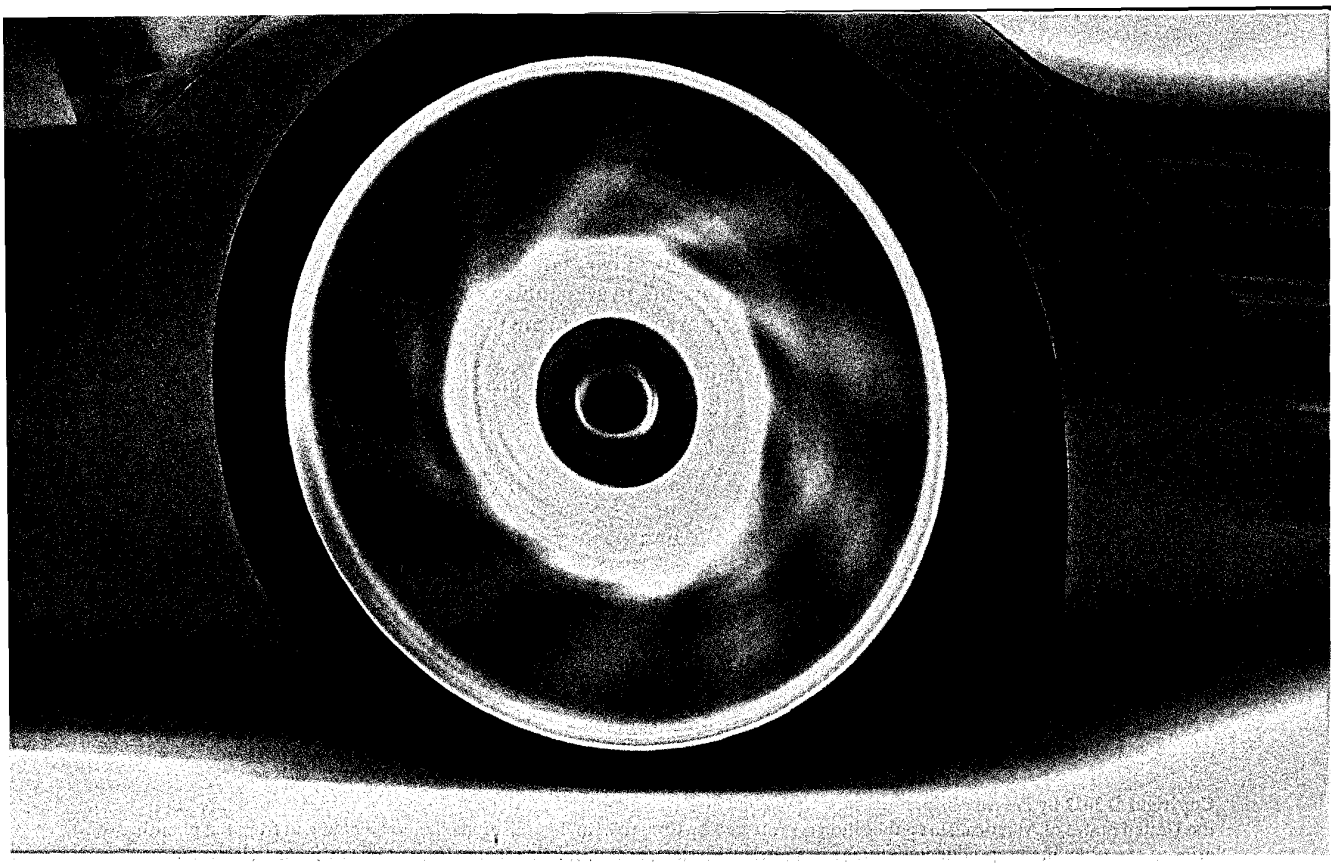

INTRODUCTION TO MECHATRONIC DESIGN

J. EDWARD CARRYER
R. MATTHEW OHLINE
THOMAS W. KENNY

Mechanical Engineering
Stanford University

Prentice Hall

Boston Columbus Indianapolis New York San Francisco Upper Saddle River
Amsterdam Cape Town Dubai London Madrid Milan Munich Paris Montreal Toronto
Delhi Mexico City Sao Paulo Sydney Hong Kong Seoul Singapore Taipei Tokyo

Vice President and Editorial Director, ECS: *Marcia J. Horton*
Senior Editor: *Tacy Quinn*
Acquisitions Editor: *Norrin Dias*
Editorial Assistant: *Coleen McDonald*
Vice President, Production: *Vince O'Brien*
Senior Managing Editor: *Scott Disanno*
Production Liaison: *Jane Bonnell*
Production Editor: *Pavithra Jayapaul, TexTech International*
Senior Operations Supervisor: *Alan Fischer*
Operations Specialist: *Lisa McDowell*
Executive Marketing Manager: *Tim Galligan*
Marketing Assistant: *Mack Patterson*
Senior Art Director and Cover Designer: *Kenny Beck*
Cover Image *Somatuscan/Shutterstock*
Art Editor: *Greg Dulles*
Media Editor: *Daniel Sandin*
Composition/Full-Service Project Management: *TexTech International*

Company and product names mentioned herein are the trademarks or registered trademarks of their respective owners. See p. xxiii for additional trademark acknowledgments.

Copyright © 2011 by Pearson Education, Inc., Upper Saddle River, New Jersey 07458. All rights reserved. Manufactured in the United States of America. This publication is protected by Copyright and permissions should be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. To obtain permission(s) to use materials from this work, please submit a written request to Pearson Higher Education, Permissions Department, 1 Lake Street, Upper Saddle River, NJ 07458.

The author and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, expressed or implied, with regard to these programs or the documentation contained in this book. The author and publisher shall not be liable in any event for incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of these programs.

Library of Congress Cataloging-in-Publication Data
Carryer, J. Edward.
Introduction to mechatronic design / J. Edward Carryer, R. Matthew Ohline, Thomas W. Kenny.
p. cm.
Includes bibliographical references and index.
ISBN-13: 978-0-13-143356-4 (alk. paper)
ISBN-10: 0-13-143356-3
1. Mechatronics. I. Ohline, R. Matthew. II. Kenny, Thomas William. III. Title.
TJ163.12.C37 2010
621—dc22 2010033713

Prentice Hall
is an imprint of

PEARSON

www.pearsonhighered.com

13 12 11 10 9 8

ISBN-13: 978-0-13-143356-4
ISBN-10: 0-13-143356-3

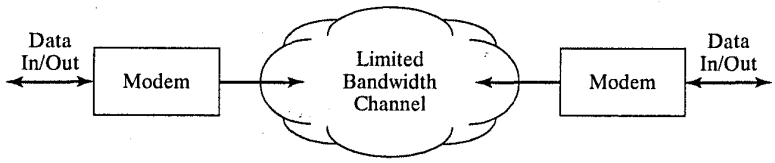

FIGURE 7.16
Communicating over a limited bandwidth channel.

these channels. Consider the case of a long idle line that might exist between transmissions on two devices using asynchronous communications. That long idle time will represent a very low frequency signal, in the limit approaching DC. This will not pass over any communications channel whose response does not extend down to DC. In order to communicate over one of these limited bandwidth channels, we will need to modify our signal so that it always fits within the bandwidth provided by the communications channel.

The way that we modify our signal is to interpose a pair of new devices between the two devices that wish to communicate (Figure 7.16).

The **modem** (MOdulator DE-Modulator) is responsible for translating our logical signal into something that fits in the frequency range provided by the channel. This translation can take place using any of a number of different modulation techniques, which we discuss in the following sections.

7.5.1.1 Modulation Techniques

The modulation techniques used over limited bandwidth channels, including both RF and the telephone system, fall mostly into a group of techniques involving amplitude modulation (AM), frequency modulation (FM), phase modulation (PM), or some combination of these. All of these techniques involve modulating (modifying) a carrier or base-band signal in order to encode digital information onto that signal. The basic operation of RF modems mimics that of the telephone modem except that the carrier frequencies are much higher.

7.5.1.2 Amplitude Modulation

In **amplitude modulation**, or **AM**, the carrier is turned on and off to indicate the changes in logical state (Figure 7.17).

As a result, this type of modulation is also referred to as **on-off-shift keying (OOSK)**, **amplitude shift keying (ASK)**, or **on-off keying (OOK)**. This simple approach to modulation is not actually employed by modems for use over the telephone system, but finds application in simple RF communications links. It serves as an excellent and easy-to-grasp starting point for a discussion on signal modulation. In an ASK system, if the carrier signal is present, it represents a "1" in the data stream. The absence of the carrier represents a "0." Asynchronous

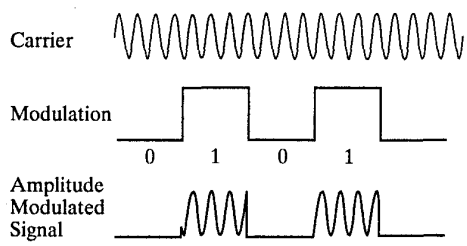

FIGURE 7.17
Amplitude modulation (AM).

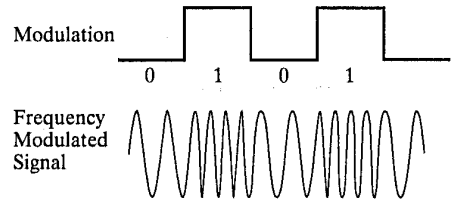

FIGURE 7.18

Frequency modulation (FM), also known as frequency shift keying (FSK).

communications is used with this and all of the other modulation techniques that we discuss, so both sides of the conversation must have agreed on what a bit time will be.

7.5.1.3 Frequency Modulation

When we modify the frequency of the carrier signal rather than its amplitude, it is called **frequency modulation** or **FM**[2]. This is also referred to as **frequency shift keying (FSK)**, and is shown in Figure 7.18.

FSK is the modulation technique used in the simplest of telephone modems. One early standard governing the function of such systems is the Bell 103 standard (released by AT&T in 1962). Under this standard, one device represents logical 1s with a 2,225 Hz signal and logical 0s with a 2,025 Hz signal. The other device in the pair represents logical 1s with a 1,270 Hz signal and logical 0s with a 1,070 Hz signal. In this way, both devices can transmit at the same time (full-duplex) and filters can be used to separate out the different frequency bands corresponding to the two devices. However, the maximum data rates allowed with this approach are quite low, since a few cycles are needed to determine the frequency present at any given time. In practice, devices employing the Bell 103 standard achieved data rates of 300 bps.

7.5.1.4 Phase Modulation

If we modulate the carrier by creating sudden phase shifts, it is referred to as **phase modulation (PM)** or **phase shift keying (PSK)**, and is depicted in Figure 7.19.

Because it is relatively easy (though outside the scope of this textbook) to detect phase shifts, even over a telephone system, it is possible to have more than two possible phases involved in the modulation. Figure 7.20 shows a system in which a pair of bits (a **dibit**) is represented by the relative shift in phase between two adjacent cells.

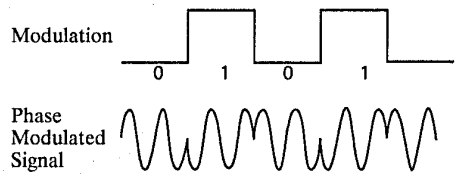

FIGURE 7.19

Phase modulation (PM) or phase shift keying (PSK).

[2]This use of FM is the same as that used in FM radio. The principal difference being that in FM radio an analog, rather than digital, voltage is being used to control the modulation. FM modems for use with phone lines operate on audio signals in the range of a few kHz, while FM radio operates with RF signals in the range of 88–108 MHz.

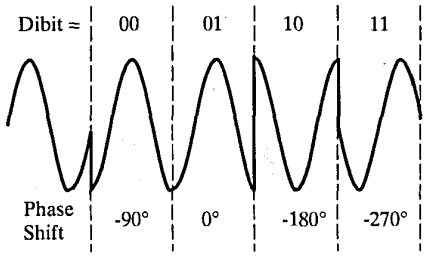

FIGURE 7.20
PM encoding 2 bits per cell.

7.5.1.5 Quadrature Amplitude Modulation

We can combine PSK with a variation on AM to encode even more information into a signal. While this approach is mathematically equivalent to modulating both phase and amplitude, the scheme is implemented by modulating only the amplitude of two carriers that are 90° out of phase (in quadrature) and then summing the two carriers together. It goes by the name of **quadrature amplitude modulation** or **QAM.** We represent this system with a plot that shows amplitude of the two carriers (I and Q) on the x- and y-axes (Figure 7.21). Another way to view this presentation is as a polar plot with the length of the vector representing its amplitude and the angle with respect to the positive x-axis, the phase.

For the increased complexity that this approach requires, it returns benefits in increased encoding efficiency (more bits transferred in a given period of time). The system of Figure 7.21 uses two amplitude levels and eight phases to give 16 distinct phase-amplitude combinations. This allows us to encode 4 bits of information into a single tone. This approach can be extended to larger numbers of phase-amplitude combinations. Telephone modems have been built with up to 512 combinations of amplitude and phase allowing 9 bits of information to be transmitted by a single tone with a unique amplitude and phase.

7.6 COMMUNICATING WITH LIGHT

Every time you use the remote control on your TV or DVD player, you are using a form of communications that employs IR light. The range of light wavelengths in the IR spectrum is

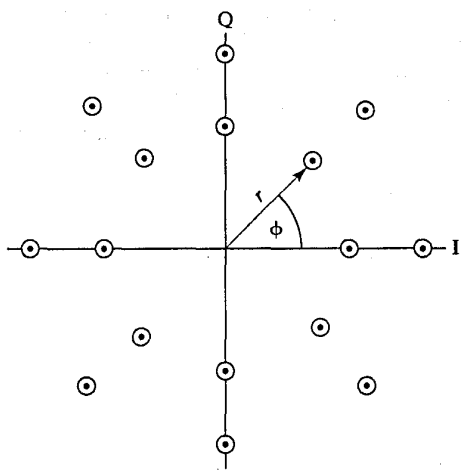

FIGURE 7.21
An example of quadrature amplitude modulation (QAM).

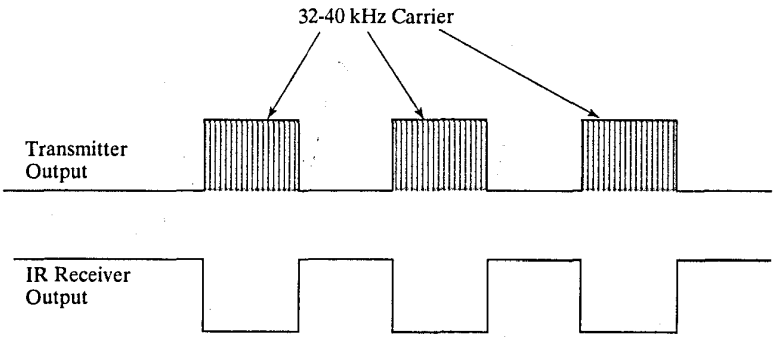

FIGURE 7.22

IR remote control modulation of the carrier.

very large, spanning from 750 to 1 mm, though most IR communications devices operate in the region between 850 and 950 nm. The remote control encodes your command as variations of the IR light level produced by an IR light emitting diode (LED).

It is not easy to modulate the frequency of the light directly, so IR remote controls generate a carrier by turning the IR LED on and off at the carrier frequency and then modulate the carrier. Most remote controls use a carrier frequency in the range of 32–40 kHz. The IR receivers used by the appliances are designed to interpret the presence of the carrier as a logical 0 and the absence of the carrier as a logical 1, see Figure 7.22.

While it might be possible to transmit an asynchronous data stream using this approach, the realities of the response of the IR receivers make this difficult. There are time shifts between the appearance of the IR signal and when the presence of the carrier is detected. These delays vary to some degree depending on ambient light levels, making it difficult to maintain framing between the transmitter and receiver. As a result, most remote controls use OOSK to modulate the carrier, but use a different encoding technique to encode the 1s and 0s. There are a number of different approaches that have been developed but all rely on incorporating timing information into every message and sending discrete messages rather than arbitrary strings of data.

The timing information is included by adding a **preamble** to every message. This preamble will consist of a predefined pattern of pulses of the carrier. The receiver uses the preamble and its knowledge of what the pattern should be to determine the bit time associated with this particular message. In this way, the receiver adapts to the transmitting rate for every message and the transmit timing only needs to be stable for a very short period of time (usually just a small fraction of a second). Figure 7.23 shows a preamble consisting of three equal pulses followed by a pulse that is twice as long that indicates the end of the preamble and the beginning of data.

To increase the reliability of the transmission, most IR remote controls do not use the simple presence or absence of a carrier to indicate logical 1s and 0s. Instead, they encode the logical data in the timing of the period that the carrier is present or absent. There

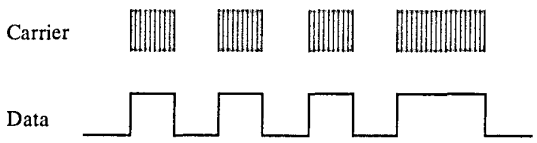

FIGURE 7.23

A preamble to an IR message.

23.3 BIDIRECTIONAL CONTROL OF MOTORS

Any of the circuits discussed above that include kickback suppression are good choices for unidirectional control of DC motors, but what about applications that will require a motor to spin in both directions (clockwise—CW, and counterclockwise—CCW)?

For the circuits shown in Figures 23.4, 23.6, 23.7, and 23.8, current can only flow from the positive terminal of the power source to ground, so the motor can only turn in one direction. Enabling a motor to turn either CW or CCW requires the capability to cause current to flow through the motor in either direction. A different, more complex drive circuit will be required to accomplish this. One possibility is to make use of a "dual" power supply, that is, one with both a positive voltage supply and a negative voltage supply. By switching one on or the other, current can flow through the motor in either direction. Figure 23.10 illustrates this configuration. Note that kickback diode protection must be included, preventing inductive kickback from exceeding the maximum collector-emitter voltage for each of the transistors.

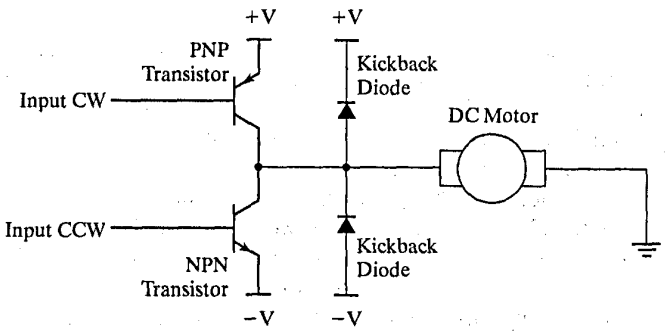

FIGURE 23.10
A circuit for bidirectional current flow through a DC motor using a dual power supply.

For this circuit, if the voltage at the Input CW node is at least 0.6 V below the positive voltage supply (+V) and the Input CCW node is held at the negative supply voltage (−V), the PNP transistor will be on and conducting, and the NPN transistor will be off, resulting in current flowing from the +V terminal through the motor to ground. This will cause the motor to spin in one direction (our labeling convention indicates that it will spin CW).

We can reverse the direction of current flow through the motor by changing which transistor is on. If Input CW is held at +V, and Input CCW has a value of at least 0.6 V above −V, the PNP transistor will be off and the NPN transistor will be on and conducting. Under these conditions, current will flow from ground (0 V) to the negative supply terminal (−V), which has a voltage potential lower than that of ground, and the motor will spin in the CCW direction.

Turning on both transistors at the same time must be avoided in this circuit, as it would result in a short circuit from +V to −V through the transistors. This would likely destroy the transistors and/or the power supply.

For this example, current flows from left to right across the motor if the upper transistor is on, and from right to left across the motor if the lower transistor is on. This is one way to accomplish bidirectional current flow, but in practice, it is not common to have a high current negative voltage supply available. It would be more practical to design a circuit capable of causing current to flow in either direction through a motor using only a single positive power supply. The most common way of accomplishing this is through the use of a network of transistors, arranged in what is known as an **H-bridge** configuration. The shape of the circuit resembles an "H," which gives rise to the distinctive name. Figure 23.11 illustrates an H-bridge.

Recall that the objective of this circuit is to enable current to flow through the motor in either direction (left-to-right or right-to-left) using only a single positive power supply

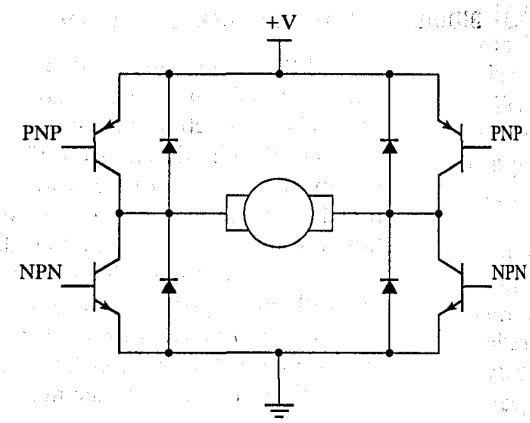

FIGURE 23.11

A typical H-bridge circuit, including kickback diodes to protect against inductive kickback.

(shown in Figure 23.11 as +V). The diodes that surround the motor are in place to protect the transistors from the inductive kickback, as described in Section 23.2. Since current may flow in either direction across the motor, the high voltage inductive kickback can occur on either motor terminal and must be snubbed to prevent the destruction of the transistors. Though any of the snubbing techniques discussed above are applicable with H-bridges, by far the most common practice for brushed DC motors is the use of standard diodes.

Examining Figure 23.12a first, we see that we have established the situation where the left PNP transistor is on, the right PNP transistor is off, the left NPN transistor is off, and the right NPN transistor is on. For these conditions, current can flow from the positive voltage terminal of the power supply, through the left PNP transistor, through the motor from left to right, and finally through the right NPN transistor to ground.

In Figure 23.12b, the opposite conditions have been established: by reversing the state of each of the transistors, current can now flow from the positive terminal of the power supply, through the PNP transistor on the right side, through the motor (this time from right to left), and finally through the left NPN transistor to ground.

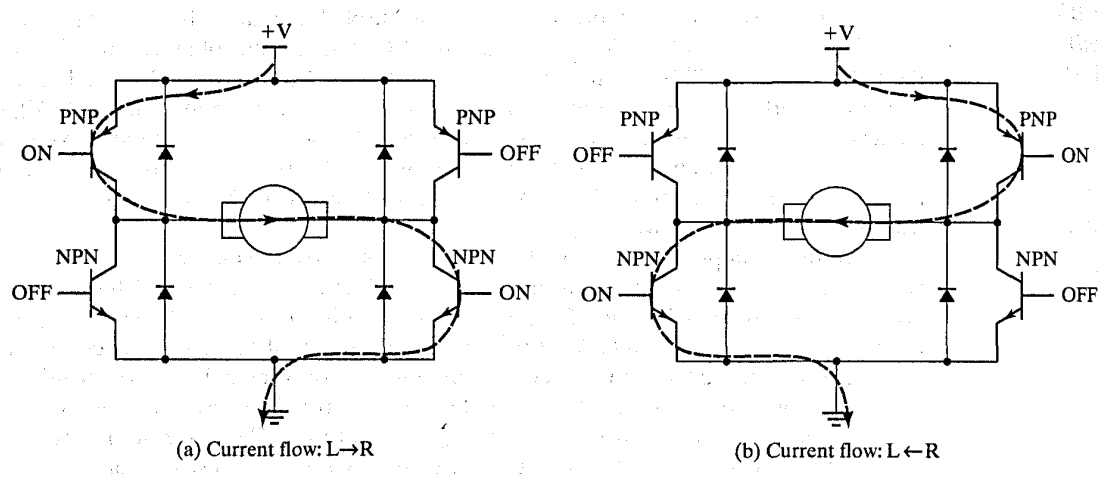

FIGURE 23.12

Using an H-bridge to control the direction of current flow in a motor.

Other possibilities exist for combinations of transistors states, of course. For instance, if both transistors on either the left side or the right side of the "H" are switched on, current will flow without any resistance, directly from power to ground, Figure 23.13. In this case, a short circuit is effectively established between power and ground, since there is no load or other component in the path to resist the flow of current. This results in very high levels of current flow. If it is sustained for any significant length of time, one of the components involved is very likely to be destroyed by the excessive current levels. This unrestrained current is referred to as **shoot-through current.** Shoot-through current can occur under normal circumstances, such as while an H-bridge is in the process of switching from one desirable configuration to another (for instance, when reversing the direction of the current flow through a motor). This occurs because the transistors don't have symmetric turn-on and turn-off times (turn-off time usually being longer than turn-on time). Under these conditions, brief bursts of shoot-through current may be experienced. Shoot-through current happens, but it is NOT desirable. The designers of H-bridges take great care to ensure that shoot-through currents are minimized, often including intentional delays, called **dead time,** into the switching of the transistors to minimize the overlap time, when the transistors on one side of the H-bridge are both conducting, that results in shoot-through. The need to control shoot-through currents is one of the significant issues that makes the design of successful H-bridges a nontrivial task.

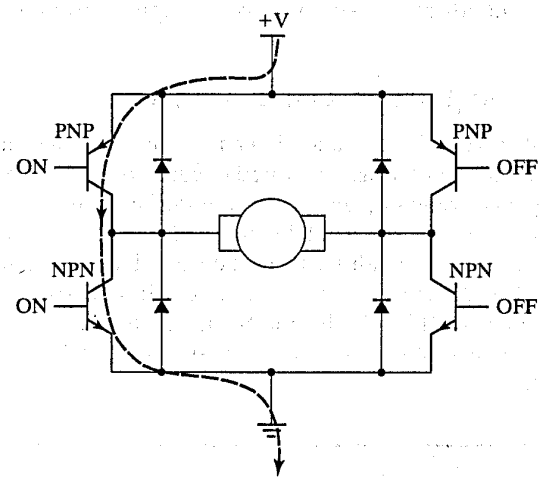

FIGURE 23.13

Shoot-through current in an H-bridge, where current flows unimpeded from power to ground, usually with disastrous effects.

Another, more useful, combination of transistors to turn on configures an H-bridge to dynamically brake, or resist the rotation of, a motor. When the leads of a spinning motor are shorted together, the current that is generated by the back-EMF (recall that the spinning motor acts as a generator) causes current to flow in a loop, as shown in Figure 23.14. The direction of current flow during **dynamic braking** depends on the direction of the motor's rotation. Current will flow from one of the motor's terminals, through one of the kickback diodes, around to one of the transistors that are in the "on" state, and finally around to the motor's other terminal. The internal resistance of the motor coils, the forward voltage drop across the diode, and the saturation voltage of the transistors serve as the load for this current, and the electrical power being generated by the rotating motor will be dissipated as heat. In the absence of a load caused by dynamic braking, the only forces acting to slow the rotation of a motor are the frictional resistance of the motor bearings and other mechanical systems coupled to the motor. Dynamic braking can be performed with an H-bridge

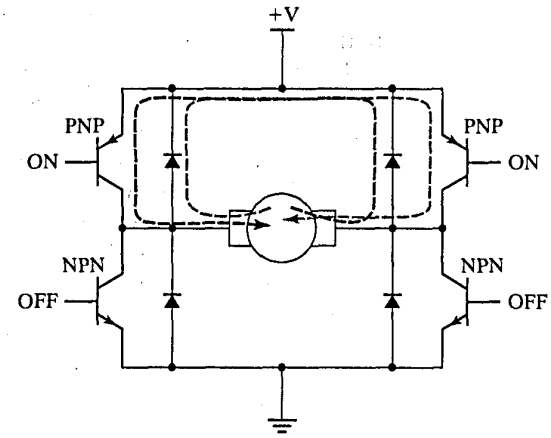

FIGURE 23.14

An H-bridge configured to short both leads of motor together, creating dynamic braking. Depending on the direction of rotation of the motor, the current flow will be either CW or CCW, as shown.

by turning on both the upper transistors while turning off both the lower transistors, or by turning off both the upper transistors while turning on both the lower transistors. This effectively shorts both leads of the motor together and results in dynamic braking. This is useful when switching a motor off in circumstances where it is desirable for the motor to stop rotating more quickly than would otherwise occur if it were simply allowed to slowly decelerate and coast to a stop.

23.3.1 Commercially Available H-Bridge Integrated Circuits

Bidirectional switching of current through a load, such as a motor, is such a common need in electronic design that there are numerous integrated circuits (ICs) in the market that implement H-bridges in a single convenient package, so that designers don't have to create one from discrete components. H-bridges are available in a wide variety of current drive capabilities, package sizes, and features. They all perform the same basic task, however: switching current through a load in two directions. Some integrated solutions you may wish to investigate are the Texas Instruments SN754410NE, the ST L293, the ST L298, and the National Semiconductor LMD18200, several of which are discussed in detail in Chapter 17: Digital Outputs and Power Drivers.

Example 23.1

Design drive circuitry based on an L293B H-bridge IC (manufactured by STMicroelectronics) that is capable of controlling a Maxon RE-13 (Part #: 118638) motor using a 12 V power supply. The motor has terminal resistance $R = 14.1\ \Omega$ and coil inductance $L = 0.48$ mH. It will be switched on and off at a maximum rate of 1 kHz. To protect the circuit from excessive inductive kickback voltages, use standard diodes. Ensure that the specifications for all components are met by the design.

In order to satisfy the constraints, the L293B must be capable of switching the maximum current possible for this motor powered at 12 V. Since the terminal resistance is given, we can solve for the stall current ($\omega = 0$):

$$I_{\text{stall}} = \frac{V}{R} = \frac{12\ \text{V}}{14.1\ \Omega} = 0.851\ \text{A}$$

The specifications for the L293B show that it is capable of handling up to 1 A per channel (continuous). Since the stall current is the maximum current the system will experience, the L293B will be able to safely handle controlling this motor.

The next issue to address is the inductive kickback. The problem specifies the use of standard diodes, so we must identify a candidate that has a suitable reverse recovery time specification and appropriate current capabilities. A diode that is very commonly employed as a kickback diode is the 1N4935, which is identified as a "fast recovery diode" in the list of features found in its data sheets. Recall that kickback diodes should have a reverse recovery time specification of a few hundred nanoseconds or less, and the 1N4935 has $t_{rr} < 200$ ns. The requirement for low reverse recovery time has been met. Recall also that good design practice dictates that the diode's specification for peak intermittent current should be above the maximum current that the circuit will encounter, which is the stall current calculated above (0.851 A). The peak intermittent current specification for the 1N4935 is 10 A, which is more than 10 times the stall current expected for this circuit. The 1N4935 is therefore a good choice in this respect. The issue at the heart of the specifications for peak (and continuous) current is power dissipation: the energy stored in the motor's coil is dissipated during kickback by the standard diode and the motor's own coil resistance. We can analyze the power dissipated by the snubbing circuit from Eq. (23.8), where we find:

$$P = \left(\frac{1}{2} I_{\text{stall}}^2 L\right) \times \text{Frequency} = \frac{1}{2}(0.851\text{ A})^2(0.48\text{ mH})(1\text{ kHz}) = 0.174\text{ W}$$

The total power dissipated during kickback is thus 0.174 W, which will be distributed between the standard diode and the resistance of the motor's coils. A conservative approach is to assume that the power dissipated in the diodes dominate the power dissipated by the motor's coils. The most conservative assumption is that *all* the power is dissipated in the diode; using this reasoning, any standard diode capable of dissipating at least $P_D = 0.174$ W will suffice. The data sheet for the 1N4935 doesn't directly state a maximum power dissipation specification; however, we know that its maximum continuous forward current is 1 A, and the maximum forward voltage drop at 1 A is 1.2 V. From this, we can infer that the maximum continuous power dissipation capability of the 1N4935 is $P_{D\,max} = VI = 1.2\text{ V} \times 1\text{ A} = 1.2$ W. This is much greater than the 0.174 W it will dissipate in this circuit, so we've verified that the 1N4935 is an acceptable choice. Thermal analysis of the 1N4935 under these conditions also shows that the junction temperature remains well below the maximum allowable levels, further supporting the choice of the 1N4935.

At this point, all of the components have been selected and the schematic may be drawn (Figure 23.15).

The inputs "Half1" and "Half2" each control one side, or leg, of the H-bridge. Each side of the H-bridge is called a **half-bridge**. A logic-level high to an input (such as the one labeled "IN1") connects the corresponding output (OUT1) to the motor supply voltage, minus a voltage drop due to

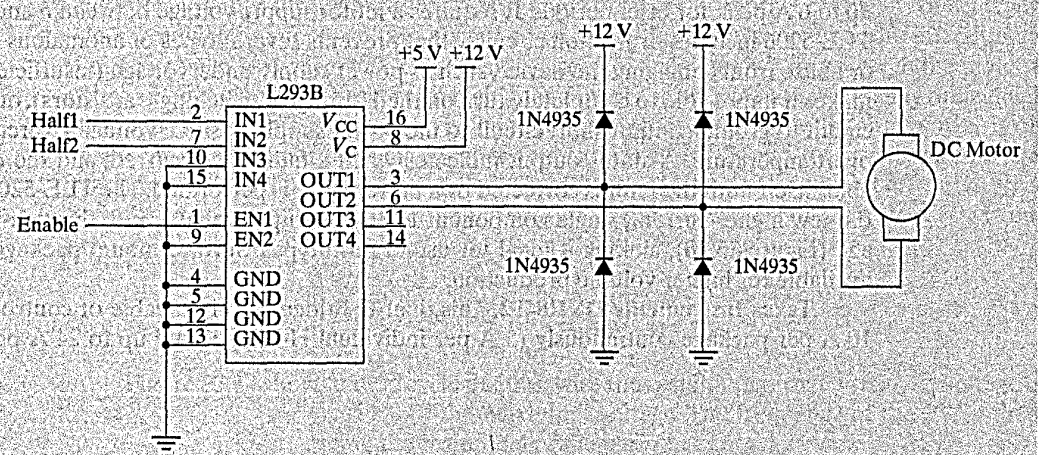

FIGURE 23.15

H-bridge circuit for Example 23.1.

losses (the saturation voltage). This drop is V_{CEsatH} (collector-emitter saturation voltage, high) and is specified to have a typical value of 1.4 V in the L293B data sheet. A logic-level low to an input connects the corresponding output to ground, again with losses due to the saturation voltage across the active switch. This is the low side saturation voltage (V_{CEsatL}) and is specified to be typically 1.2 V in the data sheet. Setting Half1 to a logic-level high and Half2 to a logic-level low will drive current through the motor in one direction and cause the motor to turn. Reversing the inputs drives the current in the opposite direction, so the motor also turns the other direction. Setting Half1 and Half2 both to logic-level low at the same time (or both to logic-level high) connects the terminals of the motor together and causes dynamic braking. The input called "Enable" controls current flow through the half-bridges. Setting the Enable pin low inhibits current flow and turns the motor off, allowing it to coast to a stop.

23.3.2 H-Bridges for Higher Current Applications

Many mechatronic devices are well-served by small, low-torque motors, and these often require only small amounts of current (e.g., 1 A or less). However, this doesn't apply in all situations: some applications call for substantial amounts of current to flow bidirectionally through a DC motor (or, more generally, any type of load). Circuits capable of managing such current levels are often little different than those designed for lower currents, but the components selected must be capable of managing the higher currents. There are a great many H-bridges available that are capable of handling between 1 and 2 A (some of which were enumerated in Section 23.3.1), and similarly there are many capable of managing 2 A up to about 10 A, as well as a category of H-bridge driver chips that allow for still greater current levels by moving the switching transistors off-board.

Once a designer determines how much current must be supplied to a motor (or other load) for a given application, the drive components for the motor must be capable of managing the required current levels. Identifying candidate components involves searching the product offerings of electronic component distributors and/or manufacturers, and then following their guidelines for the proper use of the devices. In the remainder of this section, we present a few representative devices.

Infineon TLE-5206: This device incorporates a single H-bridge driver and is thus able to drive a single motor bidirectionally. It is capable of delivering up to 5 A continuously, and up to 6 A peak for brief periods. It requires a motor supply voltage between 6 and 40 V. The TLE-5206 includes protection circuitry that prevents several types of anomalous conditions or abuse from damaging the device, such as power supply undervoltage (insufficient voltage makes it impossible to completely turn on the H-bridge's switching transistors), output short circuit to ground, output short circuit to the power supply positive voltage, overcurrent, and overtemperature. A status output indicates when a fault has occurred, and the device disables itself until the fault is corrected or a power cycle is performed. The TLE-5206 is one of the few higher-current, single component H-bridge chips available in a through-hole package (Figure 23.16), making it ideal for use in prototypes. Surface-mount packages are also available for higher-volume production.

Texas Instruments DVR8402: This dual H-bridge chip is capable of controlling up to 10 A per package continuously (5 A per individual H-bridge) and up to 24 A per package

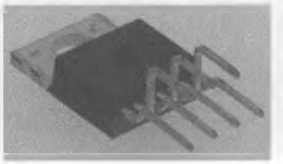

FIGURE 23.16
The Infineon TLE-5206, a 5 A H-bridge.

for brief intervals. The two H-bridges can also be wired together in parallel to provide a total of 10 A (continuous) to a single motor, or more generally, another type of load. Two separate power supplies are required by the DVR8402: a 0–50 V supply for the motor, and an 11.4–13.2 V supply for the internal logic circuits. Independent control of each of the four half-bridges on the chip is provided, enabling the use of dynamic braking. The internal MOSFETs used in the H-bridge circuits have a low $R_{DS(ON)}$ specification of 90 mΩ (recall that we discussed MOSFETs and the $R_{DS(ON)}$ specification in Chapter 10). This results in an impressively low 0.45 V drop across each transistor when the maximum 5 A is flowing. The DVR8402 provides protection against overtemperature and overcurrent conditions, making it more difficult (but certainly not impossible) to damage. This device is available only in a 36-pin surface-mount package, which can make it more challenging to use in prototype circuits.

Allegro A3985: This device isn't actually an H-bridge (a fully integrated device to which a motor or other load may be directly connected). Rather, this is an **H-bridge driver**: a type of device that combines control logic and gate driver outputs designed to be connected to the gates of external N-channel MOSFETs in both the high-side and low-side positions to create a full H-bridge. The designer specifies and connects the external MOSFETs separately. Moving the switching elements off-chip removes the most significant sources of heat within the device, and provides the designer with much greater flexibility and control over cost and performance of the motor-driving circuit. H-bridge driver chips allow for the control of much higher currents than single-chip solutions, provided the MOSFETs mated to the driver are capable of controlling the required currents. The A3985 incorporates two separate H-bridge driver modules, so each chip can control two motors. It requires a motor supply voltage between 12 and 50 V, is controlled via an SPI serial communications link (which we described in Chapter 7), and is available only in a surface-mount package. The A3985 is highly configurable, which allows designers to set it up to match the requirements of a wide variety of applications, but this often results in additional (sometimes extreme) effort to bring the circuit up and get everything working properly. However, if a circuit you're designing requires a 20 A motor, H-bridge driver chips with external MOSFETs may be the best choice.

23.4 SPEED CONTROL WITH PULSE WIDTH MODULATION

Another essential aspect of controlling motors is the ability to change the speed of rotation and the amount of torque produced. One simple but usually impractical means of achieving this is to adjust the supply voltage up and down as required. However, it is more effective and common to use the concept of PWM. This is a control technique where power to the motor is switched on and off rapidly, at rates high enough that the effects of the switching are negligible. The resulting effective voltage is then the average fraction of the time the power is on. This technique is also used in many other applications and is discussed elsewhere in the text.

Figure 23.17 illustrates the concept. The drive signal is switched on and off with a given period and is in the "on" state at voltage V_{ON} for a fixed fraction of the period. This "on" time is referred to as the **duty cycle** and is stated as a percentage, calculated as:

$$\text{Duty cycle (\%)} = \frac{\text{"On" time}}{\text{Period}} \times 100 \tag{23.9}$$

Figure 23.17 shows a duty cycle of 50%, with the resulting average *perceived* voltage of this PWM signal equal to 50% of the maximum voltage. The frequency of the PWM drive signal is calculated by taking the reciprocal of the period:

$$\text{PWM frequency} = \frac{1}{\text{Period}} \tag{23.10}$$

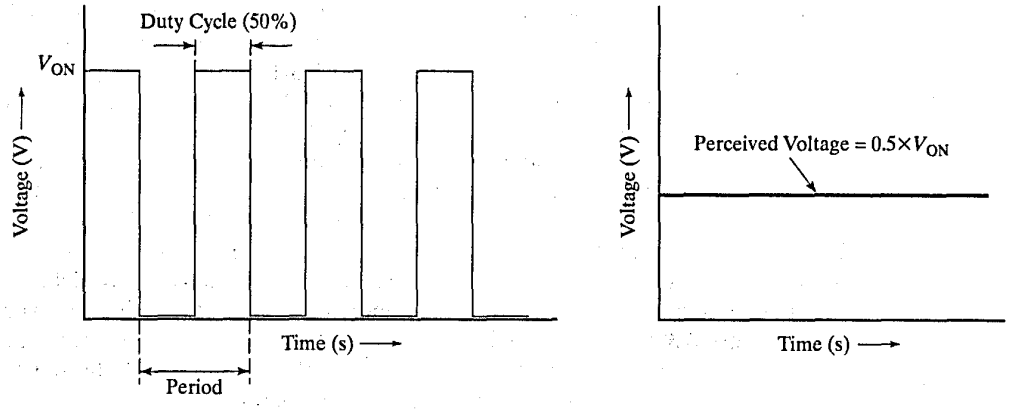

FIGURE 23.17
PWM with 50% duty cycle.

Changing the duty cycle of a PWM signal changes the average, or perceived, voltage level. For instance, adjusting the duty cycle shown in Figure 23.17 so that it is in the "on" state for 80% of the period will change the resulting perceived voltage as shown in Figure 23.18.

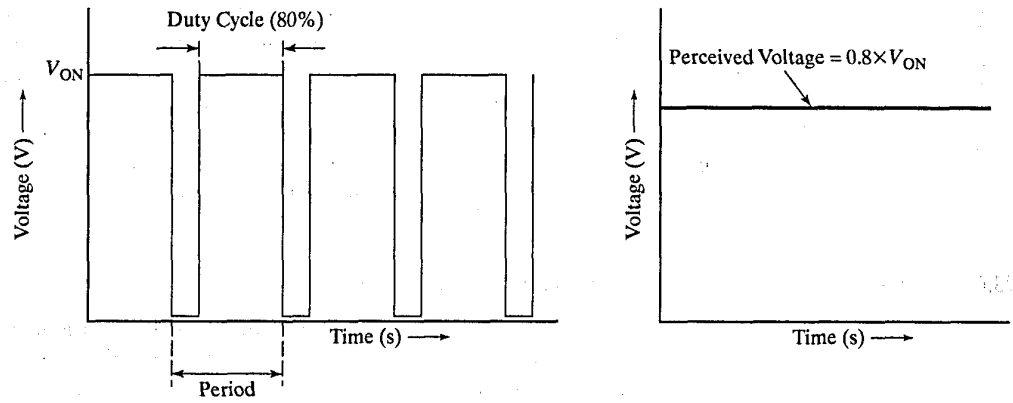

FIGURE 23.18
PWM with 80% duty cycle.

By increasing the duty cycle to 80%, the perceived voltage increases to 80% of V_{ON}. If PWM is implemented at a frequency that is too low, the result will be a jerky, stop-start response. Instead, the desired result is to approximate the "perceived voltage" as closely as possible with a minimum of perceptible ripple. When driving a permanent magnet brushed DC motor with PWM, the smoothing or filtering is very effectively performed by the physical inertia of the mechanical system and the electrical inductance of the motor windings. When done properly, the switching occurs too rapidly for the mechanical system to follow. For applications other than controlling the speed of a DC motor (e.g., approximating an analog output voltage with a microcontroller digital output pin), filtering circuits may be used to perform the required smoothing of the PWM output.

A typical mechanical time constant for a permanent magnet brushed DC motor is a few milliseconds, without a gearhead or other components attached. When selecting the frequency of the PWM drive signal, the designer must take into consideration the physical

characteristics of the mechanical drive system, as well as the electrical characteristics of the motor. Usually, a fairly wide range of PWM frequencies will give reasonable results.

At the lower limit of acceptable PWM frequencies, the mechanical time constant of the physical system (i.e., the amount of time the system requires to respond to a change in an input) will dominate. In the extreme case where a very long period is chosen—much longer than the mechanical time constant—the motor will very obviously start and stop. Consider the case of PWM period of 2 seconds (which is a 0.5 Hz PWM frequency): the motor will almost certainly accelerate and decelerate noticeably during each period. This is the undesirable **torque ripple** that will become much less noticeable if higher PWM frequencies are used.

The same effect that leads to torque ripple also affects the linearity of the response of the motor to changes in duty cycle. As an example, Figure 23.19 shows the speed of a geared DC motor vs. drive duty cycle at several PWM frequencies.

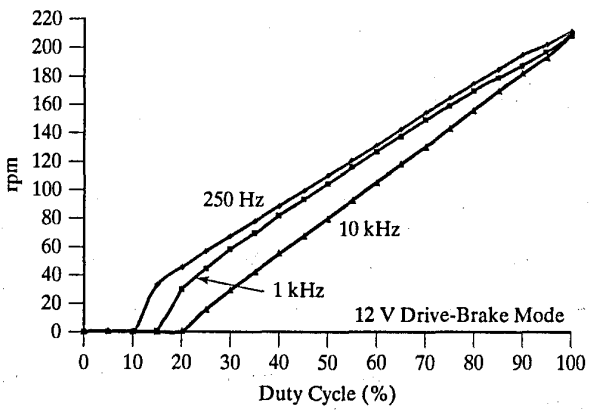

FIGURE 23.19
Effect of PWM frequency on the linearity of the duty cycle–speed relationship.

Increasing the PWM frequency results in better linearity of the relationship between duty cycle and motor speed. Usually, PWM frequencies in the range of 100 to 20,000 Hz will give workable results.

Regardless of the PWM frequency, torque ripple (and, by extension, **current ripple**) will occur to some degree. Higher PWM frequencies result in less ripple, but since current flow in the inductive motor coils is never able to follow the crisp edges of the PWM drive signal, ripple is unavoidable. Figure 23.20 illustrates how a steady state, low duty cycle PWM signal switching between a 0 V and V_{dc} with an average value of E induces a current flow in the motor's windings that has a minimum value of I_m and a maximum value of I_M. Shorter periods, T, will reduce the magnitude of the ripple, and hence the difference between I_m and I_M. Again, the rise and fall times of the current levels are governed by the motor winding's

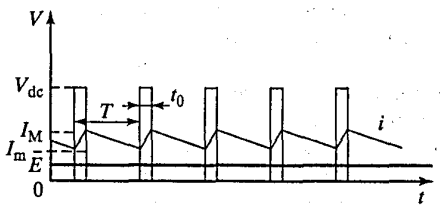

FIGURE 23.20
Current ripple characteristics for a steady state, low duty cycle PWM drive signal. (Courtesy of Maxon Motor AG.)

inductive time constant, $\tau = L/R$, Eq. (23.2). The amount of ripple current that can be tolerated will differ depending on the requirements of each application. One manufacturer, Maxon Precision Motors, recommends that ripple current be limited to 10% of their motors' maximum continuous current specifications [1].

When the PWM duty cycle is in the process of changing and isn't in steady state, the current flowing in the motor's coils will rise or fall in response to the changing average perceived voltage. The current can't immediately reflect instantaneous change in the duty cycle, just as it can't perfectly follow the crisp edges of the PWM signal as it switches on and off. Figure 23.21 shows how the motor current responds to changes in PWM duty cycle, again dictated by the time constant of the motor's coil windings.

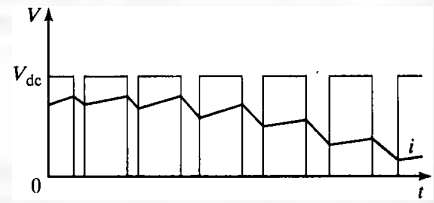

FIGURE 23.21

Motor current levels changing in response to a non-steady-state PWM signal. (Courtesy of Maxon Motor AG.)

PWM speed control can be readily combined with position or velocity feedback to create a servomotor controller (also somewhat inaccurately called a servomotor amplifier). We'll cover this concept in more detail in Chapter 28: Basic Closed-Loop Control.

Example 23.2

A motor will be required to spin continuously CW, and must supply 0.25 in · oz torque at 5,000 rpm when powered by 12 V. The motor has $K_T = 1.657$ in · oz/A, $K_e = 1.230$ V/krpm, and terminal resistance $R = 21.3\ \Omega$. Given that bidirectional control is not required, use a DS3658 IC, manufactured by National Semiconductor, and described as a "high-current peripheral driver" in the data sheet, to drive the motor.

1. Design a circuit to drive the motor, and show that none of the component specifications will be violated.
2. Specify the PWM duty cycle that will be required to obtain the specified torque and speed. Is there any uncertainty about what the resulting speed and torque will be? If so, where is this uncertainty introduced?

To answer the first question, we need to understand the function of the DS3658. It incorporates four independent switches, each of which can connect a load to ground or disconnect it from ground (low-side drive, Chapter 17), and each of which is controlled by a corresponding logic-level input. In the case of a motor or similar load, when one terminal of the load is connected to a power supply and the other terminal is connected to one of the DS3658's output pins, the motor will be turned on when the switch connects the output to ground and turned off when it is disconnected. One detail to look out for, though: the output pins are not perfect and can't connect a load to ground without some small losses. These losses manifest themselves as a slight rise in voltage from ground to the level at the output pin when it is on. This voltage rise is called the "saturation voltage," and is labeled on the data sheet as "Output Low Voltage."

The circuit we need to design must take into account that we will be using PWM drive to control the speed and torque of the motor (so this will require that we switch the motor on and off rapidly), we must protect the components from inductive kickback, and we must ensure that the components are capable of switching the current flowing through the motor. The circuit below, Figure 23.22, will be effective for controlling the switching, but the remaining specifications will need to be checked to ensure that none are being violated by the design.