DIODES AND DIODE CIRCUITS

1.25 Diodes

The circuit elements we've discussed so far (resistors, capacitors, and inductors) are all *linear*, meaning that a doubling of the applied signal (a voltage, say) produces a doubling of the response (a current, say). This is true even for the reactive devices (capacitors and inductors). These devices are also *passive*, meaning that they don't have a built-in source of power. And they are all two-terminal devices, which is self-explanatory.

Figure 1.66. Diode.

The diode (Fig. 1.66) is a very important and useful two-terminal passive *nonlinear* device. It has the *V–I* curve shown in Figure 1.67. (In keeping with the general philosophy of this book, we will not attempt to describe the solid-state physics that makes such devices possible.)

The diode's arrow (the anode terminal) points in the direction of forward current flow. For example, if the diode is in a circuit in which a current of 10mA is flowing from anode to cathode, then (from the graph) the anode is approximately 0.5 volt more positive than the cathode; this is called the "forward voltage drop." The reverse current, which is measured in the nanoamp range for a general-purpose diode (note the different scales in the graph for forward and reverse current), is almost never of any consequence until you reach the reverse breakdown voltage (also called the peak inverse voltage, PIV), typically 75 volts for a general-purpose diode like the 1N914. (Normally you never subject a diode to voltages large enough to cause reverse breakdown; the exception is the zener diode we mentioned earlier.) Frequently, also, the forward voltage drop of about 0.5 and 0.8 volt is of little concern, and the diode can be treated as a good approximation to an ideal one-way conductor. There are other important characteristics that distinguish the thousands of diode types available, e.g., maximum forward current, capacitance, leakage current, and reverse recovery time (see Table 1.1 for characteristics of some typical diodes).

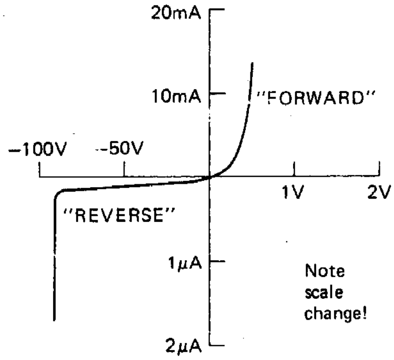

Figure 1.67. Diode *V–I* curve.

Before jumping into some circuits with diodes, we should point out two things: (a) A diode doesn't actually have a resistance (it doesn't obey Ohm's law). (b) If you put some diodes in a circuit, it won't have a Thévenin equivalent.

1.26 Rectification

A rectifier changes ac to dc; this is one of the simplest and most important applications of diodes (diodes are sometimes called rectifiers). The simplest circuit is shown in Figure 1.68. The "ac" symbol represents a source of ac voltage; in electronic circuits it is usually provided by a transformer, powered from the ac power line. For a sine-wave input that is much larger than the forward drop (about 0.6V for silicon diodes, the usual type), the output will look like that in Figure 1.69. If you think of the diode as a one-way conductor, you won't have any trouble understanding how the circuit works. This circuit is called a *half-wave rectifier*, because only half of the input waveform is used.

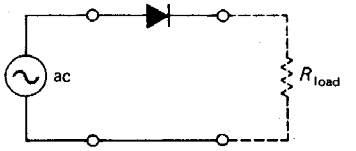

Figure 1.68. Half-wave rectifier.

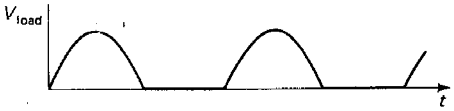

Figure 1.69

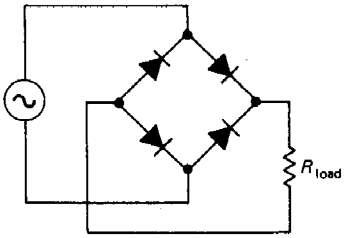

Figure 1.70. Full-wave bridge rectifier.

Figure 1.70 shows another rectifier circuit, a full-wave bridge. Figure 1.71 shows the voltage across the load for which the whole input waveform is used. The gaps at zero voltage occur because of the diodes' forward voltage drop. In this circuit, two diodes are always in series with the input; when you design low-voltage power supplies, you have to remember that.

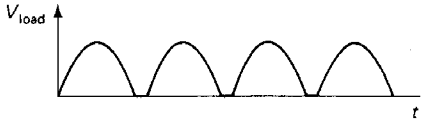

Figure 1.71

1.27 Power-supply filtering

The preceding rectified waveforms aren't good for much as they stand. They're dc only in the sense that they don't change polarity. But they still have a lot of "ripple" (periodic variations in voltage about the steady value) that has to be smoothed out in order to generate genuine dc. This we do by tacking on a low-pass filter (Fig. 1.72). Actually, the series resistor is unnecessary and is always omitted (although you sometimes see a very small resistor used to limit the peak rectifier current). The reason is that the diodes prevent flow of current back out of the capacitors, which are really serving more as energy-storage devices than as part of a classic low-pass filter. The energy stored in a capacitor is $U = \frac{1}{2}CV^2$. For C in farads and V in volts, U comes out in joules (watt-seconds).

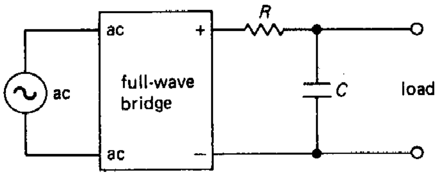

Figure 1.72

The capacitor value is chosen so that

$$R_{\text{load}}C \gg 1/f$$

(where f is the ripple frequency, here 120Hz) in order to ensure small ripple, by making the time constant for discharge much longer than the time between recharging. We will make this vague statement clearer in the next section.

Calculation of ripple voltage

It is easy to calculate the approximate ripple voltage, particularly if it is small compared with the dc (see Fig. 1.73). The load causes the capacitor to discharge somewhat between cycles (or half cycles, for full-wave rectification). If you assume that the load current stays constant (it will, for small ripple), you have

$$\Delta V = \frac{I}{C}\Delta t \qquad \left(\text{from } I = C\frac{dV}{dt}\right)$$

Just use $1/f$ (or $1/2f$ for full-wave rectification) for Δt (this estimate is a bit on the safe side, since the capacitor begins charging again in less than a half cycle). You get

$$\Delta V = \frac{I_{\text{load}}}{fC} \qquad \text{(half wave)}$$

$$\Delta V = \frac{I_{\text{load}}}{2fC} \qquad \text{(full wave)}$$

(While teaching electronics we've noticed that students love to memorize these equations! An informal poll of the authors showed that two out of two engineers don't memorize them. Please don't waste brain cells that way – instead, learn how to derive them.)

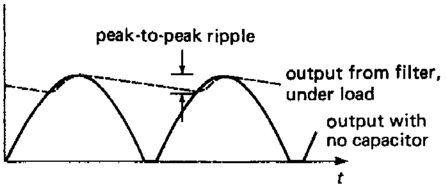

Figure 1.73. Power-supply ripple calculation.

If you wanted to do the calculation without any approximation, you would use the exact exponential discharge formula. You would be misguided in insisting on that kind of accuracy, though, for two reasons:

1. The discharge is an exponential only if the load is a resistance; many loads are not. In fact, the most common load, a *voltage regulator*, looks like a constant-current load.

2. Power supplies are built with capacitors with typical tolerances of 20% or more. Realizing the manufacturing spread, you design conservatively, allowing for the worst-case combination of component values.

In this case, viewing the initial part of the discharge as a ramp is in fact quite accurate, especially if the ripple is small, and in any case it errs in the direction of conservative design – it overestimates the ripple.

EXERCISE 1.27

Design a full-wave bridge rectifier circuit to deliver 10 volts dc with less than 0.1 volt (pp) ripple into a load drawing up to 10mA. Choose the appropriate ac input voltage, assuming 0.6 volt diode drops. Be sure to use the correct ripple frequency in your calculation.

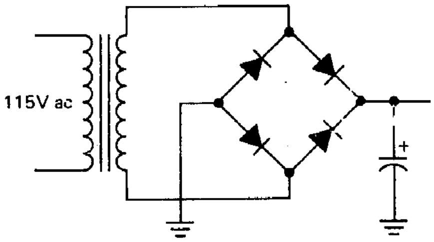

Figure 1.74. Bridge rectifier circuit. The polarity marking and curved electrode indicate a polarized capacitor, which must not be allowed to charge with the opposite polarity.

1.28 Rectifier configurations for power supplies

Full-wave bridge

A dc power supply using the bridge circuit we just discussed looks as shown in Figure 1.74. In practice, you generally buy the bridge as a prepackaged module. The smallest ones come with maximum current ratings of 1 amp average, with breakdown voltages going from 100 volts to 600 volts, or even 1000 volts. Giant bridge rectifiers are available with current ratings of 25 amps or more. Take a look at Table 6.4 for a few types.

Center-tapped full-wave rectifier

The circuit in Figure 1.75 is called a center-trapped full-wave rectifier. The output voltage is half what you get if you use a bridge rectifier. It is not the most efficient circuit in terms of transformer design, because each half of the secondary is used only half the time. Thus the current through the winding during that time is twice what it would be for a true full-wave circuit. Heating in the windings, calculated from Ohm's law, is I^2R, so you have four times the heating half the time, or twice the average heating of an equivalent full-wave bridge circuit. You would have tc choose a transformer with a current rating 1.4 (square root of 2) times as large, as compared with the (better) bridge circuit; besides costing more, the resulting supply would be bulkier and heavier.

EXERCISE 1.28

This illustration of I^2R heating may help you understand the disadvantage of the center-tapped rectifier circuit. What fuse rating (minimum) is required to pass the current waveform shown in Figure 1.76, which has 1 amp average current? Hint: A fuse "blows out" by melting (I^2R heating) a metallic link, for steady currents larger than its rating. Assume for this problem that the thermal time constant of the fusible link is much longer than the time scale of the square wave, i.e., that the fuse responds to the value of I^2 averaged over many cycles.

Split supply

A popular variation of the center-tapped full-wave circuit is shown in Figure 1.77. It gives you split supplies (equal plus and minus voltages), which many circuits need. It is an efficient circuit, because both halves of the input waveform are used in each winding section.

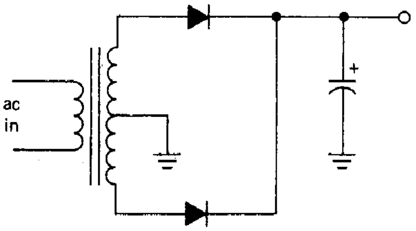

Figure 1.75. Full-wave rectifier using center-tapped transformer.

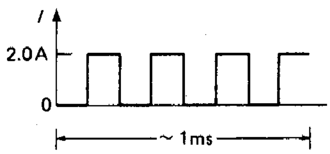

Figure 1.76

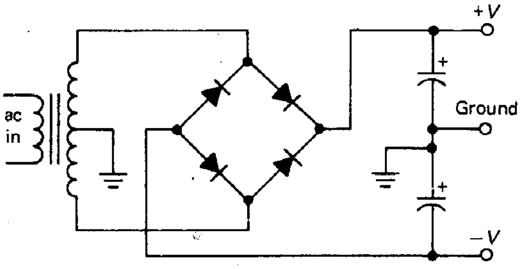

Figure 1.77. Dual-polarity (split) supply.

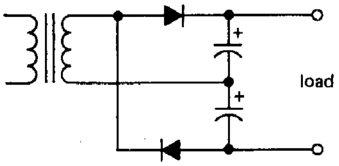

Figure 1.78. Voltage doubler.

☐ *Voltage multipliers*

The circuit shown in Figure 1.78 is called a voltage doubler. Think of it as two half-wave rectifier circuits in series. It is officially a full-wave rectifier circuit, since

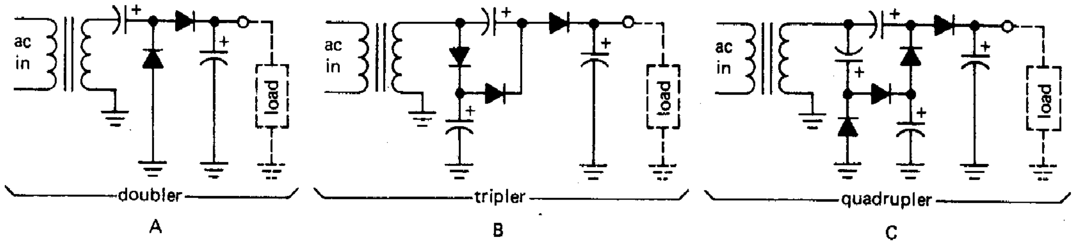

Figure 1.79. Voltage multipliers; these configurations don't require a floating voltage source.

both halves of the input waveform are used – the ripple frequency is twice the ac frequency (120Hz for the 60Hz line voltage in the United States).

Variations of this circuit exist for voltage triplers, quadruplers, etc. Figure 1.79 shows doubler, tripler, and quadrupler circuits that let you ground one side of the transformer.

1.29 Regulators

By choosing capacitors that are sufficiently large, you can reduce the ripple voltage to any desired level. This brute-force approach has two disadvantages:

1. The required capacitors may be prohibitively bulky and expensive.

2. Even with the ripple reduced to negligible levels, you still have variations of output voltage due to other causes, e.g., the dc output voltage will be roughly proportional to the ac input voltage, giving rise to fluctuations caused by input line voltage variations. In addition, changes in load current will cause the output voltage to change because of the finite internal resistances of the transformer, diode, etc. In other words, the Thévenin equivalent circuit of the dc power supply has $R > 0$.

A better approach to power-supply design is to use enough capacitance to reduce ripple to low levels (perhaps 10% of the dc voltage), then use an active *feedback circuit* to eliminate the remaining ripple. Such a feedback circuit "looks at" the output, making changes in a controllable series resistor (a transistor) as necessary to keep the output constant (Fig. 1.80).

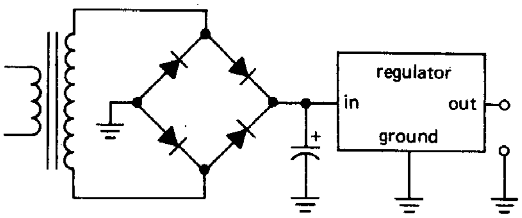

Figure 1.80. Regulated dc power supply.

These voltage regulators are used almost universally as power supplies for electronic circuits. Nowadays complete voltage regulators are available as inexpensive integrated circuits (priced under one dollar). A power supply built with a voltage regulator can be made easily adjustable and self-protecting (against short circuits, overheating, etc.), with excellent properties as a voltage source (e.g., internal resistance measured in milliohms). We will deal with regulated dc power supplies in Chapter 6.

1.30 Circuit applications of diodes

Signal rectifier

There are other occasions when you use a diode to make a waveform of one polarity only. If the input waveform isn't a sine wave, you usually don't think of it as a rectification in the sense of a power supply. For instance, you might want a train of pulses corresponding to the rising edge of a square wave. The easiest way is to rectify the differentiated wave (Fig. 1.81). Always keep in mind the 0.6 volt (approximately) forward drop of the diode. This circuit, for instance, gives no output for square waves smaller than 0.6 volt pp. If this is a problem, there are various tricks to circumvent this limitation. One possibility is to use *hot carrier diodes* (Schottky diodes), with a forward drop of about 0.25 volt (another device called a *back diode* has nearly zero forward drop, but its usefulness is limited by very low reverse breakdown voltage).

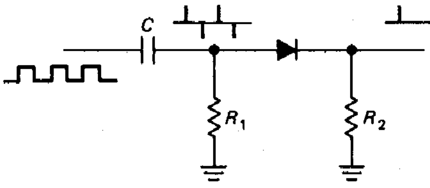

Figure 1.81

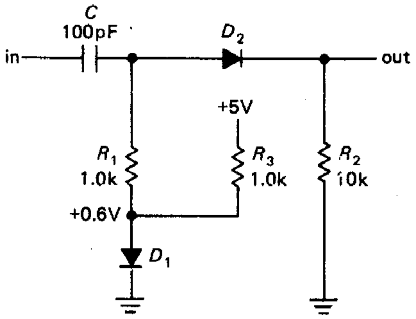

Figure 1.82. Compensating the forward voltage drop of a diode signal rectifier.

A possible circuit solution to this problem of finite diode drop is shown in Figure 1.82. Here D_1 compensates D_2's forward drop by providing 0.6 volt of *bias* to hold D_2 at the threshold of conduction. Using a diode (D_1) to provide the bias (rather than, say, a voltage divider) has several advantages: There is nothing to adjust, the compensation will be nearly perfect, and changes of the forward drop (e.g., with changing temperature) will be compensated properly. Later we will see other instances of matched-pair compensation of forward drops in diodes, transistors, and FETs: it is a simple and powerful trick.

Diode gates

Another application of diodes, which we will recognize later under the general heading of *logic*, is to pass the higher of two voltages without affecting the lower. A good example is *battery backup*, a method of keeping something running (e.g, a precision electronic clock) that must not stop when there is a power failure. Figure 1.83 shows a circuit that does the job. The battery does nothing until the power fails; then it takes over without interruption.

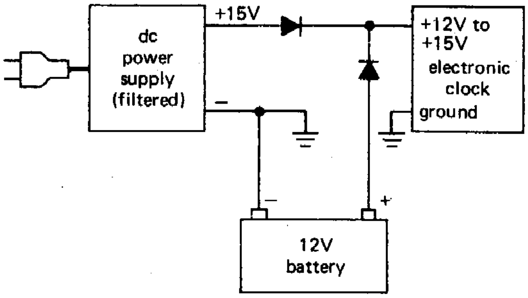

Figure 1.83. Diode OR gate: battery backup.

EXERCISE 1.29

Make a simple modification to the circuit so that the battery is charged by the dc supply (when power is on, of course) at a current of 10mA (such a circuit is necessary to maintain the battery's charge).

Diode clamps

Sometimes it is desirable to limit the range of a signal (i.e., prevent it from exceeding certain voltage limits) somewhere in a circuit. The circuit shown in Figure 1.84 will accomplish this. The diode prevents the output from exceeding about +5.6 volts, with no effect on voltages less than that (including negative voltages); the only limitation is that the input must not go so negative that the reverse breakdown voltage of the diode is exceeded (e.g., −70V for a 1N914). Diode clamps are standard equipment on all inputs in the CMOS family of digital logic. Without them, the delicate input circuits are easily destroyed by static electricity discharges during handling.

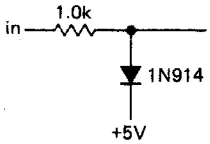

Figure 1.84. Diode voltage clamp.

EXERCISE 1.30

Design a symmetrical clamp, i.e., one that confines a signal to the range −5.6 volts to +5.6 volts.

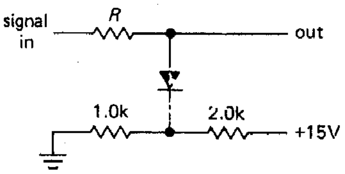

Figure 1.85

A voltage divider can provide the reference voltage for a clamp (Fig. 1.85). In this case you must ensure that the impedance looking into the voltage divider (R_{vd}) is small compared with R, because what you have looks as shown in Figure 1.86 when the voltage divider is replaced by its Thévenin equivalent circuit. When the diode conducts (input voltage exceeds clamp voltage), the output is really just the output of a voltage divider, with the Thévenin equivalent resistance of the voltage reference as the lower resistor (Fig. 1.87). So, for the values shown, the output of the clamp for a triangle-wave input would look as shown in Figure 1.88. The problem is that the voltage divider doesn't provide a stiff reference, in the language of electronics. A stiff voltage source is one that doesn't bend easily, i.e., it has low internal (Thévenin) impedance.

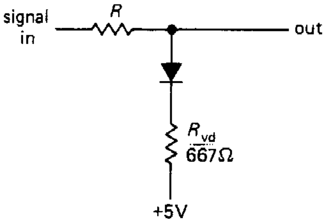

Figure 1.86

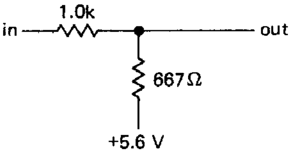

Figure 1.87

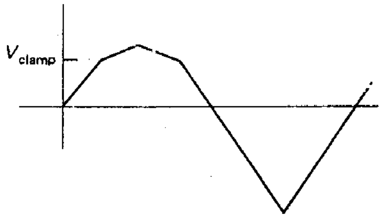

Figure 1.88

A simple way to stiffen the clamp circuit of Figure 1.85, at least for *high-frequency* signals, is to add a bypass capacitor across the 1k resistor. For example, a 15μF capacitor to ground reduces the impedance seen looking into the divider below 10 ohms for frequencies above 1kHz. (You could similarly add a bypass capacitor across D_1 in Fig. 1.82.) Of course, the effectiveness of this trick drops at low frequencies, and it does nothing at dc.

In practice, the problem of finite impedance of the voltage-divider reference can be easily solved using a transistor or operational amplifier (op-amp). This is usually a better solution than using very small resistor values, because it doesn't consume large currents, yet it provides impedances of a few ohms or less. Furthermore, there are other ways to construct a clamp, using an op-amp as part of the clamp circuit. You will see these methods in Chapter 4.

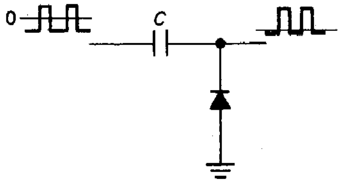

Figure 1.89. dc restoration.

One interesting clamp application is "dc restoration" of a signal that has been ac-coupled (capacitively coupled). Figure 1.89 shows the idea. This is particularly important for circuits whose inputs look like diodes (e.g., a transistor with grounded emitter); otherwise an ac-coupled signal will just fade away.

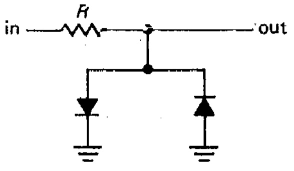

Figure 1.90. Diode limiter.

Limiter

One last clamp circuit is shown in Figure 1.90. This circuit limits the output "swing" (again, a common electronics term) to one diode drop, roughly 0.6 volt. That might seem awfully small, but if the next stage is an amplifier with large voltage amplification, its input will always be near zero volts; otherwise the output is in "saturation" (e.g., if the next stage has a gain of 1000 and operates from ±15V supplies, its input must stay in the range ±15mV in order for its output not to saturate). This clamp circuit is often used as input protection for a high-gain amplifier.

Diodes as nonlinear elements

To a good approximation the forward current through a diode is proportional to an exponential function of the voltage across it at a given temperature (for a discussion of the exact law, see Section 2.10). So you can use a diode to generate an output voltage proportional to the logarithm of a current (Fig. 1.91). Because V hovers in the region of 0.6 volt, with only small voltage changes that reflect input current variations, you can generate the input current with a resistor if the input voltage is much larger than a diode drop (Fig. 1.92).

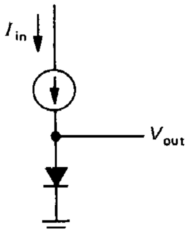

Figure 1.91. Exploiting the diode's nonlinear V–I curve: logarithmic converter.

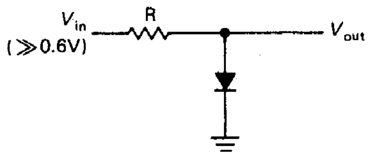

Figure 1.92

In practice, you may want an output voltage that isn't offset by the 0.6 volt diode drop. In addition, it would be nice to have a circuit that is insensitive to changes in temperature. The method of diode drop compensation is helpful here (Fig. 1.93). R_1 makes D_2 conduct, holding point A at about −0.6 volt. Point B is then near ground (making I_{in} accurately proportional to V_{in}, incidentally). As long as the two (identical) diodes are at the same temperature, there is good cancellation of the forward drops, except, of course, for the difference owing to input current through D_1, which produces the desired output. In this circuit, R_1 should be chosen so that the current through D_2 is much larger than the maximum input current, in order to keep D_2 in conduction.

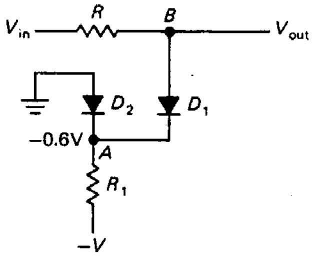

Figure 1.93. Diode drop compensation in the logarithmic converter.

In the chapter on op-amps we will examine better ways of constructing logarithmic converter circuits, along with careful methods of temperature compensation. With such methods it is possible to construct logarithmic converters accurate to a few percent over six decades or more of input current. A better understanding of diode and transistor characteristics, along with an understanding of op-amps, is necessary first. This section is meant to serve only as an introduction for things to come.

1.31 Inductive loads and diode protection

What happens if you open a switch that is providing current to an inductor? Because inductors have the property

$$V = L\frac{dI}{dt}$$

it is not possible to turn off the current suddenly, since that would imply an infinite voltage across the inductor's terminals. What happens instead is that the voltage across the inductor suddenly rises and keeps rising until it forces current to flow. Electronic devices controlling inductive loads can be easily damaged, especially the component that "breaks down" in order to satisfy the inductor's craving for continuity of current. Consider the circuit in Figure 1.94. The switch is initially closed, and current is flowing through the inductor (which might be a relay, as will be described later). When the switch is opened, the inductor "tries" to keep current flowing from A to B, as it had been. That means that terminal B goes positive relative to terminal A. In a case like this it may go 1000 volts positive before the switch contact "blows over." This shortens the life of the switch and also generates impulsive interference that may affect other circuits nearby. If the switch happens to be a transistor, it would be an understatement to say that its life is shortened; its life is ended!

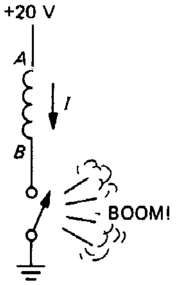

Figure 1.94. Inductive "kick."

The best solution is to put a diode across the inductor, as in Figure 1.95. When the switch is on, the diode is back-biased (from the dc drop across the inductor's winding resistance). At turn-off the diode goes into conduction, putting the switch terminal a diode drop above the positive supply voltage. The diode must be able to handle the initial diode current, which equals the steady current that had been flowing through the inductor; something like a 1N4004 is fine for nearly all cases.

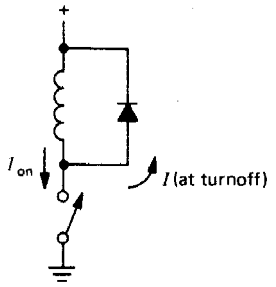

Figure 1.95. Blocking inductive kick.

The only disadvantage of this protection circuit is that it lengthens the decay of current through the inductor, since the rate of change of inductor current is proportional to the voltage across it. For applications where the current must decay quickly (high-speed impact printers, high-speed relays, etc.), it may be better to put a resistor across the inductor, choosing its value so that $V_{\text{supply}} + IR$ is less than the maximum allowed voltage across the switch. (For fastest decay with a given maximum voltage, a zener could be used instead, giving a ramp-down of current rather than an exponential decay.)

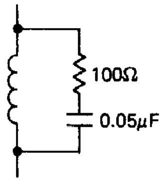

Figure 1.96. *RC* "snubber" for suppressing inductive kick.

For inductors driven from ac (transformers, ac relays), the diode protection just described will not work, since the diode will conduct on alternate half cycles when the switch is closed. In that case a good solution is an *RC* "snubber" network (Fig. 1.96). The values shown are typical for small inductive loads driven from the ac power line. Such a snubber should be included in all instruments that run from the ac power line, since a transformer is inductive. An alternative protection device is a metal-oxide varistor, or transient suppressor, an inexpensive device that looks something like a disc ceramic capacitor and behaves electrically like a bi-directional zener diode. They are available at voltage ratings from 10 to 1000 volts and can handle transient currents up to thousands of amperes (see Section 6.11 and Table 6.2). Putting a transient suppressor across the ac power-line terminals makes good sense in a piece of electronic equipment, not only to prevent inductive spike interference to other nearby instruments but also to prevent occasional large power-line spikes from damaging the instrument itself.