9

BASIC ELECTROMECHANICAL DEVICES

In this chapter, we'll take a look at a sampling of *electromechanical devices*—components that translate electronic signals into motion and vice versa. These devices are instrumental for tasks such as automatically switching other circuits on and off, playing and capturing sound, and spinning the wheels of an RC car.

As with passive components, the electromechanical category is not especially well defined. It is generally understood to include switches and motors, but knob-operated potentiometers are more commonly classified as passives (see Chapter 8). It's best not to overthink it. The categories matter only insofar as they help you pass exams and browse online catalogs.

As a matter of practicality, we'll focus on common electromechanical components that are fundamental to a wide range of electronic devices. More specialized sensors or actuators also exist, but they generally operate on the same principles. As for complex "computerized" electromechanical sensors, such as accelerometers and gyroscopes, we'll consider them in Chapter 17.

Connectors

Connectors need little introduction. They're simple mechanical devices that provide temporary coupling between circuits. Compared with permanently soldered wires, connectors make electronics easier to transport, reconfigure, and repair.

A taxonomy of every connector on the market would fill hundreds of pages—there's an endless variety. Some connectors are designed to be compact, others to be rugged, and still others to safely handle high currents and voltages. A picture is worth a thousand words, so for an engaging exploration of common designs, I recommend the highly visual book *Open Circuits* (No Starch Press, 2022). Another way to brush up on connector types is to browse the catalogs of the electronics retailers discussed in the appendix at the end of this book.

Despite the wide range of options, we tend to stick to the classics. The most common small connector in use today is probably USB-C; it supports flexible power delivery and high-speed data transfers between computers and peripherals. Another common type is the tubular barrel plug, still widely used to supply fixed DC voltages to low-power household appliances. Many readers will also recognize the various sizes of audio jacks for wired headphones, speakers, and microphones.

Inside electronic devices, we frequently encounter so-called *headers*—rows of metal pins on a 1.27 mm or 2.54 mm grid that mate with similarly shaped plastic receptacles. It's also common to see *wire-to-board* connectors that clip or crimp onto cables. Finally, compact electronics also use delicate, amber-colored *flexible printed circuit (FPC)* strips providing internal connections to displays, touchscreens, keypads, and so forth.

On the workbench, you'll find BNC connectors for signal generators and oscilloscope probes, along with single-conductor *banana plugs* used with multimeters and lab power supplies.

The side of a connector with exposed prongs or that slides into a receptacle is commonly called *male*; the other side is termed *female*. The words *plug* and *socket* refer to cable-side and device-side connections, respectively. Finally, *jack* is somewhat ambiguous, with different meanings in different parts of the world. Figure 9-1 shows several ways of representing sockets and plugs on circuit schematics.

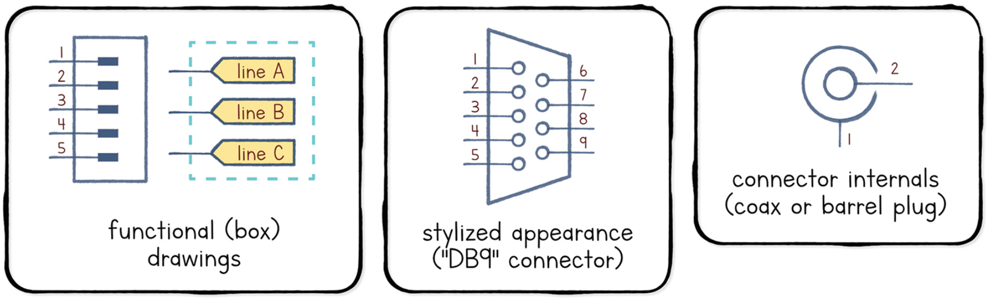

Figure 9-1: Some of the ways to draw connectors

The styles range from simple boxes to more elaborate visual metaphors. You can choose any approach you like, but make sure your drawings are unambiguous: Always include pin numbers or labels, and make the connector appear as a distinct, logical whole.

Switches

Switches should be familiar to any reader who has ever turned on a light or used a vending machine. These devices open or close electrical connections in response to a mechanical action, such as flipping a rocker, pushing a button, or turning a knob.

As with connectors, we could spend hours discussing the many form factors and styles of switches. There are miniature tactile (dome) switches that can be mounted directly onto a printed circuit board to function as reset buttons, simple keypads, and so on. There are also elevator-style illuminated pushbuttons, slide switches, rocker switches, paddle switches, and many others. Instead of rattling off all the options, let's focus on what can't be grokked from catalog photos alone.

Electrical Ratings and Architecture

Every switch is rated for a particular current and voltage. If the current limit is significantly exceeded for more than a split second, the contacts may overheat and possibly weld together. Exceeding the voltage limit usually has less dramatic consequences, but excessive sparking—which can occur briefly even in low-voltage circuits when disconnecting a large inductive load—can cause premature wear. At sufficiently high voltages, sustained electrical arcing could also become a concern.

In practice, miniature tactile switches may be rated for as little as 20 to 50 mA at 12 V; small slide switches are commonly designed to handle at least 300 mA; and larger panel-mounted rocker switches can usually carry between 5 and 20 A at 250 V. When the component's dimensions allow, these ratings are typically stamped on the side of the switch.

Current and voltage limits aside, switches are also distinguished by their mechanical action. A switch is said to have a *momentary* action if it stays in position only while an external force is applied. In contrast, *latching* or *sustained* action means that the mechanism stays in position until toggled back. A doorbell is a typical example of a momentary switch, while a wall-mounted light switch is a common sustained type. Most momentary switches are normally open (NO), but normally closed (NC) variants are also available from some manufacturers for specialized applications.

Another distinguishing feature of a switch is the number of *poles*. Each pole is a separate electrical circuit operated by a common mechanical assembly—essentially multiple units ganged together. Most switches have a single pole (SP), although double-pole (DP) mechanisms are sometimes used in outdoor lighting, motor control, and similar applications. Higher pole counts (3P, 4P, 5P, and so on) are fairly rare, partly because few applications require toggling five completely separate circuits at once.

The number of poles is often cited together with the number of *throws*. A throw refers to the set of electrical paths through which the mechanism cycles. A single-throw (ST) switch still has two positions—since a one-position switch would be useless—but it provides a single on/off path between two terminals, possibly repeated for each switch unit in a multi-pole assembly. A double-throw (DT) mechanism can toggle each pole between two electrical paths; for example, in a single-pole double-throw switch with three terminals labeled A, B, and C, the connection might alternate between A-B and A-C.

Switches with three or more throws—and thus three or more positions—are uncommon, but they crop up in some household appliances. For example, an old-fashioned fan may use a 3T switch to select among three motor speeds (low, medium, and high). Figure 9-2 shows some of the most common schematic symbols for switches.

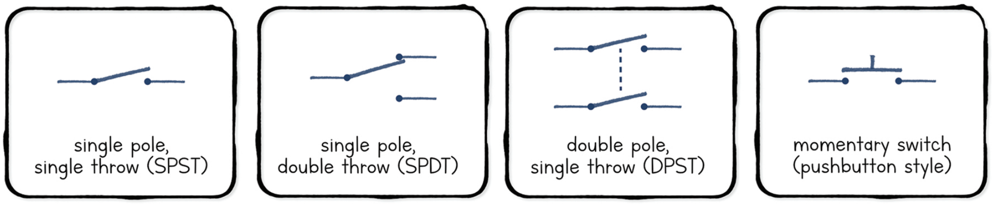

Figure 9-2: Schematic symbols for switches

In the third panel (DPST), note the use of a dashed line to indicate a mechanical linkage between the two electrically separate paths in a single package.

Switch Bounce

If you pry open a typical switch, you'll probably find that its internals are quite complex. Inside, there might be a small spring-loaded roller that travels over a tiny hump or an intricately shaped metal leaf that "pops" when pressed firmly enough. The purpose of these mechanisms is to keep the switch firmly in place until sufficient force is applied, at which point it snaps immediately into the new position. This nonlinear behavior provides satisfying tactile feedback, reduces sensitivity to vibration and other accidental inputs, and prevents the switch from resting in any state other than "fully off" or "fully on."

The flip side of this approach is that once the mechanism gives in, things happen quickly: The contacts are pulled apart or slammed together at a considerable speed. Because metal is elastic, this motion produces some vibration that manifests as electrical glitches lasting from tens of microseconds to a few milliseconds (Figure 9-3). The effect is known as *switch bounce*.

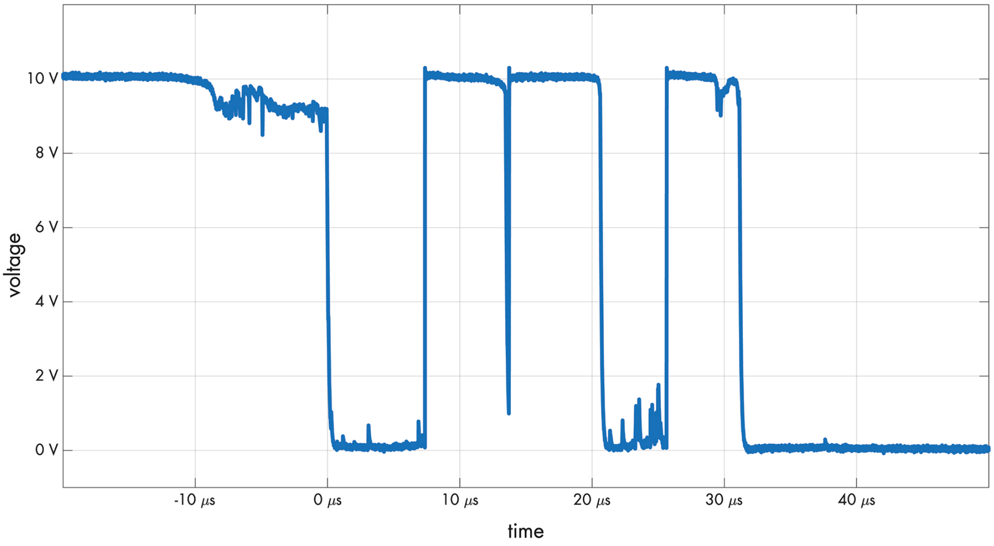

Figure 9-3: A fairly typical pattern of switch bounce when a switch goes from on to off

The glitches are imperceptible to humans and inconsequential in most analog applications, but they last an eternity from the perspective of a digital chip. If the bounce pattern isn't taken into account, embedded computers can misread a single input as a series of independent key presses, producing confusing results and sending novice circuit designers on a wild-goose chase.

In earlier days, we used to incorporate electronic filters to suppress high-frequency switch noise before it reached a high-speed chip. Nowadays, the problem is more commonly handled in software. In most cases, requiring around 10 milliseconds to pass between successive inputs is enough to filter out switch bounce without adding any human-perceptible lag.

Switch-Based Sensors

Switches are a common building block for simple sensors. For example, reed switches, used in home security systems, contain a permanent magnet and are activated by bringing another magnet nearby. Similarly, mechanical thermostats use a bimetallic strip that deflects in response to changes in temperature, bringing contacts together or moving them apart.

Another interesting type of sensor in this category is the *incremental rotary encoder*. As the name implies, rotary encoders convert rotation into electrical signals; their applications range from car radio volume knobs to position sensors in industrial machinery. In such uses, they serve as longer-lasting alternatives to delicate potentiometers. The downside is that encoders output pulses, not analog readings, so the signal needs to be interpreted by an embedded computer to convert the pulses into angular position.

Rotary encoders can be constructed in a number of ways, but the simplest version consists of two stationary pins that glide across a rotating, rosette-shaped conductive disk. Because the pins are slightly offset, the order in which they contact the disk can be used to infer the direction of rotation (Figure 9-4). Meanwhile, the frequency of the generated pulses represents the rotation speed.

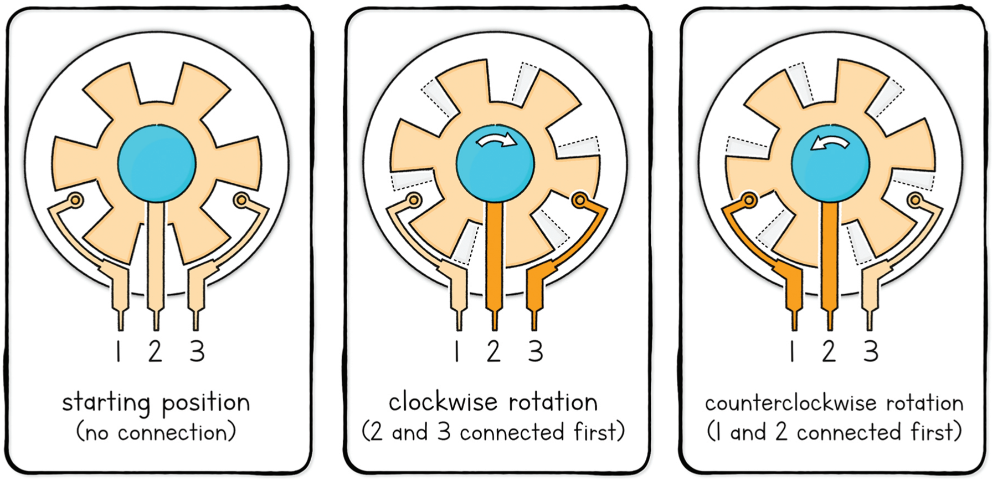

Figure 9-4: A conceptual illustration of a simple rotary encoder

There's no well-established schematic symbol for this class of devices. Components that don't have dedicated symbols are often illustrated as blank boxes with a descriptive label on the side. That said, some authors may opt to include additional visual cues of their choice.

Simple Magnetic Actuators

In Chapter 6, we noted that the flow of electricity through a coil of wire can produce a coherent magnetic field. This field can attract or repel permanent magnets, and it can also more weakly attract ferromagnetic materials that don't produce coherent magnetic fields of their own but contain microscopic, reorientable magnetic domains in the crystal lattice. The materials that exhibit this property include iron, cobalt, nickel, and some grades of steel.

A coil of wire used to attract metals is known as an *electromagnet*. An electromagnet can be used on its own—for example, to sort scrap metal—but it's more common to combine it with a captive plunger that's pushed or pulled by the field. This contraption, usually referred to as a *solenoid*, can be found in some electronically controlled door locks, hydraulic valves, and vintage pinball machines.

The coil of a typical electromagnet consists of very thin wire wrapped around the electromagnet hundreds or thousands of times. This construction serves two purposes. First, it maximizes the strength of the resulting magnetic field. Second, the length of the conductor gives the coil substantial resistance, limiting the current to a safe maximum at the rated supply voltage. Most electromagnets can be operated at higher voltages to deliver more punch, but the duty cycle must be short to allow the coils to cool down.

Solenoids are a niche product, but they share their operating principle with two additional common electronic components: speakers and relays.

Speakers

A speaker is a solenoid with a twist: Instead of a fixed coil that actuates a plunger, there's a permanent magnet affixed to the frame and an electromagnet coil mounted on a lightweight diaphragm made of paper or plastic. Because of the great difference in mass, the coil doesn't move the magnet. Instead, it moves itself back and forth in response to the applied current. This motion also moves the diaphragm in a similar fashion. Because the diaphragm has a fairly large surface area, it scoops up a decent volume of air and produces audible sound waves (Figure 9-5).

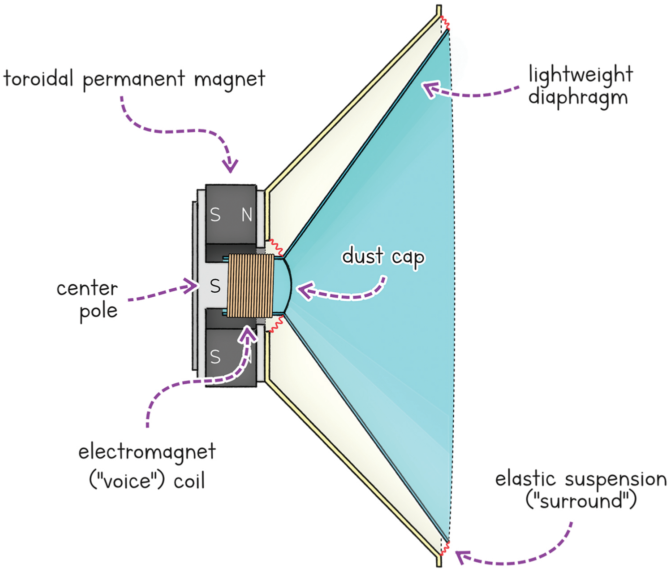

Figure 9-5: The internals of a simple speaker

Like solenoids, speakers are designed with a nominal coil resistance, usually 8 Ω, although 4 Ω and 16 Ω coils aren't unheard of. Another key parameter is the maximum continuous power rating. For small speakers, this usually falls between 250 mW and 2 W; home audio components may be rated for 50 W or more. As expected, exceeding the rating risks overheating the coil.

Acoustically, speakers are complicated beasts. Size matters: Small diaphragms have less inertia, and small coils have less inductance, so compact units reproduce high-frequency sounds more effectively. On the flip side, tiny speakers can't move a whole lot of air, so they can't deliver good low-frequency rumble. This problem can be mitigated to some extent by electronically boosting some audio frequencies and attenuating others (see Chapter 12). Still, the hi-fi solution is to spend a bit more money and have multiple speakers of different designs working together.

To get a sense of how speakers work and how well they handle different frequencies, you can try hooking up a small speaker to a function generator. Crank up the voltage to 5 to 10 V, and then try outputting sine waves at different frequencies;

around 1 kHz is a sensible starting point. Note how the apparent loudness changes with frequency. Some of this reflects the nature of human hearing, but you will also notice marked differences between speakers of different sizes. Next, check out square waves. They should sound much louder due to their fast-rising and fast-falling edges, but the sound will also have a much harsher, tinnier quality.

You can also connect the speaker to an oscilloscope and see what happens if you speak into it or push the diaphragm with your finger. You should see AC voltages across the terminals. We'll discuss why this happens shortly.

Relays

A *relay* is a solenoid-operated switch. Like normal switches, relays can have a varying number of poles and throws. Most relays feature momentary switching action, though several varieties of latching mechanisms are also available. The devices range from miniature models used in precision electronic instruments to fist-sized *contactors* found in high-voltage appliances, such as whole-house generators and air conditioners. Figure 9-6 shows some common schematic symbols for relays and speakers.

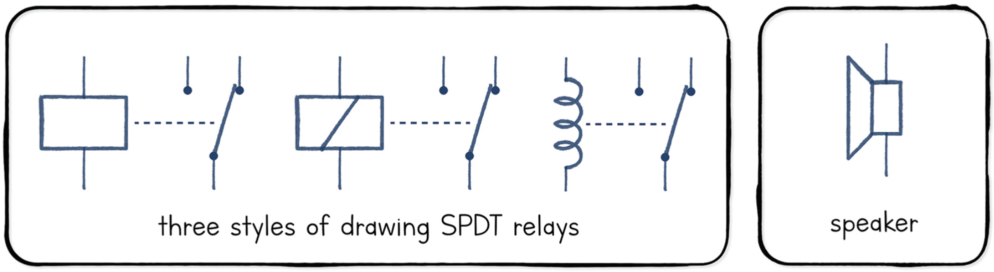

Figure 9-6: Symbols for relays and speakers

If you have a miniature non-latching relay and a capacitor in the vicinity of 100 μF, there's an interesting experiment you can conduct (Figure 9-7).

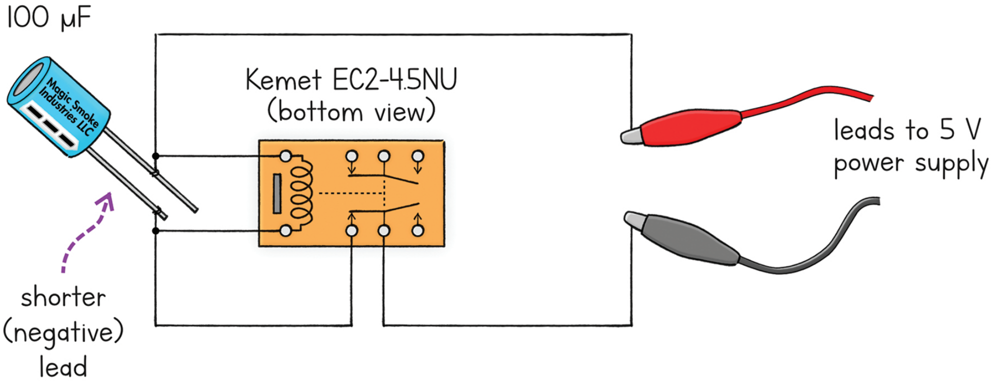

Figure 9-7: An experiment with a miniature relay (bottom view)

The circuit uses a DPDT relay made by Kemet and labeled EC2-4.5NU. Other relays might have a different pinout, but most of them should work fine if the necessary wiring adjustments are made.

Once the circuit is connected to a 4 to 5 V supply at 500 mA, the relay should start clicking at a frequency of several hertz. See if you can figure out why—connecting an oscilloscope to the leads of the capacitor may help.

If you're stumped, note that when first powered on, there's a connection between the negative supply, the capacitor, and the bottom relay control pin. The capacitor quickly charges, and the relay is energized. After a brief moment for the mechanism to move, this interrupts the connection and opens the circuit. The relay doesn't turn off immediately because of the energy stored in the capacitor. Only after the capacitor sufficiently discharges does the electromagnet release, and the cycle restarts. We've constructed a very simple electromechanical oscillator.

Relays are among the oldest electronic components that can be harnessed for digital computation. They've been used to construct some of the earliest computers and telephony systems in the world. Still, the devices suffer from three significant limitations. First, their switching speed is limited by the time it takes for a mass to travel from one position to another. Second, the components are fairly power-hungry—even the most efficient miniature relays require around 30 mA just to hold in place. Above all, relays are prone to mechanical wear and fail after far fewer cycles than a transistor would. They keep working for decades in devices such as wall thermostats, where they're actuated only several times a day, but if we upped this to several times per second, the relay would probably be toast within a month.

Motors

Whereas solenoids produce single nudges, motors are designed to convert electric energy into continuous motion. This can be achieved in a couple of ways, but the most practical approach uses an arrangement of electromagnets to spin a metal shaft. The rotating shaft can then be coupled to gears, timing belts, hydraulic pumps, rack-and-pinion mechanisms, and much more.

The usual design of a small DC motor features a combination of permanent magnets and electromagnet coils. If the electromagnets are built into the rotor, as in Figure 9-8, the permanent magnets are integrated with the motor frame (the *stator*). It's also possible to reverse this arrangement.

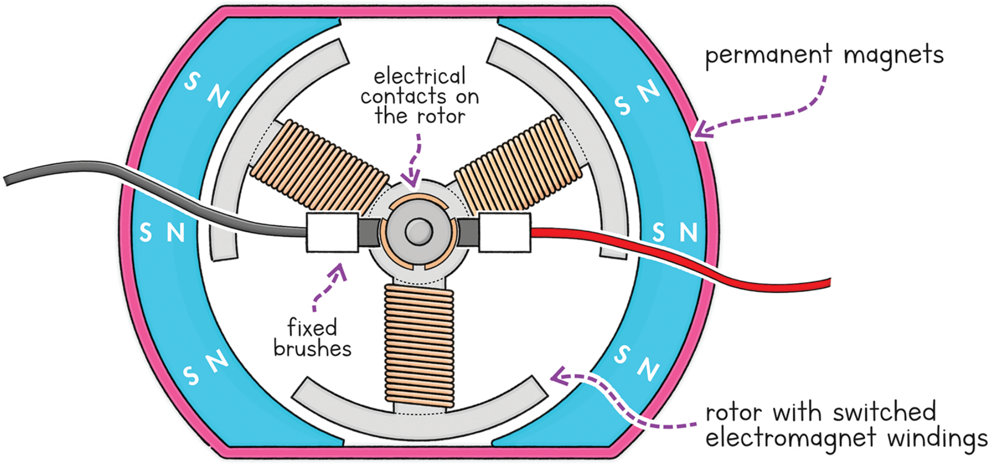

Figure 9-8: The structure of a brushed DC motor

Regardless of the layout, the idea is to energize the electromagnets in a way that pulls the rotor toward one of the stationary magnets. Then, just as the poles start to align, we break the circuit and send the current to another coil, making the rotor want to move to another configuration that requires turning farther along. We keep doing this to achieve the electromagnetic version of a dog forever chasing its own tail.

Brushed DC Motors

In a typical *brushed motor* like the one shown in Figure 9-8, the switching of coils on the rotor is accomplished by a mechanical device known as a *commutator*. The commutator consists of a circular pattern of copper contacts mounted on the rotor and two smooth, conductive brushes affixed to the motor frame. This energizes different coils depending on the angular position of the rotor assembly.

High-speed DC motors exhibit an interesting property: They don't use more power than necessary. From Chapter 6, recall that if the current flowing through an inductor (a coil) is changing, some of the energy is soaked up by a magnetic field. This manifests as a momentary pushback voltage that's proportional to the rate of change in the supplied current.

Importantly, the effect works both ways. If the magnetic field is changing for some external reason—for example, because the coil is spinning near a permanent magnet—electrons in the coil are pushed around and a voltage develops. This is what we observed with a manually actuated speaker, and it's what happens in a motor. The faster the rotor turns, the greater the rate of change of the magnetic field acting on the coils and thus the higher the pushback voltage. This phenomenon, known in motor design as *back electromotive force (back EMF)*, eventually reaches an equilibrium with the supply voltage. At that point, the rotor stops accelerating, and the device admits only as much current as needed to make up for friction and other energy losses. Yet, as soon as the rotor is connected to a load and starts spinning more slowly, back EMF is reduced, allowing more current to flow.

Because of this desirable property, there's some nuance to parsing the specs of brushed motors. Motors are advertised based on their RPM and torque, but you can never get the maximum of both. Without load, the motor spins fastest, but the available torque is zero—it can muster only enough energy to keep the rotor itself running at that speed. To get maximum torque, you must make back EMF disappear—and that happens only at stall. Between these extremes lies a roughly linear relationship among speed, torque, and current. For example, if you want to spin something at 80 percent of the motor's no-load RPM, you must settle for about 20 percent of the advertised stall torque. On the upside, you only need to supply roughly 20 percent of the stall current given in the motor spec.

Brushless DC Motors

Unlike their brushed counterparts, *brushless DC motors (BLDCs)* don't rely on mechanical commutation to switch current to the coils; they use electronic switching instead. Brushless designs are meant to overcome the disadvantages of brushed motors, such as relatively noisy operation and the tendency for brushes to wear. BLDCs can also be made lighter and more compact by making the control circuitry sense overload and kill power in a matter of milliseconds. In a brushed motor, the windings must be made bulky enough to survive stall currents indefinitely—or at least long enough for the operator to react or a fuse to blow.

Figure 9-9 shows some typical schematic symbols for motors. The left panel of the figure shows the usual symbol for a two-terminal motor. In a DC circuit, this motor would be brushed. There's no consistent symbol for BLDCs, but some common choices are shown on the right.

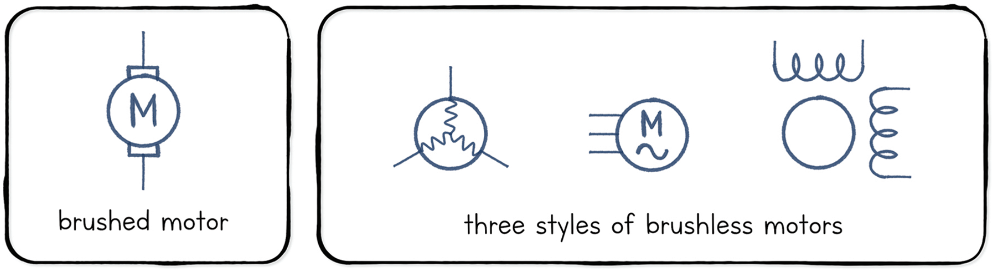

Figure 9-9: Common schematic symbols for motors

Constructing the control circuit for an electronically switched motor isn't trivial. Without knowing the position of the rotor, we might end up switching the coils too slowly or too quickly for the load. The former would result in jerky motion and excessive power consumption, and the latter would make the rotor lose steps and helplessly flail around.

Stepper motors were an early solution to this problem—or more precisely, an electronically simple solution to a related use case of actuating precision mechanisms at lower speeds. These motors have a high number of coils and magnet poles, producing relatively smooth motion without requiring perfectly timed coil switching. Further, they are meant to be operated at low RPM so that there is no need to adjust switching speed for loads that stay within the motor's spec.

Stepper motors usually run no faster than several hundred rotations per minute—about two orders of magnitude slower than a brushed motor of a comparable size. The benefits of back EMF largely disappear, so steppers are essentially constant-current devices, power-hungry even when not doing any useful work. Still, their simplicity, near silence, and precise angular positioning—due to the high number of coils—make them valuable in mains-powered equipment. For example, most hobby 3D printers use them for positioning.

A refinement of this concept is the modern brushless DC motor, which incorporates some form of position sensing. A reliable method uses a magnetic *Hall effect sensor* that detects small disturbances in the flow of current through a semiconductor substrate when a magnet passes nearby. A cheaper but more complex alternative is to monitor the current flowing through the motor coils; the amount of back EMF is proportional to rotor speed, so rudimentary coil timing corrections can be done by monitoring how successful the power supply is at pushing electrons through the device.

A BLDC motor has at least three terminals for the coils. Because these coils need to be energized in a specific, carefully timed sequence, an external digital controller is almost always necessary. However, as noted earlier, the motors themselves are lighter, quieter, and more durable than their brushed counterparts, so they find use in performance gear, ranging from high-end drones to premium battery-operated workshop tools.

Piezoelectric Crystals

Piezoelectric crystals are electronic components made from substances that change dimensions in response to an applied electric field. This might sound somewhat mystifying, so to understand this phenomenon, let's start by having another look at a more familiar component: the capacitor.

Consider a simple capacitor consisting of two metal plates and an inert insulator sandwiched in between. In Chapter 1, we noted that dissimilar charges attract, so in a charged capacitor, there must be an attractive force between the plates that's proportional to the applied voltage. This force is quite strong; it squishes the insulator layer, changing the dimensions of the device to some microscopic extent.

This relationship goes both ways. If we put a charged capacitor in a vise and apply external mechanical force, this brings the plates closer together, increasing capacitance (C). In Chapter 5, we established that capacitance is equal to charge (Q) divided by voltage (V); if C increases and Q remains the same, the voltage across the terminals of a mechanically squished capacitor must drop.

In real capacitors, this effect is usually negligible, but a special class of insulators—piezoelectrics—amplifies it through permanently aligned electric dipoles. These crystals, when squeezed, can produce sufficient voltage swings to give off sparks, so they're used as igniters in cigarette lighters and some gas appliances. In the same vein, although their dimensions change too little for general-purpose speakers, the crystals can generate high-intensity ultrasound or high-pitched beeps in smoke alarms. This is the operating principle of *piezoelectric transducers*.

Beyond fire-starting and noise-making, the most common use of these materials is as timing devices in modern electronics, from wall clocks to radios to computers. The resulting timing component is simply called a *crystal* or *quartz* (hence "quartz watch"). To explain the design, we can lean on the analogy of a playground swing. You'd be right to treat such analogies with suspicion, but in this instance, we're dealing with mechanical oscillations, so the comparison is apt.

A swing has a natural resonant frequency: If you give it a gentle push, it moves away a tiny bit and then comes back after a while. If you time the next push correctly, the amplitude of motion increases; if you time it incorrectly, it doesn't. The same is true for a piezoelectric crystal. We can set it in motion with electrical nudges, and if the frequency of the nudges matches the physics of the crystal, the amplitude of the oscillations will grow. It's possible to design a circuit that zeros in on maximum signal amplitude and thus naturally settles at the crystal's resonant frequency. This works because, as noted earlier, the phenomenon is bidirectional: Mechanical oscillations induced by small nudges eventually begin to manifest as measurable voltage swings.

Figure 9-10 shows the response of a crystal trimmed to a natural resonant frequency of 2 MHz. In the top portion of the plot, we can see that as the nudges from a signal generator approach the resonant frequency of the crystal, the number of electrons moving back and forth increases. This tells us that mechanical oscillations are increasing in amplitude, sympathetically pushing charges back and forth. At precisely 2 MHz, we notice a dip in impedance, meaning the admitted back-and-forth current is at its peak in relation to the constant-amplitude driving voltage.

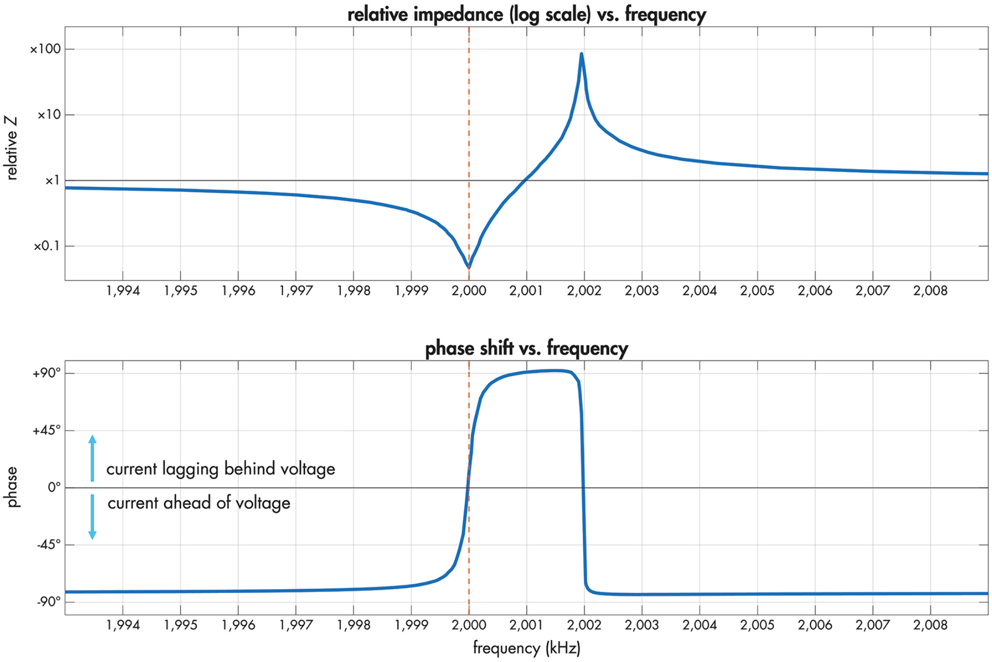

Figure 9-10: Real-world behavior of a 2 MHz crystal

In the same plot, we see that if the driving frequency is slightly higher than 2 MHz, the impedance briefly spikes. This is the anti-resonant frequency. The spike is most easily explained with the swing analogy: If you push too early—while the swing is still moving toward you—you're not helping at all. Instead, you're removing energy from the system and forcing the swing to stop.

One difference between a swing and a piezoelectric crystal is that in electronic circuits, we usually don't drive the crystal with discrete pulses. Instead, we use a continuous sine wave. At the resonant frequency, the driving signal and the motion of the crystal are perfectly in sync, so there's no phase shift between the input voltage and the resulting current (Figure 9-10, bottom half). At most other frequencies, the crystal isn't resonating, so we're looking at two pieces of metal with an insulator in between—a run-of-the-mill capacitor with its associated −90° phase lag. However, in the narrow range of frequencies—the "active braking" region between resonance and anti-resonance—the phase shift is flipped around. I'll spare you the math that explains it, since the phenomenon is relevant to just one circuit design: a particularly minimalistic oscillator we'll review in Chapter 23. If you're curious, though, exploring a discrete-time model should offer good insights.

In this day and age, it might seem odd that we're still using mechanical devices as timing sources. The reason is simple: They work exceptionally well. The specific resonant frequency of a crystal can be laser-trimmed during the manufacturing process with an accuracy measured in parts per million. Equally important, the frequency remains stable over time. Even better performance can be achieved by controlling for temperature, either with bundled sensors or with a miniature PCB-mounted "oven" housing the crystal.

Crystals are commonly available for frequencies between 2 and 60 MHz or so. A separate range of slow crystals, tuned to 32.768 kHz, serves applications such as low-power real-time clocks. Figure 9-11 shows symbols for piezoelectric elements and our next subject, microphones.

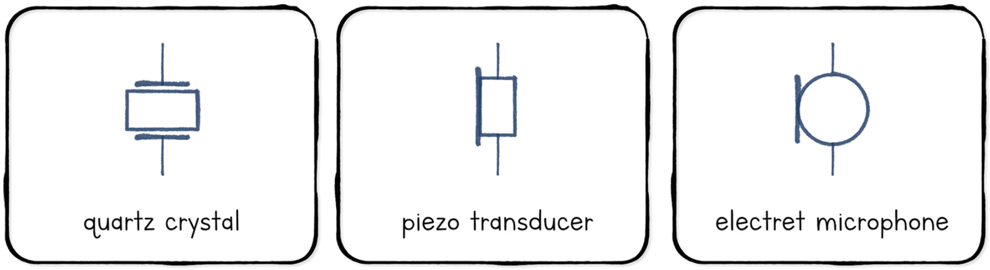

Figure 9-11: Symbols for piezoelectric elements and microphones

Apart from crystals, there's also a competing range of *micro-electromechanical system (MEMS)* devices that are manufactured by etching the resonant structure directly onto silicon, alongside the other components needed for an electronic oscillator. The advantage of this relatively new process is space savings, so MEMS oscillators can be found in smartphones and other highly miniaturized electronics where every square millimeter counts. Still, it's just a spin on the same theme, and when size isn't the overriding concern, piezoelectric crystals currently perform equally well or better at lower cost.

There's not much we can do with a crystal just yet, but as noted earlier, we'll revisit practical timing circuits in Chapter 23.

Microphones

In this chapter, we've discussed two ways of making noise: speakers and piezoelectric transducers. In principle, both can be used in reverse to capture sound waves, but they aren't particularly good at that. Conventional piezoelectric elements don't have sufficient movement range to capture low frequencies. A speaker performs better—the coil produces back EMF voltage if moved by air pressure across a wide range of frequencies—but the component is needlessly bulky and fragile for this use.

Instead, to record audio, we typically use a device called a *condenser* or *electret microphone*. The design consists of a permanently charged membrane—the electret—made of plastic in which an electrostatic charge is trapped during the manufacturing process. This membrane forms one plate of a capacitor; when it vibrates in response to sound, its motion relative to the other plate induces the movement of electrons. This resembles our earlier discussion of how squishing or stretching a ceramic capacitor produces swings in the voltage across its terminals.

In Chapter 2, we introduced field-effect transistors as voltage-controlled electronic valves in which the voltage applied to the gate moderates the flow of current between the two remaining terminals. In a condenser microphone, the second plate of the capacitor is connected to the FET gate. The induced microscopic currents change the gate voltage, so the current admitted through the transistor's main conduction channel changes accordingly.

As with piezoelectric crystals, the traditional design is facing some competition from the newfangled MEMS manufacturing process. MEMS microphones use active circuitry to detect the motion of a micromachined structure—typically a suspended membrane-shaped capacitor. As of this writing, MEMS microphones lack exceptional acoustic specs, but their small size and compatibility with industrial reflow soldering make them attractive for applications such as smartphones and hearing aids. For everyone else, electrets are still fine.

The electrical ratings of electret microphones are usually straightforward. The component can often tolerate a wide range of voltages across its terminals. For MEMS models, the active circuitry requires a more precise supply regulation, but the electrical side remains uncomplicated. The more important consideration is acoustic performance. Every microphone has a specific spatial pickup pattern, which can be plotted by spinning a sound source around the device under test. Like speakers, microphones also have a frequency response curve that should ideally remain flat across the audible spectrum, from 20 Hz to 15 kHz. Electrets typically come close to this ideal, but for hi-fi applications, it's always wise to check.

So far, we have covered the easy stuff. However, the most important acoustic parameters are sensitivity and noise performance, and to make sense of these specs, we regrettably need to talk about decibels.

Deciphering Decibels

In electronics and some adjacent disciplines, we often compare magnitudes. Differences in magnitude can become large, so we need shorthand to avoid spelling out that a quantity has changed by a factor of, say, 1,000,000.

A simple solution is to use exponents: We can write 1,000,000 as 10^6, take the base of the multiplier (10) as a given, and then use a keyword to indicate that we're talking only about the exponent of the multiplier. For example, we could call 10^6 an increase of 6 "clerts." In the same vein, a decrease by a factor of 10,000 could henceforth be known as −4 clerts. This notation also simplifies some circuit math. For example, if we have a long cable that attenuates signals by x clerts and another by y clerts, the combined clert effect is a simple sum, because $10^x \cdot 10^y = 10^{x+y}$.

Alas, the clert is made-up; the *decibel (dB)* is real. More precisely, the original pseudo-unit is the *bel*. The bel works much like our clert: +1 bel is a tenfold increase, while −2 bels is a decrease by a factor of 100. So far so good. But there's a twist: The bel was originally designed specifically for measuring power. Recall from Chapter 3 that power sometimes has a quadratic relationship to applied voltage. Specifically, for purely resistive loads, we have Joule's law: $P = V^2 / R$.

Seeing this, someone decided that 1 bel should always represent a hypothetical tenfold increase in power, even when applied to another base unit and even without resistive loads present. This means that if you're talking about watts, +1 bel is an increase of 10×, but if you're talking about volts, you're supposed to interpret +1 bel as an increase of $\sqrt{10}\times$ (circa 3.1622×). This ensures that squaring the voltage later yields the correct power via Joule's law. For other units, where the relationship to power is less obvious than for volts, figuring out which kind of bel you're supposed to be using can be a major headache.

The weirdness doesn't end there. At some point, we decided that the bel is too big to use—but only slightly so. Instead of putting up with the occasional fraction or switching to another base for the exponent, we divided the bel into 10 steps, hence *deci*bel. In effect, we started raising 10 to fractional powers, producing irrational per-decibel multipliers. For power, 1 decibel is a multiplier of $10^{1/10} \approx 1.258925$; for voltage, it's $\sqrt{10}^{1/10} = 10^{1/20} \approx 1.122018$. In all cases, 0 dB means a multiplier of 1—signifying no change.

The original bel is gone; we now use decibels exclusively.

Quantifying Microphone Sensitivity and Noise

Knowing the meaning of the decibel isn't enough. When someone uses the term, we also need to know the base unit and the reference point they have in mind. In acoustics, decibels are used to measure sound level; 0 dB is the loudness of a 1 kHz acoustic wave that's calibrated to exert 20 micropascals of pressure. This is roughly the threshold of human hearing.

From this 0 dB origin, we can derive several acoustic scales. One popular choice measures the absolute sound pressure level with no regard to frequency, often denoted as dB SPL. Other scales can be weighted to mimic human hearing, typically peaking at 3 kHz and then tapering sharply below 200 Hz and above 10 kHz. The hearing-weighted loudness scale is typically denoted as dBA.

In a microphone spec sheet, the sensitivity might be given as −45 dB. This isn't an acoustic decibel: Here, the base unit is volts (dBV). The zero point on the scale is a theoretical microphone that would output 1 V_{RMS} in response to a reference acoustic wave. If your head is spinning already, it gets worse: The reference sound isn't 0 dBA but a 94 dB squeal at 1 kHz, roughly the loudness of a gas-powered lawnmower. If a microphone has a sensitivity of −45 dBV, it means the lawnmower-induced swing is 45 (voltage) decibels lower: $1\ V_{RMS} \cdot 10^{-45/20} \approx 6\ mV_{RMS}$.

Needless to say, you don't need to memorize all this detail. I'm just giving you a quick overview in case you need to figure out how the numbers in the spec translate into the voltages on the oscilloscope screen.

Sensitivity aside, another important characteristic of a microphone is its noise floor. There are two ways to measure it, but the more common approach is the signal-to-noise ratio (SNR). If the SNR is given as 60 dB, it means there's 60 decibels of headroom between the noise floor and the aforementioned 94 dB lawnmower reference point. In more logical terms, this means that the internal "hiss" contributed by the electronics of the microphone is equivalent to a 94 dB − 60 dB = 34 dBA source of static noise placed in the vicinity of an ideal microphone. This is known as *equivalent input noise (EIN)* and is a more reasonable, if less common, way to report microphone performance. A positive EIN means the hiss is louder than the threshold of human hearing; 34 dBA is enough to drown out whispering or nature sounds.

An average electret microphone might have −45 dBV sensitivity and an SNR of 60 dBA. Spending $2 more will get you a top-of-the-line module with −25 dBV sensitivity and an SNR of 80 dBA. For low-fidelity applications, such as gaming headsets, this doesn't matter. For everything else, a good microphone makes a dramatic difference. If you're curious, you can listen to a quick demo at the following URL: *http://lcamtuf.coredump.cx/slc/mic-compare*.

NOTE *The astute reader may have picked up that I'm not a fan of decibels. Although exponent notation is useful, this particular flavor of it is a source of needless pain. Suffixes help, but their landscape is chaotic too. In communications, for example, dBm references 1 milliwatt, but dBµ denotes 1 microvolt; in other disciplines, the same acronyms may mean something else. We can't escape decibels entirely, but apart from component specs, we won't be leaning on them throughout the book.*

Circuit Protection Components

Circuit protection devices (Figure 9-12) are electromechanical only in the loosest sense: They're meant to "pop" before a more expensive portion of the circuit catches fire. Still, we need a place to discuss them briefly. If you squint your eyes, the blowing of a fuse can be seen as an act of electromechanical sacrifice.

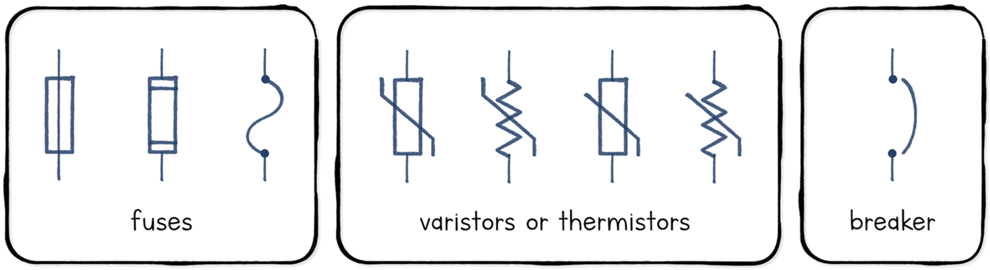

Figure 9-12: Symbols for circuit protection devices

Fundamentally, a *fuse* is a current-protection device. It serves two functions. First, it shields equipment from transient overloads (for example, an unexpected motor stall). Equally important, it protects surroundings from electrical malfunctions caused by internal short circuits. An unchecked short could ignite flammable plastics if the device overheats.

Basic fuses consist of a strip of (slightly) resistive material and essentially exploit the squared term in one form of Joule's law: $P = I^2R$. They generate manageable heat up to a certain current (I), but the material melts quickly—interrupting the circuit—if that threshold is exceeded.

A more user-friendly variant of this approach is the *breaker*, a latching switch that releases automatically if the current exceeds a certain threshold. Small pushbutton breakers, found on some appliances and power strips, are usually thermal and work like thermostats. Larger breakers use a solenoid that pulls a linkage if the current flowing through the breaker is higher than it's supposed to be.

Simple fuses are necessarily somewhat approximate. They trip quickly if the current is completely out of whack, but they might tolerate as much as 150 percent of their rated current for minutes or even hours. This can be a curse or a blessing. Some fuses are deliberately designed to have a higher thermal mass to withstand current surges when, for example, a large motor is spooling up. Reputable manufacturers provide detailed response curves for their products, so you can pick the right fuse for your needs.

In addition to traditional fuses and breakers, we have *resettable fuses*, also known as *polyfuses* or (more generically) *thermistors*. Once again, these devices are designed to be heated by the flow of current through a resistive material. A *positive temperature coefficient (PTC)* thermistor consists of conductive particles embedded in a nonconductive substrate that expands with temperature and eventually pulls the conductive particles apart. This typically reduces—but does not interrupt—current flow until the fuse is allowed to cool.

The last circuit protection devices on our list are *metal-oxide varistors (MOVs)*. Unlike fuses, they're meant to protect the device not against current but against voltage spikes that may follow lightning strikes, downed power lines, and so on. MOVs are made of ceramic materials with a carefully controlled breakdown voltage. The components usually straddle mains supply lines and are supposed to act as a crowbar if the voltage gets too high. If tripped, they tend to explode rather spectacularly. The hope is just that the MOV's fiery death gives enough time for the fault to clear or a standard fuse to take notice and trip.

Leaving passives and electromechanical devices in the rearview mirror, we're now ready to enter the kingdom of semiconducting devices.