


ELECTRONICS, CLEARLY

A motor-first guide to understanding, modifying and troubleshooting electrical machines.

A complete route through the motor



Close the gap to complete the route. Open either wire to interrupt it.

Concept and Instructional Design by Creaky
Written by Jim (the one with all the issues)

Free to download and print from creakyscoots.com

17 chapters + The Bench Book

Before you start

This book begins with a battery and a motor. Each new component appears because the previous circuit leaves a problem to solve. The goal is to understand what a machine is doing well enough to choose a useful measurement.

Read the chapters in order the first time. The Bench Book collects the units, equations, symbols and diagnostic patterns for later lookup.

Examples and real machines

Motor M and the numerical fault cases are constructed teaching models. They are not measurements of an identified commercial product. Changing the component, temperature, load or controller can change the answer.

The DeWalt parallel/series scooter account in Chapters 4-6 is Creaky's workshop history. The switch failure is real; its exact cause was not instrumented. The original diagrams in this edition explain circuit principles, not the complete wiring or protection of that build.

Physical practice

You can follow every lesson without buying equipment. For hands-on practice, use a purpose-built educational kit with its matched, protected supply. Disconnect power before changing connections and guard moving parts. Scooter batteries and controller capacitors can release substantial energy even when called low voltage.

The scooter, charger and controller cases are supplied diagnostic examples. They are not instructions to open energized mains equipment or bypass protection. Real testing depends on the exact machine, instrument ratings, suitable probes and a controlled setup.

Using this edition

This free CSS edition is formatted for US Letter paper. Print at actual size on Letter, or fit to the printable area on A4. Black-and-white printing preserves all diagram information. Use the document outline to navigate the chapters.

Written source links appear at the back.

Manufacturer examples explain particular mechanisms; they do not certify an unrelated component or modification.

Contents

Electronics, Clearly

1. Make something move
2. Four things happening in the wires
3. The motor that apparently breaks the equation
4. More than one road
5. Why the switch melted
6. The current doesn't stop politely
7. The terrible speed controller
8. Turn it on and off really fast
9. Little circuits running big circuits
10. Components are materials behaving predictably
11. Why brushes exist, and why we got rid of them
12. Three wires, three magnetic fields
13. Inside a real scooter controller
14. What happens when you change the rules
15. Failure has a small vocabulary
16. Stop guessing
17. Here's something you've never seen before

The Bench Book

Written sources

CHAPTER 1

1. Make something move

A small motor sits on the table. A short metal shaft sticks out of one end. That is the part that turns. At the other end are two metal connection points. Beside it sits a battery, with one connection marked + for positive and the other marked – for negative.

Those marks identify the two battery connections; they are not an instruction to add or subtract. The drawing below labels the source 6 V. V is the unit symbol for volts; what a volt measures comes in Chapter 2. We will learn what makes those connections different. First, we want the motor to move.

You do not need these parts in front of you. We will run the experiments on the page and supply the results. Our example motor, Motor M, is a teaching model rather than a product with guaranteed readings.

For physical practice, use a purpose-built small-motor educational kit with its matched, protected power source and instructions. Keep moving parts guarded and disconnect power before changing wires. The battery in the drawing explains a circuit; it is not an instruction to experiment with a scooter, car or power-tool battery.

One connection is not enough

Connect the battery's positive connection to either one of the motor's connections. Leave the remaining motor connection alone. The result: the motor does not start.

You might expect one wire to be enough because it has reached the motor. But this motor cannot keep running from that one connection. Add a second wire between the unused motor connection and the battery's negative connection. Now the shaft turns.

The second wire completed an electrical route through the motor between the battery's two connections. That arrangement is a circuit. The motor needs the return connection as well as the first one. In this simple circuit, removing either wire stops the battery from continuously driving it. [1]

The battery is the source: it supplies energy. Energy is what can be transferred to make a change, such as moving something or heating it. The motor is the device, also called the electrical load: it uses that energy to do a job. The exposed connection points are terminals.


These are names for things already on the table, not a vocabulary test. When someone says "check the motor terminals," they mean the two places where you connected the wires.

Give the route a switch

Disconnect the source. Put a switch into one of the wires, so the route now has to pass through the switch. Inside a simple switch, moving metal pieces either touch or separate. Those touching pieces are its contacts.

With the contacts touching, the motor runs. With the contacts separated, the motor loses its supply and coasts to a stop. The switch has not created the energy. It has allowed or interrupted the connection.

A complete route through the motor



Close the gap to complete the route. Open either wire to interrupt it.

Figure 1. A source, an open switch and Motor M. Educational source protection is not drawn.

Straight lines represent wires. The long and short parallel marks represent the source; + identifies its positive side and – its negative side. The circle containing M represents the motor. The small circles and tilted line show separated switch contacts. This is a connection map, not a physical parts layout.

The switch is open when its contacts are separated and closed when they touch. You might expect “open” to let electricity through, like an open door. Think of closing a gap in a bridge instead: close the gap and the route is complete.

Move the switch into the other wire, with the source disconnected while you make the change. Will it still control the motor? Yes. In this one-route circuit, either interruption does the same job.

Opening the switch is not the same as applying a brake. The shaft was moving; it can continue moving briefly after the source stops driving it. “Switched off” describes the electrical command, not a promise that everything is already motionless.

Backwards is not always broken

Disconnect the source and let the shaft stop. Now exchange the two wires at the motor: the wire that reached one motor terminal goes to the other, and vice versa. Reconnect the source and close the switch.

The motor turns the other way. You might have expected it to burn out because its connections are backwards. Our chosen motor is a simple kind with fixed magnets and sliding electrical contacts inside. Reversing its supply reverses its turning direction. Its internal workings come later. [2]

The arrangement of positive and negative connections is called polarity. We have reversed the motor’s polarity without changing the battery itself.

This is not a trick to try on every device with two wires. A motor assembly with a circuit board inside may be damaged by reversed power. Even an added protective component can change what happens. Our answer belongs to the bare motor in this example, not everything sold as a motor.

It looks connected. It is not.

We restore the original wiring. This time the motor does not turn at all. It worked before we moved the wires. Has it failed?

Not enough evidence yet. “It does not turn” is the symptom, the behavior we can observe. “The motor failed” is a proposed diagnosis, an explanation of the cause. They are not the same statement.

Look closely at the moved connection. The clip is gripping the wire's plastic covering instead of the bare metal terminal.

The metal inside an ordinary wire is a conductor, a material that allows electrical flow readily. Its plastic jacket is insulation, material intended to prevent unwanted electrical contact. A clip can hold the plastic firmly while making no useful electrical connection.

Disconnect the source, move the clip onto the terminal's metal, then reconnect. The motor turns again. The repair was a connection, not a replacement motor.

You might have blamed the motor because it was the part refusing to move. The failed connection explains the same symptom. The successful repair supports that explanation for this example.

A metal motor case is not automatically a substitute terminal, either. Connect to the designated connection points, not whichever metal happens to be nearest. A part's mechanical job and its electrical job can be different.

Same source. Harder job.

So far, the motor has only had to spin its own shaft. Now imagine it driving a small wheel inside a guarded test fixture. A fixture is simply an arrangement that holds equipment in place for a particular test. Ours includes a little adjustable brake pressing against the wheel.

This is a supplied experiment, not an instruction to grip a running shaft or tire. The fixture applies the opposing force; fingers do not.

We begin with the brake released. The wheel spins freely. Then the fixture presses the brake lightly against it. The wheel slows but keeps turning. Release the brake, and it returns toward its earlier speed.

The source has stayed the same. The wires have stayed the same. The switch is still closed. What changed was the mechanical job.

The motor is the electrical load on the source. The brake is a mechanical load on the motor: something opposing its motion. The word "load" describes a relationship, so always ask what is loading what.

On a scooter, that mechanical job includes moving the rider and overcoming the forces opposing travel. In our small experiment, the brake gives us a controllable way to make the job harder without changing the electrical connections.

You might expect the same source to make the motor turn at exactly the same speed regardless of its job. Our result says otherwise. A motor connected directly to a source is not automatically a machine that holds one speed against any load. [2]

This is also where speed and effort separate. Turn an easy screw quickly; then try a tighter screw. Your hand may turn more slowly while pushing harder. A motor's turning effort has a name: torque. The word tells us about the twist it produces, not how fast the shaft goes around. [3]

Increase the brake far enough in our paper experiment and the shaft eventually stops. That condition is a stall: the source is still connected, but the motor is prevented from turning. We disconnect it rather than leave it stalled.

You might think a stopped motor must be using nothing. A stalled motor can still take energy from the source and become dangerously hot. Opening the switch and holding the shaft still are not equivalent: one interrupts the supply, while the other leaves it connected to a motor that cannot turn. No movement does not prove no electrical activity. [2]

A missing route and a wrong route

The loose clip interrupted the motor's route. We call that an open circuit. Now imagine a different mistake: an extra wire connects the source's positive and negative terminals directly, bypassing the motor.

That is a short circuit: an unintended route with very little opposition to electrical flow. "Short" does not mean the wire is physically short. A long wire can create the same kind of fault.

Would that be a useful shortcut that makes the motor faster? No. It bypasses the device we wanted to operate. The extra route can allow far more electrical flow than intended, heating the wires and source and causing damage. A protected source may shut down or reduce its output instead. The actual outcome depends on the source and the fault. [2]

We keep this experiment on paper. Do not join battery terminals to see what happens. The lesson is about recognizing a route, not earning a burn.

Notice why "it stopped" cannot name the fault. Our motor could stop because a required connection was missing, because a mechanical load held it still, or because an unwanted connection made the source's protection act. The visible result can be similar while the underlying problem is very different.

The question that comes next

Return to the working circuit. Source connected. Switch closed. Wheel turning. The guarded brake presses against the wheel again, and the motor slows.

We can see the mechanical change, but we cannot see what is happening inside the wires. We need a way to distinguish what the source is offering from what the motor is actually doing.

That is the next chapter's job. We will put a meter into this same story, supply every reading, and explain the marks on its display as we encounter them. The source did not change. The job did. Now we find out what changed electrically.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [1], [2], [3]. Full titles and links appear in the source list.

CHAPTER 2

2. Four things happening in the wires

Motor M is turning the little wheel from Chapter 1. The guarded brake presses against it. The wheel slows. Nothing about the source or the wiring has been changed.

We can see what happened mechanically. Now we want to see what changed electrically. Put a measuring instrument into the story, and the motor gives us something more useful than “it sounds different.” It gives us numbers.

You do not need the instrument. Every reading below is supplied. Motor M and its readings are a constructed teaching example, not a particular product or a record of physical tests. A different real motor is not faulty merely because it gives different numbers. Physical practice still belongs inside the matched, protected educational kit described in Chapter 1.

Start with the number, not the dictionary

Our instrument is a multimeter: one meter with several measuring jobs. Its two flexible test wires are called leads. Their pointed ends are probes. The dial selects a job, the sockets select the connections inside the meter, and the probe positions select the thing being measured. All three matter.

For our first reading, the black lead is in the socket marked COM, short for common. This is the reference connection: the point we compare the other probe with. The red lead is in the socket marked V. Set the meter to DC voltage. DC means direct current; in this example, the source keeps its positive and negative connections the same way around. On many meters the DC mark is a solid horizontal line above a broken one. The actual meter's manual identifies its settings. [5]

With the motor switch open, put the black probe on source negative and the red probe on source positive. Our display reads 6.00 V. Say “six volts.” The decimal point separates whole volts from fractions; the trailing zeros do not make this more voltage than 6 V. They show how finely the display is reporting, not that its answer is perfectly accurate.

A volt is the unit used to measure voltage. A unit is an agreed measuring amount, like an inch for length or a second for time. Voltage is an electrical difference between two places. “Electrical push” is a useful first comparison: pressure can exist across a closed valve even though no water is flowing. Likewise, our source has six volts between its terminals even while the motor's route is open.

You might expect zero because the motor is off. The answer is still 6.00 V across the source. Having an electrical difference available and having current flowing through the motor are different conditions. The meter takes a tiny current to make its voltage measurement; we leave that out of our simple motor-path numbers.

A voltmeter always compares two places

Leave the source alone and exchange the probe positions. The display reads –6.00 V. That minus sign means the red probe is now on the electrically lower side of the comparison. Nothing has happened to the source. We reversed our measuring direction, not the motor's wires. [5]

Touch both probes to the same terminal. The result is approximately 0 V. A place has no voltage difference from itself. Zero on a meter is only meaningful when we also know what two places it compared.

Now close the switch and measure across the motor's two terminals. Our source holds its output steady, and the model ignores the tiny losses in healthy wires and contacts. The motor receives 6.0 V. "Across the motor" means one probe at each of its two terminals. We have not cut a wire or put the meter in the motor's main route.

Keep the motor running and compare the two sides of the closed switch instead. The answer is close to 0 V. You might expect six volts because the switch is carrying the motor's supply. But the meter is comparing the entrance with the exit. A healthy closed switch loses very little voltage between them.

That is your first hint that a useful measurement needs an address. Six volts across the source and almost zero across its closed switch can both be correct at the same time.

Now count what passes

The other measuring job we need is current. Current tells us how much electric charge passes a place each second. "Rate" means an amount for each interval of time. A traffic counter asks how many cars pass in a minute; it is not the same instrument as a speedometer.

But what is charge? It is a property of matter that makes electrical attraction and repulsion possible. Electrons are tiny charged particles already present in a metal wire. Some can move through the metal. When we talk about current here, we are counting the amount of charge passing, not trying to follow one particular electron from battery to motor. [1]


Charge has its own measuring unit, the coulomb, written C. Think of it as a standard-sized counting bundle of charge, not a new component inside the wire. One coulomb passing a point in one second is one ampere, usually shortened to one amp. The unit symbol is A. "Amperage" is another word people use for current. [1, 4]

This makes a meter reading of 0.05 A less mysterious. It means five hundredths of an amp, not five amps. The same reading can be written 50 mA, fifty milliamps. Milli, abbreviated with a small m, means one thousandth. There are one thousand milliamps in one amp. The smaller unit gives the same flow a larger-looking number, just as inches give a larger count than feet for the same length. [4]


To measure flow, enter the route

A voltage measurement compared two places. An ordinary multimeter measuring current must instead become part of a suitable current path. We disconnect the source, open one connection in the motor circuit, and fit the correctly configured meter into that gap. Its current socket and current setting let the path continue through the instrument. Putting things one after another in the same route is called series; that is all the word means in this drawing. [6]

Voltage: across



Current: in series



Two separate setups. Use the correct meter sockets, mode and range.

Figure 2. Voltage is measured across the motor. Current is measured by opening its route and inserting the correctly configured meter.

M marks the motor. V marks the voltage measurement across the motor. A marks the current measurement in the motor's route. Solid dots mark joined wires. The two panels show separate measurement setups, not instructions to connect two meters at once.

Never place a meter connected for current straight across a battery. Its current-measuring path offers very little resistance, so that connection can create the short circuit from Chapter 1. It does not safely reveal "how many amps the battery has." The correct socket, fuse, range and allowed measuring time depend on the meter. After a current test, disconnect power and return the red lead to the voltage socket before measuring voltage again. [5, 6]

A real meter can also slightly change the circuit it measures. For now, our page experiment leaves that disturbance out. We are learning which question each connection answers, not certifying a particular instrument for a scooter battery.

The motor changes its appetite

Reconnect the page circuit with the meter in its current path. Keep the source at 6.0 V. Release the guarded brake, let the motor settle into its free-running speed, and read 0.05 A.

Apply our specified mechanical load with the fixture. The shaft slows. The current rises to 0.25 A, a quarter of an amp, or 250 mA. Release the brake and the current returns toward 0.05 A as the shaft returns toward its earlier speed.

You might have expected the source to deliver the same current because it is still the same six-volt source. The answer is different because the motor's operating condition changed. While it turns more slowly against the brake, more current passes through its winding, the coil of wire inside the motor, and it develops more turning effort. That effort is the torque we met in Chapter 1. This relationship belongs to our motor; it is not a rule that every electronic device draws more current whenever its job becomes harder. [2, 3]

Here is the complete record we will carry into the next chapter:

Motor M condition	Motor voltage	Motor-path current
Switch open; shaft stopped	0 V	0 A

Motor M condition	Motor voltage	Motor-path current
Free-running; switch closed	6.0 V	0.05 A
Specified load; shaft turning	6.0 V	0.25 A
Energized stall; settled model only	6.0 V	0.50 A

Constructed values. The source remains at 6.0 V in every row; with the switch open, that voltage does not appear across the stopped motor.

The energized-stall row is a paper result, not an instruction to hold a shaft. “Settled” means after the brief electrical change at startup. It does not mean safe to leave it that way. These numbers assume the source still maintains six volts and the winding has not heated enough to change its resistance.

You might expect the stalled motor to take no current because it is doing no visible work. It takes 0.50 A, half an amp: ten times our free-running reading. A powered stall is not an open switch. The source remains connected even though the shaft cannot move. [2]

In this simple single-route circuit, a current measurement in the return wire gives the same 0.05 A during free running. The motor has not eaten some of the charge and sent less back. It transfers energy into motion and heat. Later, a controller will create additional paths, and its battery current and motor current will need separate attention.

More available amps is not more forced amps

Consider a source marked 6 V, 2 A. Here its rating means it can maintain six volts while supplying up to two amps under its specified conditions. Two amps is a capability, not a command to every device attached to it.

Connect Motor M in its free-running condition. Its current is still 0.05 A, not 2 A. Substitute a source capable of 5 A that maintains the same six volts, with the same motor and load. The answer remains 0.05 A in our model. The first source was already adequate. A bigger ceiling does not make the motor reach it. [2]

This is a separate question from increasing voltage. On paper, turn an adjustable source from three volts to six while the same bare motor has little mechanical load. The motor settles at a faster speed. It is not guaranteed to turn exactly twice as fast, and we have not specified an exact shaft speed for this model. We changed the applied voltage, not merely the source’s ability to supply current. The next chapter explains the difference. [3]

On the battery side, a stated maximum discharge current means how much current the battery is permitted to supply under stated conditions. It is not how much it forces through the motor. If a battery label gives capacity in amp-hours (Ah), the h means hours. Amp-hours describe an amount of charge available over time, not a maximum current. We will work through battery capacity and connections in Chapter 4.

On the motor side, a rated current describes an operating condition or limit. It is not necessarily enforced inside the motor. The winding can take more current than it can survive if the connected source and conditions permit it. A protective circuit may intervene; the printed label cannot. [2]

When a source cannot supply the required current, its voltage may fall or its protection may stop the output. Therefore the voltage written on a battery is not proof of the voltage actually reaching a working motor. Keep the source label and the measurement separate.

Resistance is not the brake pressing harder

Disconnect Motor M from the source and the rest of the circuit. Let it stop. Verify in voltage mode that no voltage remains at its terminals. Now the resistance measurement belongs to the isolated motor, not to a powered circuit.

Resistance is opposition to electrical current in a material or component. Its unit is the ohm, written Ω . That mark is the Greek letter omega being used as a unit symbol, not an advanced calculation. A resistor that passes one amp when one volt is across it has one ohm of resistance. A resistor is simply a part deliberately made to provide a chosen resistance. We will use one in Chapter 3. [4]

For our model, the corrected resistance reading at the stopped motor terminals is 12 Ω , twelve ohms. We treat this as the resistance of its winding, the wire wound inside it, and leave the extra effects of its sliding contacts out of the first calculations.

The meter performs this job by applying a small test signal of its own and observing the response. That is why the source must be disconnected. Other attached components can also provide extra routes and change the answer. A resistance measurement is not a second way to measure a running motor's current. [7]

You might expect a slower, more heavily loaded motor to have “more resistance.” There are two meanings tangled together. The brake mechanically resists rotation. The winding electrically resists current. Pressing the brake does not insert a bigger electrical resistor into the copper.

For the comparisons in our record, the winding resistance stays 12 Ω while current changes with the motor's movement. Real windings change resistance as they warm. Real brushed-motor readings can also change with shaft position, contact condition and the resistance of the test leads. We have held those complications still so the main effect is visible.

A resistance reading also needs its mode understood. With our resistance-test probes separated, the meter shows OL, meaning beyond its measuring range in this mode. Touching clean probe tips together gives a reading near zero instead. OL is not zero ohms. Separated probes offer no complete test path; touching probes offer very little opposition. The meter's manual explains its particular display. [7]

And twelve ohms alone is not a certificate that a motor is healthy. It answers one question about one electrical path while the motor is stopped. It does not tell us that the shaft turns freely or that the motor behaves correctly while running.

Watts tell us what is happening to energy

Current is not the same as power. The motor could have the same current at two different applied voltages and receive different amounts of energy each second. We need both parts of the description.

Chapter 1 introduced energy as what can be transferred to make something happen, such as lifting a weight or warming metal. Power tells us how quickly that energy is being transferred. Lifting the same weight to the same shelf in two seconds instead of four takes the same lifting energy, but twice the average power during the lift.

The energy unit is the joule, written J. The power unit is the watt, written W. One watt means one joule transferred each second. Ten watts sustained for three seconds transfers thirty joules. Did you get thirty? That is ten for the first second, ten for the second, and ten for the third. Watts describe the pace; joules describe the amount transferred. [4, 8]

Now the earlier description of voltage can become more precise. One volt means one joule of energy per coulomb of charge. “Per” means for each. Current tells us how many of those coulombs pass each

second. Voltage tells us the energy available for each coulomb. Together they tell us the energy transferred each second. [4, 8]

For Motor M at six volts, the supplied electrical-input powers are 0.30 W free-running, 1.50 W with our specified load, and 3.0 W at the modeled powered stall. Chapter 3 shows every step of the arithmetic. These are calculated teaching values, not readings from a physical power meter.

You might expect the three-watt condition to give the most useful motion. The shaft is stopped. Its mechanical output power is zero. The electrical input instead becomes heat in our stalled-motor model. A stationary motor can produce turning effort without delivering motion, just as pushing hard against something that will not move does not make it travel.

The loaded motor's 1.50 W input is not a promise of 1.50 W at its shaft, either. Some energy becomes unwanted heat before the useful work gets done. And a watt reading is not a temperature: how hot the motor gets also depends on how long it runs and how well it sheds heat. [8]

Same motor, now with evidence

Here is a supplied fault record. The motor has stopped. The source measures 6.0 V across its terminals. Across the motor we find 0 V, and in its main path we find 0 A. Across the switch contacts, with the switch supposedly closed, we find 6.0 V.

You might blame the motor because it is the part that stopped. But remember the working switch: its entrance and exit were at almost the same voltage. This switch has the full source voltage across its two sides instead. In our one-route circuit, that points to an interruption at the switch rather than a powered stall.

For this teaching case, inspection with the source disconnected confirms that the switch contacts are not touching even though the lever is in its ON position. Replacing the faulty switch restores 6.0 V across the running motor, approximately 0 V across the closed switch, and 0.05 A through the free-running circuit. The motor was not the failed part. These readings and the successful repair support the diagnosis together.

That conclusion has boundaries. A different circuit, an additional broken connection, or a probe touching the wrong place can change the interpretation. We named the circuit, the operating condition and both measurement points because the bare number was not enough.

We now have the record Chapter 1 was missing: the source stayed at six volts, the load changed, and the current changed. The winding's modeled twelve-ohm resistance did not. Next we will make those numbers explain one another, and discover why a spinning motor cannot be treated as only a piece of resistive wire.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [1], [2], [3], [4], [5], [6], [7], [8]. Full titles and links appear in the source list.

CHAPTER 3

3. The motor that apparently breaks the equation

Motor M has already done something useful for us: it refused to behave like a simple number on a label. At 6.0 volts it drew 0.05 amp while spinning freely, 0.25 amp when we made it work, and our paper model predicts 0.50 amp if it is stationary and energized. The winding itself did not change from twelve ohms to something else each time.

This chapter gives those readings a mathematical skeleton. The equations are not commandments. They are short ways of writing relationships we can already observe. When one seems to give the wrong answer, we will not hide the wrong answer. We will use it to find the missing behavior.

Every numerical reading is supplied. Motor M is a constructed teaching model. The stationary energized case remains a calculation, not an instruction to hold a motor stalled.

Start with something that really is mostly resistance

Before asking a motor to obey an equation, give the equation an easier job. Imagine a fixed resistor marked 120 Ω . A resistor is a component deliberately made to provide a known amount of electrical resistance. Unlike Motor M, it has no shaft, magnets or useful motion to complicate the first comparison.

Across our ideal 6.0 V source, the 120 Ω resistor carries 0.050 A, or 50 mA. Change only the source to 3.0 V and the current becomes 0.025 A, or 25 mA. Half the electrical push through the same resistance gives half the current.

You do not have to discover those answers with a meter. They are supplied so we can concentrate on the relationship.

The first equation, with nobody left outside the gate

The relationship is called Ohm's law. We will use I for current, V for voltage and R for resistance. In the equation below, = means has the same value as and $\div$ means divide:

$$I = V \div R$$

I = current, measured in amperes or amps; unit symbol A.

V = voltage, measured in volts; unit symbol V.

R = resistance, measured in ohms; unit symbol Ω .

$\times$ means multiply. Ω is the Greek letter omega, used as the symbol for the unit ohm.

The same letter V is used for the quantity voltage and as the unit symbol for volt. Context tells us which job it is doing.

$$I = 6.0 \div 120 = 0.050 \text{ A}$$

$$I = 3.0 \div 120 = 0.025 \text{ A}$$

You might have expected the current to fall by six because the voltage fell by three volts. The actual current falls by half because the voltage itself fell by half: from six volts to three.

The same relationship can answer different questions

Sometimes current is unknown. Sometimes we know current and resistance and want voltage. Sometimes we know voltage and current and want resistance. We do not need three laws; we rearrange the same relationship.

$$V = I \times R$$

If 0.050 A flows through 120 Ω , the voltage difference is $0.050 \times 120 = 6.0$ V.

$$R = V \div I$$

If a device has 6.0 V across it and 0.050 A through it, the simple resistive relationship gives $6.0 \div 0.050 = 120$ Ω .

These are tools, not an algebra exam. Pick the form that leaves the thing you want by itself, then attach the unit. A naked 120 is unfinished; 120 Ω tells us what the number means.

Now let the motor embarrass our first model

Motor M measured 12 Ω with all external power removed and its shaft stopped. Put that resistance into the relationship with a 6.0 V source:

$$I = 6.0 \div 12 = 0.50 \text{ A}$$

But Chapter 2 gave us the free-running current: 0.05 A.

You might think you used Ohm's law incorrectly. You did not. The arithmetic is correct for a situation in which all six volts are driving current through twelve ohms of resistance. The running motor is not that situation.

A correct equation can still answer the wrong physical question. Before distrusting the arithmetic, ask what the model left out.

A motor is quietly becoming a generator

When Motor M turns, conductors in its windings move through a magnetic field. That motion generates a voltage. We will build the magnetic explanation carefully in Chapter 6. For now we need only the consequence: the generated voltage points against the voltage making the motor turn.

It is called back electromotive force, usually shortened to back-EMF. Despite the word force, back-EMF is a voltage.

Our free-running Motor M is supplied with 6.0 V and generates 5.4 V of back-EMF. Those voltages oppose each other, so only 0.6 V remains to drive current through the winding resistance.

$$6.0 \text{ V} - 5.4 \text{ V} = 0.6 \text{ V}$$

$$I = 0.6 \div 12 = 0.05 \text{ A}$$

There is our supplied free-running current. The motor did not violate Ohm's law. We had originally applied the full source voltage to the resistive part of a system that was generating an opposing voltage of its own.

Put the missing behavior into the equation

Let E stand for the generated back-EMF. We can write the motor relationship more honestly as:

$$I = (V - E) \div R$$

I = motor current. V = applied motor-terminal voltage. E = generated back-EMF. R = winding resistance. The parentheses mean do the subtraction first.

Free-running:

$$I = (6.0 - 5.4) \div 12 = 0.05 \text{ A}$$


Under our specified mechanical load, Motor M slows and its back-EMF falls to 3.0 V:

$$I = (6.0 - 3.0) \div 12 = 0.25 \text{ A}$$

At the modeled stationary condition, rotational back-EMF is zero:

$$I = (6.0 - 0) \div 12 = 0.50 \text{ A}$$

One motor terminal voltage, two internal voltage terms



Settled model only. Inductance and brush voltage drop are omitted.

Figure 3. In the settled Motor M model, the terminal voltage accounts for winding voltage drop and generated back-EMF.

Now the Chapter 2 readings stop looking like unrelated facts. Slower rotation means less opposing generated voltage. Less opposition leaves more of the applied voltage across the winding resistance. More current follows.

That current produces more electromagnetic torque. The motor is not intelligently asking the source for help. Its physical state changed the electrical relationship, and current changed as a consequence.

Startup and stall are related, but not identical

At the instant a stopped motor is energized, rotational back-EMF is also zero. That helps explain large startup current. But our 0.50 A figure is a settled paper value using twelve ohms. A real winding also has inductance, which resists sudden changes in current. Chapter 6 will put time back into the story.

Do not read 0.50 A as a measured instantaneous spike, a continuous safe rating, or permission to hold a motor stalled. It is the answer to one deliberately simplified question.

Three different changes people call "more power"

Suppose somebody says, "Give the motor more amps." That sentence is incomplete.

Raise applied voltage: available speed can rise, and current can change because the electrical balance changed. How much depends on speed, load, controller behavior and the rest of the system.

Increase mechanical load at the same applied voltage: the motor slows, back-EMF falls and current rises. In Motor M, our move from free-running to the specified load raised current from 0.05 A to 0.25 A.

Raise a controller's current ceiling: that does not guarantee the motor will always draw the new maximum. It permits more current when operating conditions demand it and the rest of the system can supply it.

A battery with greater current capability follows the same important distinction: capability is permission, not a command. A source capable of 20 A does not automatically ram 20 A through a motor operating at 0.25 A.

Watts tell us how fast energy is moving

Chapter 2 gave power a unit: the watt. Let P stand for electrical power. Now we can write the relationship:

$$P = V \times I$$

P = electrical power, measured in watts W. V = voltage across the device. I = current through the device. $\times$ means multiply.

$$\text{Free-running: } 6.0 \times 0.05 = 0.30 \text{ W}$$

$$\text{Loaded: } 6.0 \times 0.25 = 1.50 \text{ W}$$

$$\text{Stationary paper case: } 6.0 \times 0.50 = 3.0 \text{ W}$$

You might expect the three-watt condition to be the most useful because it has the largest input number. The shaft is stationary, so useful shaft power is zero. The electrical energy is going somewhere else.

Current gets expensive quickly

Let $P(\text{copper})$ mean power becoming heat in the copper winding, and let I^2 mean current squared, or current multiplied by itself. Some electrical energy becomes heat simply because current passes through winding resistance:

$$P(\text{copper}) = I^2 \times R$$

R is winding resistance.

$$\text{Loaded: } 0.25 \times 0.25 \times 12 = 0.75 \text{ W of copper heat}$$

Half of the 1.50 W electrical input in this simplified loaded case becomes winding heat. The other 0.75 W is converted electromagnetically before mechanical and other rotational losses take their share. We are not claiming a measured 50-percent shaft efficiency.

$$\text{Free-running: } 0.05 \times 0.05 \times 12 = 0.03 \text{ W of copper heat}$$

The remaining 0.27 W of the 0.30 W input is converted and, because there is no useful external shaft load in our model, consumed by internal rotational losses.

$$\text{Stationary: } 0.50 \times 0.50 \times 12 = 3.0 \text{ W of copper heat}$$

All three watts become winding heat in this simplified stationary model while shaft power remains zero. A motor can be doing no useful mechanical work and still be having a terrible day.

Why doubling current is not a small change

Keep resistance unchanged for the moment. Double current:

$$2^2 = 2 \times 2 = 4$$

The current part of resistive heating becomes four times larger. Triple current:

$$3^2 = 3 \times 3 = 9$$

It becomes nine times larger.

You might expect twice the current to mean twice the heat. The actual relationship squares current. Real copper resistance also rises as the winding gets hotter, so unchanged resistance is a teaching condition, not a promise about a cooking motor.

Power tells us the rate at which energy is being transferred. Temperature tells us how hot something is. They are not the same quantity. A component can tolerate a large heating rate briefly and fail under a smaller one held long enough. Thermal mass, how much heat energy a part can absorb for a given temperature rise, and cooling will matter later.

Use equations to narrow a fault, not to invent certainty

Suppose a $120\ \Omega$ resistor has 6.0 V across it. The simple model predicts 0.050 A. If the measured current is near that value, the result agrees with the model. If it is zero, we have evidence of an open path somewhere in the measurement boundary. If it is much larger, either the resistance is not what we thought, the voltage is not what we thought, the connection is not what we thought, or the model does not describe the actual circuit.

Now return to Motor M. A stopped resistance reading of $12\ \Omega$ does not tell us its free-running current by itself because the running motor generates back-EMF. Likewise, a healthy 6.0 V measured at the source does not prove that 6.0 V reaches the motor while the circuit is loaded.

This is the beginning of troubleshooting rather than arithmetic. Every calculation has a boundary. Every measurement has an address. When they disagree, the disagreement tells us where to ask the next question.

One route was easy. Next we give electricity choices.

So far Motor M has lived in a simple circuit with one route. Real machines branch. They run lights beside controllers, sensors beside motors, and sometimes combine more than one source.

Chapter 4 adds a second road. Two identical $120\ \Omega$ resistors in series across 6.0 V will carry 0.025 A and, in this ideal equal case, each will have 3.0 V across it. Put those same resistors in parallel and each branch will see the full 6.0 V, each will carry 0.050 A, and the source will supply 0.100 A total.

Those answers sound contradictory until the routes are visible. Then they become almost obvious.

And once the routes are visible, we are finally ready for a particularly Creaky idea: two batteries that started in parallel and shifted into series like an electrical transmission. It worked. Then the switch melted. Both parts of that story are about to make sense.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [2], [3], [4], [8]. Full titles and links appear in the source list.

4. More than one road

Until now, nearly every circuit in this book has had one route. That was deliberate. Real machines branch. The moment we add a second route, two words become useful: series and parallel. We will see the behavior first and name it second.

Two identical resistors, one road

Take two identical $120\ \Omega$ resistors and an ideal $6.0\ \text{V}$ source. Connect the resistors end to end so current has only one route.

$$R_{\text{total}} = 120\ \Omega + 120\ \Omega = 240\ \Omega$$

$$I = 6.0\ \text{V} \div 240\ \Omega = 0.025\ \text{A}$$

That same $0.025\ \text{A}$ passes through both resistors. There is nowhere for it to split.

$$V_{\text{across each}} = 0.025\ \text{A} \times 120\ \Omega = 3.0\ \text{V}$$

Three volts across the first resistor plus three volts across the second accounts for the six-volt source. This one-route arrangement is called series.

Current is not used up

You might expect the first resistor to consume some current and leave less for the second. The actual answer is $0.025\ \text{A}$ through both. A resistor transfers electrical energy; it does not consume charge flow.

Replace the second resistor with $240\ \Omega$. Total resistance becomes $360\ \Omega$.

The symbol $\approx$ means approximately equal to, so $I = 6.0\ \text{V} \div 360\ \Omega \approx 0.0167\ \text{A}$.

$120\ \Omega$ resistor: about $2.0\ \text{V}$

$240\ \Omega$ resistor: about $4.0\ \text{V}$

You might expect two components to get three volts each. They do not. Equal components divided the voltage equally. Unequal resistances do not. The same current passes through both, so the larger resistance gets the larger voltage difference.

Two roads between the same two places

Return to two identical $120\ \Omega$ resistors. Connect each one directly between the source's positive and negative electrical points. Current now has two routes. This is parallel.

Each branch begins and ends at the same two electrical points, so each branch has the full $6.0\ \text{V}$ across it.

$$\text{Branch 1: } 6.0\ \text{V} \div 120\ \Omega = 0.050\ \text{A}$$

$$\text{Branch 2: } 6.0\ \text{V} \div 120\ \Omega = 0.050\ \text{A}$$

$$\text{Source current} = 0.050\ \text{A} + 0.050\ \text{A} = 0.100\ \text{A}$$

You might expect the six volts to split into three volts per branch. The actual answer is six volts on each branch. What split was the route, not the voltage between the common endpoints.

Adding resistance can reduce total resistance

Two $120\ \Omega$ resistors in parallel behave, from the source's point of view, like $60\ \Omega$.

$$R_{\text{equivalent}} = 6.0 \text{ V} \div 0.100 \text{ A} = 60 \Omega$$

That sounds backward until you stop picturing 'more resistors' and picture 'more roads.' The source now has two paths through which current can flow.

The shorthand

Series resistors: $R_{\text{total}} = R_1 + R_2 + \dots$, where the small numbers label individual resistors and ... means continue the pattern.

Parallel branches: $I_{\text{total}} = I_1 + I_2 + \dots$

Parallel resistors: $1 \div R_{\text{total}} = (1 \div R_1) + (1 \div R_2) + \dots$

The expression $1 \div R$ is called a reciprocal. A calculator can do it. The physical picture matters more: series gives one route; parallel gives branches.




Figure 4. The same two 120Ω resistors give different current and voltage relationships when the connection changes.

Motors in series are not two resistors wearing shafts

Put two identical Motor M units in series across 12 V. If they are equally loaded and spinning alike, a balanced example could place about 6 V across each.

Now load one motor harder. It slows. Chapter 3 tells us its back-EMF falls. Its electrical behavior changes, so the two motor terminal voltages no longer have to divide equally.

The rule 'two identical devices in series get half the voltage' only worked because the operating conditions made the devices electrically alike. Motors can stop being alike while their labels remain identical.

A real source notices another branch

Our ideal six-volt source never sags. A real battery or power supply can. One 120Ω branch asks for 0.050 A. Add an identical parallel branch and the source is asked for 0.100 A.

A capable source may stay close to 6.0 V. A weak source may fall below 6.0 V. Nothing about the first branch changed. The total job changed.

This is why a battery can show normal voltage with almost no load and sag badly when a motor asks for current.

Batteries in series

Sources deserve their own example because series and parallel mean something different when the things being connected are supplying energy.

Two compatible ideal 6 V sources in series, with polarity assisting, make about 12 V available.

$$6\text{ V} + 6\text{ V} = 12\text{ V}$$

If each is a 5 Ah battery, the ideal series string is still 5 Ah. The voltage doubled; the amp-hour figure did not.

Energy did increase because there are two batteries. One ideal 6 V, 5 Ah battery represents 30 watt-hours (Wh). Two in series represent $12\text{ V} \times 5\text{ Ah} = 60\text{ Wh}$.

Watt-hours are a unit for an amount of energy. Amp-hours describe charge capacity; watt-hours also account for voltage. Multiplying nominal voltage by rated amp-hours gives an approximate energy figure because real battery voltage changes during discharge.

Batteries in parallel

Two compatible, equal-voltage ideal 6 V, 5 Ah sources in parallel remain about 6 V while their combined capacity becomes about 10 Ah.

In a matched example, they can share load current. A 4 A load might be supplied at roughly 2 A from each source.

You might expect two six-volt batteries to make twelve volts because there are two of them. In parallel they do not. Their positive terminals share one electrical point and their negative terminals share another. The voltage between those common points remains about six volts.

Do not parallel batteries that disagree

The previous paragraph is not permission to connect arbitrary lithium packs together.

If one source is at a higher voltage than the other, paralleling them creates a path from the higher-voltage source into the lower-voltage source. Equalization current can be very large because battery and wiring resistance can be very small.

Chemistry, state of charge, cell count, battery management system (BMS) behavior, pack condition and manufacturer design matter. A BMS monitors battery conditions and may disconnect charging or discharging when a limit is crossed. It does not make arbitrary pack combinations compatible. Our battery examples are paper models, not a beginner lithium-pack exercise.

Creaky's two-speed battery

Now the old scooter experiment makes sense.

Two DeWalt tool batteries powered a small brushed scooter. A center-off double-pole, double-throw (DPDT) switch reconfigured them. For startup they were placed in parallel. For higher speed they were switched into series.

Double-pole means two electrical paths are controlled together. Double-throw means each path can be connected in either of two ways. This switch also had a center OFF position.

Parallel kept approximately one-pack voltage while allowing the sources to share the burden. Series added their voltages and gave the motor substantially more available voltage for speed.

Nothing mechanical changed ratio, but the way the sources were connected changed. Engineers call that connection arrangement the circuit topology. Calling it an electrical transmission is fair.

Creaky reports that the scooter reached about 18 mph. Then the little switch melted.

The surviving facts do not establish one exact forensic cause. We know the switch melted and the wires were manually disconnected to stop the system. Chapter 3 has already given us a strong category to investigate: current through resistance makes heat, and real switch contacts are not zero-ohm lines.

Topology is already a diagnostic tool

Imagine two lamps in parallel. Lamp A works and lamp B is dead. Because A works, the common source path before the split cannot be completely open. The fault is more likely in branch B or a connection unique to B.

Now put two lamps in series. One connection opens. Current stops everywhere in that route and both lamps go dark.

A circuit diagram is already becoming a map of which failures are capable of producing which symptoms.

One last answered example

Two 120 Ω resistors are in parallel across an ideal 6.0 V source. One resistor quietly goes open circuit.

You might expect the remaining resistor to receive more than six volts because it is now the only load. It does not. The ideal source still establishes 6.0 V across the same two points.

Before the fault, each branch carried 0.050 A and source current was 0.100 A. After one branch opens, the healthy branch still carries 0.050 A while source current falls to 0.050 A.

Voltage can remain normal while total current changes because a parallel path disappeared.

Next: make the circuit map real

Our drawings have quietly lied about something. Wires have been perfect. Closed switches have had zero resistance. Connectors have never loosened, oxidized or heated.

Real copper and real contacts are less cooperative.

Chapter 5 starts with the melted DPDT switch and asks a narrow question: how can a resistance so small that we normally ignore it become hot enough to destroy a connection?

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment. The DeWalt parallel/series account is Creaky's confirmed workshop history; the precise switch-failure mechanism was not measured.

Written references: [4], [8], [9], [23]. Full titles and links appear in the source list.

CHAPTER 5

5. Why the switch melted

Chapter 4 ended with a switch that successfully changed two batteries from parallel to series and eventually melted anyway. That is a useful kind of failure because the circuit idea and the physical hardware can be judged separately.

A schematic is a map of electrical connections rather than a picture of where the physical parts sit. On a schematic, a closed switch is usually drawn as though it becomes a perfect piece of wire. Real switches are made from metal contacts pressed together with finite area, finite pressure, finite conductivity, and finite ability to shed heat. Real wires, crimps and connectors have the same problem in smaller doses.

This chapter makes those nearly invisible resistances visible.

The schematic has been lying by omission

Imagine Motor M operating at the Chapter 3 loaded condition: 6.0 V at the motor and 0.25 A through it. In our earlier drawings, the wire and closed switch had exactly zero resistance.

A real conductor does not. Copper is very conductive, which means its resistance can be small enough to ignore for many jobs. Small is not the same as zero.

Switch contacts, connector pins, crimps and solder joints can add resistance too. A crimp is a metal terminal squeezed tightly around a conductor. A solder joint joins conductors with solder. When current is small, the effect may be tiny. When current becomes large, the same resistance can become a heater.

Meet the milliohm

High-current connections are often discussed in milliohms because one whole ohm would be enormous for a good power connection.

1 Ω	one ohm
1 m Ω	one milliohm = 0.001 Ω
m	milli, meaning one thousandth

A connection measuring 20 m Ω has a resistance of 0.020 Ω . That looks almost like nothing. Keep it in mind.

Voltage drop lets us see resistance while the circuit works

Suppose a healthy closed switch has 0.020 Ω of resistance and Motor M is drawing 0.25 A.

$$V \text{ drop} = I \times R$$

$$V \text{ drop} = 0.25 \text{ A} \times 0.020 \Omega = 0.005 \text{ V}$$

Only five thousandths of a volt appears across the closed switch. In this small example, ignoring it was reasonable.

Now imagine a damaged or badly chosen contact with 0.50 Ω resistance at the same 0.25 A:

$$V \text{ drop} = 0.25 \text{ A} \times 0.50 \Omega = 0.125 \text{ V}$$

The connection is still physically closed, but it is no longer electrically insignificant.

This is why voltage-drop testing is so useful. Measuring a tiny resistance directly can be difficult, especially through probes and connections. Measuring the voltage that appears across the suspect connection while real current flows can reveal the problem under the conditions that create it.

Tiny resistance, large current

Our little Motor M does not make the danger dramatic. A larger motor circuit does.

Take the earlier 0.020 Ω switch contact and pass 10 A through it. Chapter 3 already gave us the heating relationship:

$$P_{\text{heat}} = I^2 \times R$$
$$10 \times 10 \times 0.020 = 2 \text{ W}$$

Two watts concentrated inside a small switch contact can matter.

Double the current to 20 A:

$$20 \times 20 \times 0.020 = 8 \text{ W}$$

The current doubled. The resistive heating became four times larger.

Now suppose looseness, oxidation, damage or inadequate contact pressure raises the connection resistance to 0.10 Ω while 20 A flows:

$$20 \times 20 \times 0.10 = 40 \text{ W}$$

You might have looked at 0.10 Ω and thought the resistance was too small to matter. The actual answer is forty watts of heating in this simplified example. The small resistance became important because current was large and current is squared in the heating relationship.

A bad connection steals voltage too

Heating is only half the symptom. The voltage lost across the connection is voltage the load does not receive.

At 20 A through 0.10 Ω :

$$V_{\text{drop}} = 20 \text{ A} \times 0.10 \Omega = 2.0 \text{ V}$$

A motor supplied through that connection sees two volts less than it otherwise would, while the connection itself dissipates forty watts as heat.

This is a common real-world combination: the machine becomes weak and something upstream gets hot.

Mechanically connected is not the same as electrically healthy

A connector can click into place and still be bad electrically. A crimp can hold the wire but grip only a few strands. A bolted lug can look tight while oxidation sits between conducting surfaces. A spring contact can lose pressure. Vibration can turn a good connection into an intermittent one.

Electrical current does not care whether a connection looks finished. It cares about the actual conductive path.

Useful things to inspect include wire cross-section, length, terminal contact area, crimp quality, solder quality where solder is appropriate, connector spring pressure, corrosion, oxidation, heat discoloration, looseness and vibration damage.

Wire size and length are part of the circuit

A longer wire has more resistance than a shorter wire of the same material and cross-section. A thinner wire has more resistance than a thicker wire of the same material and length.

That does not mean 'always use the biggest wire possible.' Larger wire costs space, mass, flexibility and money. It means wire is selected for the current, acceptable voltage drop, temperature, installation and protection of the actual circuit.

The same reasoning applies to printed circuit board (PCB) traces, the flat conductive paths on a board; bus bars, heavy conductors used to distribute substantial current; and connector pins. Geometry is electrical.

Ratings are conditions, not force fields

A switch marked with a current rating is not promising that everything below that number is harmless and everything one milliamp above it explodes.

A rating belongs to test conditions: voltage, current type, contact material, duty, meaning how long and how often it carries or switches current, temperature, enclosure, switching frequency and sometimes the kind of load.

Continuous current and momentary current are different jobs. Carrying current through already-closed contacts and interrupting that current are also different jobs.

Voltage rating matters because opening contacts must withstand the voltage across the widening gap. Current rating matters because contacts must carry and interrupt current without unacceptable heating, erosion or welding.

DC is stubborn when you try to open it

Alternating current (AC) reverses direction and naturally passes through zero current every cycle. Direct current (DC) does not provide that recurring zero crossing.

When contacts open under DC load, an arc can persist across the separating contacts more readily than it would in an otherwise comparable AC switching situation. That is one reason a switch can have different AC and DC ratings.

We are deliberately stopping short of the nastier part. Motors and relay coils can generate additional voltage when their current is interrupted. Chapter 6 is about that.

What a fuse does, and what it does not do

A fuse is a deliberately weak section of conductor intended to open the circuit when current exceeds its designed time-current behavior.

Its main job is to prevent wiring and the rest of the protected circuit from carrying dangerous fault current indefinitely. In many DC systems it is placed near the source so as much downstream wiring as practical is protected.

A fuse does not make an undersized switch suitable for normal operation. If a switch overheats at 20 A while the fuse is correctly designed to tolerate that normal 20 A load, the fuse may never open. The switch can cook below the fuse's fault threshold.

A fuse is also not a motor controller. It is protection, not regulation.

Heat can make the connection worse

A connection begins with some resistance. Current heats it. Heat can soften surrounding plastic, change contact pressure, accelerate oxidation, damage plating or alter the joint mechanically. If those changes increase resistance, the next moment produces still more heating.

That feedback can run away if current remains high: heat damages the connection, resistance rises, and the joint dissipates still more power. In other circuits the added resistance reduces current enough to limit heating. The outcome depends on the source and load as well as the joint.

This is one form of thermal runaway. The phrase does not mean every warm connector inevitably runs away. It describes a feedback condition in which rising temperature makes the cause of heating worse.

Return to the little DPDT

The historical setup used a small center-off DPDT switch to reconfigure two DeWalt batteries between parallel and series operation on a brushed scooter.

The switch eventually melted. The rider had to disconnect wires manually to stop the system.

We can now name several plausible contributors: contact resistance, current beyond what the switch could continuously carry, heating during repeated switching, DC arcing, deteriorating contact pressure, or some combination. The surviving story does not give us measurements that prove which mechanism dominated.

That distinction matters. Failure analysis is not improved by pretending a plausible explanation is a measured one.

What we can say confidently is simpler: the switch was part of the motor-current path, and a real switch must survive both the current it carries and the electrical conditions created when it changes state.

Answered diagnostic case: the motor is weak

A motor system is supposed to receive about 6.0 V under load. The battery terminals measure 6.0 V while the motor is running, but the motor terminals measure only 4.8 V.

You might suspect the motor because the symptom is weak motor output. The measurements tell us something more useful: 1.2 V is being lost somewhere between the source terminals and the motor terminals.

$$6.0 \text{ V} - 4.8 \text{ V} = 1.2 \text{ V missing in the path}$$

Measure voltage directly across each connection while the motor remains loaded.

Wire section A: 0.03 V. Closed switch: 0.04 V. Connector C: 1.10 V. Remaining wiring: about 0.03 V.

Connector C is the standout. A healthy closed connection is supposed to have very little voltage across it. Here, 1.10 V appears across the connector because current is encountering substantial resistance there.

If the loaded current is 2.0 A, its effective resistance under that condition is:


$$R = 1.10 \text{ V} \div 2.0 \text{ A} = 0.55 \Omega$$

And its heating is:

$$P = 1.10 \text{ V} \times 2.0 \text{ A} = 2.2 \text{ W}$$

You did not need to measure 0.55 Ω directly. The loaded voltage-drop measurement exposed the fault where it actually mattered.

Measure across the suspect connection while loaded



$$1.10 \text{ V} \times 2.0 \text{ A} = 2.2 \text{ W at this one connection}$$

Source: 6.0 V. Motor: 4.8 V. Other path losses total 0.10 V.

Figure 5. A loaded voltage measurement exposes a resistive joint without relying on a low-ohms meter reading.

After repairing or replacing the connection, the correct final step is not 'it looks better.' Run the system under load again and verify that the voltage drop has returned to a small value and the motor receives the expected voltage.

The next problem survives even a perfect switch

Suppose we replace the melted switch with one that comfortably carries the current, uses appropriate wire and connectors, and stays cool.

We still have a motor.

Its windings create a magnetic field while current flows. When the switch opens, that stored magnetic energy does not politely vanish just because the contacts separated.

Chapter 6 asks where that energy goes. That brings us back to induction, and introduces the diode, capacitor and snubber because now we have a problem that needs them.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment. The DeWalt parallel/series account is Creaky's confirmed workshop history; the precise switch-failure mechanism was not measured.

Written references: [8], [13], [22], [25], [27]. Full titles and links appear in the source list.

CHAPTER 6

6. The current doesn't stop politely

Chapter 5 explained how an ordinary connection can heat itself to death. But suppose the switch is properly sized, the wire is adequate, every crimp is good, and nothing is overheating.

We still have a motor.

Open the switch while current is flowing through a motor winding and the current does not simply agree to become zero. The reason is not mechanical momentum in the electrons. The useful analogy is momentum because something in the system resists abrupt change. What is actually stored is energy in a magnetic field.

This chapter follows that stored energy from the moment a coil is energized until the moment we try to make it disappear.

Start with a relay because it makes the invisible visible

A relay contains a coil of wire and a movable set of contacts. Apply voltage to the coil and current flows through the wire. The coil creates a magnetic field. That field pulls a small mechanical armature. You hear a click.

Remove power and the magnetic field collapses. A spring returns the armature. You hear another click.

The relay has turned an electrical event into something your ears can detect.

A motor winding is more complicated mechanically, but electrically it belongs to the same family of behavior: current through wound wire creates a magnetic field.

Current creates magnetism

A current-carrying conductor produces a magnetic field around it. Wind the conductor into a coil and the fields from many turns reinforce one another. Put suitable magnetic material inside the coil and the field can be concentrated and strengthened.

That is enough physics for the moment. We do not need a field equation to understand the problem at the switch.

While current flows, energy is stored in that magnetic field.

A coil resists sudden change in current

A coil has a property called inductance. Inductance describes, among other things, how strongly the circuit opposes a change in current by generating a voltage.

An inductor is simply a component deliberately built to exploit that property. But a relay coil and motor winding have inductance whether or not anyone calls them inductors.

This resolves a loose end from Chapter 3. We calculated a settled stationary Motor M current of 0.50 A from 6.0 V and 12 Ω . We warned that this was not necessarily the current at the first instant after the switch closes.

At switch-on, the magnetic field is building. The winding's inductance opposes the rapid rise of current. Current takes time to approach the value set by the rest of the circuit.

The simple $V \div R$ calculation was not wrong. It described the eventual resistive limit of our simplified stationary model after the transient behavior was omitted.

Now open the switch

This is where the coil becomes rude.

Just before the switch opens, current is flowing and the magnetic field contains stored energy. The contacts separate. The old conductive path disappears.

But the collapsing magnetic field generates a voltage that tries to keep current moving in the same direction it had been moving.

If no easy path exists, the generated voltage can become much larger than the original supply voltage until some path is found: across an air gap, through an electronic component that normally blocks it, through insulation, or through some other part that was not volunteering.

That switch-off voltage is commonly called flyback, inductive kick, or a back-voltage transient. A transient is a brief electrical event that appears during a change.

A supplied example, not a universal number

Imagine a teaching relay coil operating from 6.0 V and drawing 0.10 A while energized. Its switch is between the coil and source negative. In our illustrative unprotected case, opening the switch produces a brief +45 V at the coil-switch junction relative to source negative.

You might expect a six-volt circuit to be incapable of producing forty-five volts. The actual answer is that the supply is not directly producing the forty-five volts. The collapsing magnetic field is converting stored magnetic energy back into electrical energy.

Forty-five volts is an illustrative number for this lesson, not a prediction for every six-volt relay. Actual transient voltage depends on inductance, current, switching speed, capacitance, insulation, arcing, suppression and the rest of the circuit.

Induction is the common thread

The process of a changing magnetic field producing voltage is electromagnetic induction.

We already met one version without giving it the full family name. In Chapter 3, a spinning motor generated back-EMF. Motion changed the relationship between conductors and magnetic field, and a voltage appeared.

Here, the magnetic field itself is collapsing because current was interrupted, and a voltage appears.

These are related consequences of induction, but they are not the same event. Running back-EMF is part of normal motor rotation. Flyback is a transient associated with changing current, especially interruption.

We need somewhere for the current to go

The problem suggests the next component.

When the switch opens, we want the coil current to have a safe temporary path instead of forcing the voltage high enough to invent one.

A diode can provide that path.

Meet the diode

A diode is a component that conducts readily in one direction and blocks strongly in the other under normal intended conditions.

For now, think of it as an electrical one-way door. Chapter 10 will explain how semiconductor material makes that behavior possible.

Place a diode across a DC relay coil in the correct direction. While the relay is energized normally, the supply voltage tries to push through the diode in its blocking direction, so very little current flows through that path.

When the switch opens, the coil's generated voltage reverses polarity as necessary to keep its current moving. That new polarity makes the diode conduct.

The current now circulates through the coil and diode while the stored magnetic energy decays.




Figure 6. A diode across a fixed-polarity DC coil. K marks the cathode and A the anode, the diode's two terminals. At switch-off the coil current keeps its direction; the diode completes a local loop. This is not a reversible-motor protection circuit.

What the diode changes

Without the diode, our teaching example produced a brief 45 V transient.

With an ordinary silicon flyback diode in the simplified model, the junction between the coil and switch, called the switch node, can be held near a much safer voltage instead of rising by tens of volts. Engineers call this holding action clamping.

You might expect the diode to make the stored energy disappear. It does not. It gives the current a lower-voltage route while resistance in the coil and diode turns the stored energy into heat over time.

There is a tradeoff: lower clamp voltage generally means the current decays more slowly, so a relay may release more slowly. Protection choices affect behavior.

The diode has a direction

A diode used for flyback must be oriented correctly. Reverse it accidentally across the supply and it can behave almost like a short circuit when power is applied.

That is why diode polarity markings matter. A common physical diode has a band marking the cathode end. Schematic symbols use a line on the cathode side.

You do not need to memorize every diode symbol variant now. You do need to understand that the component is directional and that direction is part of the circuit.

What if we need to catch a fast voltage change without a one-way loop?

Another component can temporarily accept electrical energy into an electric field: the capacitor.

At its simplest, a capacitor consists of two conductive regions separated by insulating material. Charge can accumulate on the conductive regions, establishing an electric field between them.

Because its stored energy is associated with voltage, a capacitor can oppose very rapid changes in voltage by accepting or releasing charge.

That makes capacitors useful around switching circuits, power supplies and motors. It does not mean 'put a capacitor anywhere a spike exists.' Value, voltage rating, resistance, inductance, placement and the desired transient behavior all matter.

Meet the snubber

A common way to tame a switching transient is a snubber. One simple form uses a resistor and capacitor together, often called a resistor-capacitor (RC) snubber.

The capacitor provides a temporary path for the rapid voltage change. The resistor limits the resulting current and dissipates stored energy rather than allowing the capacitor and circuit inductance to ring indefinitely.

We are not yet calculating snubber values. The point is to understand the job: give transient energy a controlled place to go and control how quickly voltage changes.

Return to the DeWalt experiment

The parallel-start/series-run scooter used a snubber as protection against switching transients. That was an early electronics experiment motivated by a practical concern: DeWalt batteries were too expensive to casually expose to ugly electrical events.

The snubber did not make the little DPDT switch capable of carrying unlimited motor current. Chapter 5 dealt with that separate problem.

Nor do we have measurements proving exactly what transient amplitude occurred or exactly what the snubber prevented. The defensible statement is that the snubber was included to manage switching transients and protect the battery/electrical system from inductive kick.

That distinction between two simultaneous problems matters: one component can fail from carrying current while another protection network is successfully handling voltage transients.

Relays show us control without carrying control current through the control

Look again at the relay.

The coil circuit can be small. Its magnetic field moves a separate set of contacts. Those contacts can switch another circuit.

That gives us our first clean example of a low-power control action commanding a separate power path.

The idea will become extremely important. We do not always want the knob, sensor or logic circuit to carry the motor current it controls.

Symptoms that can point toward inductive trouble

A switch that arcs badly when opening a coil, an electronic switch that repeatedly dies when a relay is released, a controller that resets when a motor is disconnected, or electrical noise that appears exactly at switching events can all make transient behavior worth investigating.

None of those symptoms proves flyback by itself. Troubleshooting remains evidence-based. But the reader now has another failure class to consider: stored magnetic energy was forced through a path that could not tolerate it.

One answered comparison

A six-volt relay works normally while energized. Every time it is switched off, a nearby control circuit resets.

You might suspect the relay is drawing too much current while on. That is possible in general, but the timing gives us a better clue: the reset happens at switch-off, not during the energized period.

That makes an inductive transient a stronger hypothesis. The next useful measurement would be the switching-node voltage during release, ideally with an oscilloscope, an instrument that displays voltage changing over time, or another suitable transient-measurement method.

If adding correctly designed flyback suppression removes the switch-off transient and the resets stop, the evidence becomes much stronger.

The lesson is not 'every reset needs a diode.' The lesson is to match the timing of the symptom to the physical event that changed.

The next bad idea is useful

We can now switch a motor on and off without pretending the winding is just a resistor. But our motor still has only two speeds: running and stopped.

Suppose we want half speed.

A very natural idea is to put an adjustable resistor in series with the motor and turn a knob until the motor slows down.

It works.

That is the problem.

Chapter 7 builds the terrible speed controller, watches where the energy goes, and discovers why a control knob should not have to wrestle the motor's power directly.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment. The DeWalt parallel/series account is Creaky's confirmed workshop history; the precise switch-failure mechanism was not measured.

Written references: [10], [11], [16], [27]. Full titles and links appear in the source list.

CHAPTER 7

7. The terrible speed controller

We can now make Motor M turn, measure what it is doing, predict its current, understand why load changes that current, wire more than one route, and stop an inductive load without pretending its magnetic field vanishes instantly.

There is still a painfully ordinary problem: the motor is either on or off.

Suppose full speed is too fast. We want a knob.

The first solution is obvious, understandable, and bad enough to teach us exactly what a real controller needs to do.

Put resistance in the way

Chapter 3 showed that resistance restricts current. So put an adjustable resistance in series with the motor.

As the added resistance increases, some of the source voltage appears across that resistance instead of across the motor. At first, current and electromagnetic torque fall. The motor slows and its back-EMF changes. The eventual current depends on how the mechanical load changes with speed.

It works.

That is why this idea is worth taking seriously before rejecting it.

Meet the potentiometer

A potentiometer, often shortened to pot, is a resistor with a movable electrical contact called a wiper.

An ordinary potentiometer has three electrical connections: the two ends of its resistive track and the movable wiper. On many panel pots the wiper is the middle terminal, but package layouts vary. Use the exact terminal diagram.

Turn the shaft or slide the control and the wiper moves along the resistive material.

potentiometer	an adjustable resistive component
resistive track	the material that provides resistance
wiper	the movable contact that touches the track
three terminals	the two ends of the track plus the wiper connection

Use one end terminal and the wiper and the potentiometer can act as a variable resistor. Use all three terminals and it can divide a voltage, which becomes much more useful later.

Our deliberately bad controller

Use a constructed operating point rather than pretending Motor M is a fixed resistor.

Suppose Motor M is connected to 6.0 V under a particular mechanical load and is operating at the Chapter 3 loaded condition: 0.25 A.

Now insert an adjustable series resistor and turn it until the motor settles at a slower operating point. For this example, the required load torque falls as speed falls. Current is now 0.20 A, the motor has 4.0 V across its terminals, and the added resistance has the remaining 2.0 V across it. A constant-torque load would not necessarily settle at this lower current.

The source still provides 6.0 V:

$$4.0 \text{ V across motor} + 2.0 \text{ V across control resistor} = 6.0 \text{ V}$$

The series resistor's effective value at this setting is:

$$R = 2.0 \text{ V} \div 0.20 \text{ A} = 10 \, \Omega$$

The motor slowed because we deliberately prevented it from receiving the full source voltage.

Where did the missing electrical power go?

The added resistor is not storing the unwanted voltage for later. Current passes through it while a voltage exists across it, so it consumes electrical power.

$$P = V \times I$$

$$P_{\text{resistor}} = 2.0 \text{ V} \times 0.20 \text{ A} = 0.40 \text{ W}$$

That 0.40 W becomes heat in the series resistance.

The motor receives:

$$P_{\text{motor input}} = 4.0 \text{ V} \times 0.20 \text{ A} = 0.80 \text{ W}$$

The source supplies:

$$P_{\text{source}} = 6.0 \text{ V} \times 0.20 \text{ A} = 1.20 \text{ W}$$

Our accounting closes: 0.80 W reaches the motor and 0.40 W is deliberately burned in the control resistance.

You might think slowing the motor means the source simply provides less power. In this circuit, part of the source power is being turned into heat solely because we chose resistance as the control method.

Scale the same idea up

The little teaching motor makes 0.40 W of wasted heat sound harmless. Use the same principle on a larger motor and the problem becomes obvious.

Imagine a motor circuit drawing 10 A while an added speed-control resistance drops 12 V.

$$P_{\text{wasted}} = 12 \text{ V} \times 10 \text{ A} = 120 \text{ W}$$

That control element is now a 120-watt heater.

If the rider asks for more torque and current rises, the heating problem becomes worse. This is why high-power resistance control needs physically large, deliberately cooled resistors or other specialized hardware. A little panel potentiometer is not secretly a motor controller.

Why the potentiometer may die

A typical small potentiometer is designed to handle a small amount of power in its resistive track and wiper. Put motor current directly through it and two weak points appear.

The resistive track itself can overheat. The wiper contact can also overheat or arc because all the motor current must pass through a small moving contact area.

A potentiometer can be perfectly suitable for telling a controller what the rider wants while being completely unsuitable for carrying the motor current that produces it.

That distinction is the point of this chapter.

The speed setting moves when the load moves

There is another problem even if we build a giant resistor that never burns.

Motor M does not behave like a fixed resistor while running. Its winding resistance can remain unchanged while its back-EMF changes with speed.

Suppose we choose a series resistance that gives the speed we want while the wheel is unloaded. Then the motor encounters a hill.

Mechanical load increases. Motor speed falls. Back-EMF falls. Current rises. The voltage dropped across the series resistor changes because its voltage drop is $I \times R$.

The motor's terminal voltage therefore changes too.

You might expect a fixed knob position to mean a fixed motor speed. It does not. The resistor controls an electrical condition, while the motor's electrical condition changes with mechanical load.

The resistor hurts exactly when the motor needs help

At startup, the motor is not yet spinning, so rotational back-EMF is near zero. That is when a motor naturally wants substantial current to make starting torque.

The series resistor restricts that current.

So our crude speed controller steals voltage and torque during the moment when the motor is trying hardest to get moving.

A scooter using this scheme could feel weak from a stop, change speed badly with hills, and waste significant battery energy as heat.

A knob does not need to carry what it controls

Chapter 6 already gave us a clue.

A relay coil can be operated by a relatively small control circuit. The coil's magnetic field moves a separate set of contacts. Those contacts switch another circuit.

The control circuit and the controlled power path are related without being the same path.

That is a much more useful architecture.

What if our speed knob carried almost no motor power at all? What if it merely produced a small electrical command saying 'a little,' 'half,' or 'full,' while another device handled the motor current?

The potentiometer gets a better job

Use all three potentiometer terminals across a small control voltage.

For example, connect the two ends of the resistive track across 5.0 V and 0 V. The wiper can then provide an adjustable voltage somewhere between those endpoints.

At one position the wiper might read 1.0 V. At another, 2.5 V. At another, 4.0 V.

Very little current needs to flow through the wiper if the receiving control circuit has high input resistance.

Now the potentiometer is not throttling the motor by burning its power. It is creating information: a voltage that represents what the knob position means.

Same resistive track, two different jobs




Figure 7. A potentiometer can carry load current as a series resistance or provide a small command signal as a divider. These are functional connections, not the physical terminal arrangement, or pinout, of a particular package.

This is our first clean separation between a command signal and the power being commanded.

Why not just use a relay?

A relay can separate control power from motor power, but it gives us another limitation.

Its contacts are mechanical. They move in milliseconds, bounce when they touch, wear when switched, and can arc.

A relay is excellent for many jobs where a circuit needs to be on or off occasionally. It is a poor choice if we want smooth speed control by switching a motor thousands of times every second.

We need the relay's basic idea without the moving contacts: a small electrical command controlling a much larger electrical path, very quickly.

The better question

Our resistor controller slowed the motor by continuously throwing away some source voltage as heat.

What if we stop trying to create a smaller continuous voltage?

What if the motor receives the full source voltage for a short time, then receives nothing for a short time, then full voltage again?

If we repeat that fast enough, the motor's mechanical system may respond to the average effect instead of lurching through every individual electrical event.

That would let the power-handling device spend most of its time in two efficient states: strongly conducting or strongly blocking, rather than deliberately sitting in the hot middle.

One answered comparison

Controller A uses a series resistor that drops 2.0 V at 0.20 A. We already calculated 0.40 W of intentional resistor heat.

Controller B is hypothetical for now. Imagine an ideal switch that is either completely open or has exactly zero resistance when closed. In the ideal model, the open switch carries zero current and the closed switch has zero voltage across it.

$$\text{Open: } P = V \times 0 \text{ A} = 0 \text{ W}$$

$$\text{Closed ideal switch: } P = 0 \text{ V} \times I = 0 \text{ W}$$

You might conclude that real electronic switching therefore wastes zero power. It does not. Real switches have resistance, take time to change state, and need control energy.

But the ideal comparison shows the direction we want to go: stop wasting large amounts of energy by intentionally operating a control element as a heater.

Next: turn it on and off really fast

Chapter 8 answers the better question.

We will give Motor M full voltage in pulses. We will name the fraction of time it is on. We will learn why switching frequency matters. Then the relay will fail our speed requirement because it is mechanical.

That failure introduces a metal-oxide-semiconductor field-effect transistor (MOSFET): an electronic power switch fast enough to make the terrible speed controller obsolete.

The potentiometer survives. It simply gets promoted from carrying motor current to telling the controller what we want.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [2], [3], [8], [12]. Full titles and links appear in the source list.

CHAPTER 8

8. Turn it on and off really fast

Chapter 7 built a speed controller that worked by wasting power in an adjustable resistance. It slowed the motor, heated the control element, weakened starting torque, and changed behavior when the mechanical load changed.

Then we asked a better question.

What if the motor receives the full source voltage for a short time, then nothing for a short time, over and over?

That question leads to pulse-width modulation, usually shortened to PWM.

Full on, full off, then somewhere between

Start with Motor M and the same 6.0 V teaching source.

Switch closed continuously: the motor receives the full source voltage.

Switch open continuously: the source no longer provides sustained drive. Any winding current present at switch-off must decay through the protection path.

Now imagine closing and opening the switch repeatedly: on for a while, off for a while, on again, off again.

If we do that slowly, the motor surges and coasts. If we do it rapidly enough, the rotating part of the motor, called the rotor, and the attached mechanical system cannot change speed much during each individual pulse. The motion becomes comparatively smooth while the electrical power is still being switched in distinct events.

One cycle

One complete ON interval plus one complete OFF interval is one cycle.

Suppose each cycle lasts 10 milliseconds. The switch is ON for 5 milliseconds and OFF for 5 milliseconds.

ms	millisecond; one thousandth of a second
cycle	one complete repeating ON/OFF pattern
ON time	the part of the cycle during which the power path conducts
OFF time	the part during which the power path blocks

The motor does not receive a mysterious new 3 V source. In a simple one-direction drive, the source supplies pulses. During the OFF intervals, winding current can keep circulating through the flyback path. No source drive is not the same as no motor current.

Duty cycle

The fraction of each cycle spent ON is called duty cycle.

$$\text{Duty cycle} = \text{ON time} \div \text{total cycle time}$$


We will often write duty cycle as a percentage; % means parts per hundred. If the switch is ON for 5 ms during a 10 ms cycle:

$$5 \text{ ms} \div 10 \text{ ms} = 0.50 = 50\% \text{ duty cycle}$$

At 25% duty, the switch is ON for one quarter of each cycle. At 75%, it is ON for three quarters. At 100%, it never enters the OFF portion. At 0%, it never enters the ON portion.

So 50% means fifty parts out of one hundred, or one half.

Same frequency; different ON time



Gate command versus time. OFF does not imply zero winding current.

Figure 8. Duty cycle changes the fraction of each cycle spent ON. These are command waveforms, not a claim that motor current instantly falls to zero.

Do not turn duty cycle into a fake speed equation

You might expect 50% duty cycle to guarantee exactly 50% motor speed.

It does not.

Duty cycle controls the pattern of applied electrical drive. Actual motor speed also depends on mechanical load, motor construction, back-EMF, winding resistance, controller losses, friction, supply sag and other conditions.

For a simple motor under a stable load, more duty cycle generally produces more average drive and therefore more speed. But '50% duty = 50% revolutions per minute (RPM)' is not a universal law.

Frequency: how often we repeat the decision

Duty cycle tells us how much of each cycle is ON. Frequency tells us how many complete cycles happen each second.

Frequency is measured in hertz. The prefix kilo, abbreviated k, means one thousand, so kHz means kilohertz.

Hz	hertz; cycles per second
1 Hz	one cycle each second
100 Hz	one hundred cycles each second
1 kHz	one thousand cycles each second
k	kilo, meaning one thousand

At very low PWM frequency, a motor can respond to individual pulses with obvious surging, vibration or noise.

Raise frequency and the mechanical system averages the pulses more smoothly. But higher frequency also means more switching events every second, and real electronic switches lose some energy during each transition.

So higher is not automatically better. Frequency is chosen as part of a system.

Why this can waste less power than the terrible resistor

The series-resistor controller deliberately spent much of its time with substantial current flowing through a component that also had substantial voltage across it. That combination means power becomes heat.

$$P = V \times I$$

An ideal switch has two wonderfully boring states.

OFF: it can have voltage across it, but current through it is zero.

$$P_{\text{ideal OFF}} = V \times 0 \text{ A} = 0 \text{ W}$$

ON: it can carry current, but voltage across it is zero.

$$P_{\text{ideal ON}} = 0 \text{ V} \times I = 0 \text{ W}$$

A real switch is not ideal. But if we can make it spend most of its time close to those two states, it can waste much less power than a resistor deliberately held in the hot middle.

The relay almost solves the problem

Chapter 6 gave us a device that already separates control from power: the relay.

A small coil circuit moves contacts that can switch another circuit. In principle we could turn a relay on and off to pulse the motor.

In practice, a relay has moving metal parts. Its contacts take time to move, bounce when they land, wear, and can arc.

Ask it to switch thousands of times per second and the relay becomes the next terrible speed controller.

We need an electronic switch.

Meet the MOSFET

A MOSFET is a type of transistor, an electronic device that controls a current path, commonly used as a fast electronic switch in power electronics.

MOSFET is short for metal-oxide-semiconductor field-effect transistor. A semiconductor is a material whose electrical conductivity can be controlled. The long name describes the device's construction and operating principle. You do not need to memorize the expansion yet. Chapter 10 will make the semiconductor part less mysterious.

For now, a common power MOSFET gives us three important terminals:

gate	the control terminal
drain	one end of the main controlled current path
source	the other end of the main controlled current path

A suitable voltage between gate and source changes how strongly the MOSFET conducts between drain and source.

That means a relatively low-power control signal can command a much larger motor-current path without the control knob itself carrying the motor current.

A simple low-side switch

One basic arrangement puts the motor between battery positive and the MOSFET drain, with the MOSFET source connected toward battery negative. Because the switch sits on the return side of the load, this is called low-side switching.

Drive the gate appropriately and the MOSFET conducts. Motor current flows.

Discharge the gate to its intended OFF voltage and the MOSFET blocks the main switch path. Winding current transfers into the flyback path and decays; it does not instantly vanish.

Because the motor winding is inductive, the circuit still needs a safe current path during switching. The flyback behavior from Chapter 6 did not disappear just because the switch became electronic.

The gate is not a tiny motor-current pipe

A common beginner picture is that gate current somehow flows through the MOSFET and becomes drain current.

That is not the useful model.

The gate behaves largely like a small capacitive control node. Charging it to the required gate-to-source voltage changes the electric field inside the device and therefore the conductivity of the drain-source path.

Once held at a steady voltage, an ideal MOSFET gate would need essentially no continuous DC current. A real gate must still be charged and discharged every switching cycle, so the control circuit has work to do.

Chapter 9 and later controller chapters will give that control circuitry a proper home.

The potentiometer gets promoted

The potentiometer from Chapter 7 survives, but it no longer sits in the motor-current path.

Put its resistive track across a small control supply. Its wiper now produces an adjustable command voltage.

For example, a controller might interpret 1.0 V as a low command, 2.5 V as a middle command, and 4.0 V as a high command.

The control circuit converts that requested level into PWM duty cycle.

Now the chain looks like this. In these functional chains, → means 'leads to the next stage':

potentiometer → command voltage → PWM → MOSFET → motor

The knob carries information. The MOSFET carries power.

Real MOSFETs still get hot

The ideal switch wasted no power. A real MOSFET has losses.

First, when fully ON it still has some resistance between drain and source. Datasheets call this $R_{DS(on)}$, meaning drain-to-source resistance in the ON state.

Current through that resistance creates the same I^2R heating we already know.

$$P_{\text{conduction}} \approx I^2 \times R_{DS(on)}$$

Second, the MOSFET cannot change from fully OFF to fully ON in zero time. During each transition there is a brief interval when appreciable voltage exists across the device while appreciable current also flows through it.

That creates switching loss.

More switching events per second can therefore increase switching loss even if the motor's average commanded power stays similar.

The hot middle returns if the MOSFET is driven badly

A MOSFET is efficient as a switch when it is driven decisively into the state the circuit was designed for.

If the gate voltage is too low for that particular MOSFET and load, the device may conduct only partially.

Now it can have substantial current through it and substantial voltage across it at the same time.

That makes it behave suspiciously like the series resistor we were trying to escape.

Answered diagnostic case: weak motor, hot MOSFET

A PWM controller is connected to a motor. The battery remains at 12.0 V under load. The motor is weak and the power MOSFET becomes hot quickly.

The PWM command is present, but measurement shows the gate reaches only 3.0 V above the source during the ON portion. The selected MOSFET's datasheet, the manufacturer's technical specification sheet for the part, specifies its low ON-resistance at a substantially higher gate-to-source drive voltage.

You might suspect the motor because it is weak. The battery measurement says the source is holding up. The hot MOSFET plus inadequate gate drive gives us a stronger hypothesis: the switch is not being driven fully into its intended low-resistance ON state.

If 4.0 V is being lost across the MOSFET while 5.0 A flows:

$$P_{\text{MOSFET}} = 4.0 \text{ V} \times 5.0 \text{ A} = 20 \text{ W}$$

That explains both symptoms at once: during those ON intervals the motor is missing four volts and the MOSFET dissipates twenty watts. Average loss over the PWM cycle also depends on duty cycle and switching transitions.

The repair is not automatically 'apply more gate voltage.' Gate voltage has absolute limits and must be appropriate for the specific device and circuit. The diagnostic conclusion is that gate-drive conditions do not match the MOSFET's required operating conditions.

PWM does not magically lower the battery voltage

This distinction becomes important later in scooter controllers.

A 60 V battery feeding a PWM motor stage does not become a 30 V battery at 50% duty cycle. During ON intervals, the switching stage is still dealing with the actual bus voltage, meaning the main supply voltage feeding the switching stage.

Duty cycle changes how long the motor is connected to that bus during each cycle. Component voltage ratings must survive the real instantaneous electrical conditions, not merely an imagined average.

This is one reason controller voltage matching will matter when we eventually discuss overvolting.

We have accidentally built the skeleton of a controller

Our circuit now has two layers.

The power layer contains the source, motor, MOSFET, conductors and transient-current path.

The control layer contains the potentiometer command and whatever circuitry creates the gate-switching pattern.

We know what the control layer must do, but we have not opened it yet.

Where does its 5 V come from if the motor battery is 24 V? What creates PWM? What is an integrated circuit (IC)? How does a Hall throttle replace the potentiometer? Why might a relay, transistor and MOSFET all appear in the same machine?

Chapter 9 opens that layer.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [13], [19], [27]. Full titles and links appear in the source list.

CHAPTER 9

9. Little circuits running big circuits

Chapter 8 left us with a useful skeleton:

potentiometer → command voltage → PWM → MOSFET → motor

The power path makes sense. The MOSFET switches it. But something still has to power the control electronics, read the command, and create PWM. If the motor battery is 24 V or more, the little control circuit probably does not want that voltage directly.

One machine can contain several electrical worlds

A scooter may have a 24 V motor bus, a 5 V throttle supply, and logic running at 5 V or 3.3 V. These circuits can share information while handling very different amounts of energy.

Control electronics exist so small, precise electrical behavior can command a larger power path.

The relay was our first example

A small relay-coil circuit creates a magnetic field that moves separate contacts. Those contacts switch another circuit.

small control action → separate power action

The relay is slow and mechanical, but it makes the architecture obvious.

A transistor removes the moving contact

A transistor lets one electrical condition control another current path. We need two names for now.

A BJT, or bipolar junction transistor, uses base drive to control a larger collector-to-emitter current.

base	BJT control terminal
collector	one end of the main controlled path
emitter	the other end

A MOSFET uses gate-to-source voltage to control its drain-to-source path. Both can switch. Their physical mechanisms and drive requirements differ. Chapter 10 will explain why.

Something has to make decisions

Many useful electronic functions are manufactured together inside an integrated circuit, usually shortened to IC.

An IC is not one particular job. It is many electronic structures built together on a small piece of semiconductor called a die. Pins connect selected internal points to the outside circuit.

One IC may be a timer, another a comparator that decides which of two electrical levels is higher, another a regulator, another a logic circuit. A microcontroller is an IC that executes stored instructions and interacts with inputs and outputs.

The package shape does not tell you the function. Part number and circuit context do.

The control circuit needs its own voltage

Suppose the motor battery is 24 V but the throttle and control electronics need 5 V. Feeding them 24 V does not make them stronger. It can destroy them.

A voltage regulator maintains an output near a specified voltage while input, load and temperature remain inside its operating limits.

The simple linear regulator

A linear regulator can make 5 V from a higher input by continuously controlling a semiconductor element. The unused voltage difference largely becomes heat.

Teaching example: 12 V in, 5 V out, 0.10 A load.

$$\text{Voltage difference} = 12 \text{ V} - 5 \text{ V} = 7 \text{ V}$$

$$\text{Regulator heat} \approx 7 \text{ V} \times 0.10 \text{ A} = 0.70 \text{ W}$$

Seven tenths of a watt can be significant in a small package. Low current does not guarantee low heat when a large voltage difference is being discarded.

Buck: make a lower voltage by moving energy

A buck converter is a switching power converter that produces a lower DC voltage from a higher DC voltage.

Instead of continuously burning the whole voltage difference, it switches energy through components including an inductor and capacitor, then filters the result into a useful lower DC voltage.

Real buck converters still have losses. The important distinction is that the conversion mechanism can be far more efficient than deliberately dissipating the excess voltage.

Boost: the other direction

A boost converter can produce a DC output voltage higher than its input.

It does not create energy. Higher output voltage must still be paid for by input power, output current relationships, and conversion losses.

Ground is not automatically dirt

A control circuit needs a reference point from which voltages are measured. That reference may be labeled ground, GND (short for ground), 0 V, negative or common.

Circuit ground does not automatically mean a physical connection to earth.

If battery negative is our common reference, a signal of 2.5 V means 2.5 V relative to that point. Every voltage measurement is between two points, even when one is casually called ground.

Logic voltage and reference voltage

A logic supply powers the circuitry that makes decisions. Common examples include 5 V and 3.3 V, but neither is universal.

A reference voltage is a deliberately known or stable voltage used as a measuring yardstick or as part of a sensor circuit.

Do not assume every wire labeled 5 V has the same purpose or current capability.

Replace the potentiometer with a Hall throttle

A potentiometer reports position with a rubbing wiper. A Hall-effect sensor can report position without that sliding electrical contact.

At the behavioral level, a Hall sensor changes an electrical output in response to magnetic field. In a throttle, a magnet moves relative to the sensor.

A simple three-wire Hall throttle often has:

supply	regulated voltage provided to the sensor
ground/reference	reference for supply and signal
signal	output voltage that changes with position

Constructed example: 5.0 V supply, 1.0 V signal at rest, about 2.5 V mid-command, and 4.0 V at full command.

Those are teaching values. Real throttle ranges vary, and controllers may deliberately reject signals outside their expected window.


The signal does not carry motor power

At full throttle the signal may still be only a few volts and tiny current.

The control circuit reads it, chooses PWM duty cycle, and commands the MOSFET.

Hall throttle → control IC → PWM → MOSFET → motor

Power and information meet at the switch



Arrows show roles and flow, not individual wires or a complete wiring diagram.

Figure 9. A functional map separates the motor power path from the low-voltage supply and command path. All required returns and protection are not shown.

The rider's thumb reports intent. The power stage does the work.

A controller as functional blocks

Power path:

battery → protection → MOSFET power stage → motor → return

Control-power path:

battery → regulator/converter → low-voltage control supply

Command path:

throttle/sensor → control IC or microcontroller → PWM/gate command

Later controllers add current sensing, temperature sensing and communications. The split is already visible.

Measure, decide, command

Once an IC or microcontroller exists, the controller can compare measured current with a limit, watch temperature, refuse operation below a voltage threshold, or ramp commands.

measure → decide → command

That loop is useful for troubleshooting. Ask which stage stopped making sense.

Answered diagnostic case: zero throttle signal

A 24 V brushed scooter has a healthy 24.6 V battery but no motor response.

The Hall throttle is expected to receive 5.0 V between supply and ground. Measurement shows 0.0 V. Its signal is also 0.0 V.

You might condemn the throttle because its signal is zero. That would be premature. A sensor without its required supply cannot be expected to produce a meaningful output.

The next question moves upstream: why is the 5 V control supply missing?

If the regulator input has battery voltage but its output is 0 V, the regulator or something loading that rail, a shared supply line held near a particular voltage, becomes a strong suspect. If regulator input is also 0 V, move farther upstream.

One measurement changed the diagnostic direction because we understood the block before replacing the part.

A valid signal does not guarantee motion

Suppose instead the throttle gets 5.0 V and its signal changes smoothly from 1.0 V to 4.0 V.

That supports the conclusion that the throttle is behaving as expected in our model. It does not prove the controller must run the motor.

A brake input, fault state, low-battery condition, overtemperature state, missing enable or internal failure may still block output.

A good measurement at one boundary lets us move to the next boundary.

We have used a lot of black boxes

Resistors, potentiometers, capacitors, coils, diodes, relays, BJTs, MOSFETs, Hall sensors, regulators, buck and boost converters, and ICs have all appeared only when a problem needed them.

Now there are enough reasons to care about what they physically are.

Next: open the parts

Chapter 10 asks what these components are made from and why their construction produces their behavior.

Copper, resistive films, dielectric layers, wound wire, magnetic cores and doped silicon will collapse a huge parts catalog into a handful of physical ideas.

By now the reader has enough electrical language that the book can become denser without becoming cryptic.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [13], [14], [15], [17], [21]. Full titles and links appear in the source list.

CHAPTER 10

10. Components are materials behaving predictably

We have spent nine chapters meeting components only when a problem demanded them. A resistor appeared because current needed restricting. A diode appeared because a coil needed somewhere to put its current. A capacitor appeared because voltage was changing too quickly. A MOSFET appeared because a relay could not switch fast enough.

Now we have enough reasons to ask what these things actually are.

The useful surprise is that hundreds of component types are built from a much smaller set of physical ideas: conductive materials, resistive materials, insulators, electric fields, magnetic fields, semiconductor junctions and mechanical motion.

what it is made from → what that does electrically → useful behavior → limitation → common failure

The conductive and resistive family

Copper conducts well because its structure lets electrical charge carriers move comparatively easily. Aluminum does too, although its conductivity, oxide layer, mechanical behavior and joining requirements differ.

A conductor is not a zero-resistance path. It is a material whose resistance is low enough to make charge movement useful.

same material: longer conductor → more resistance

same material: smaller cross-section → more resistance

That is why a long thin wire and a short thick wire made from the same copper are not the same electrical component.

A resistor is a conductor chosen to resist

A resistor deliberately combines material and geometry to produce a useful amount of resistance.

A metal-film resistor may use a thin resistive film deposited on an insulating body, with a spiral cut making the current path longer. A carbon resistor uses carbon-based resistive material. A high-power wirewound resistor uses resistance wire wound around a heat-tolerant core.

Construction changes tolerance, power handling, noise, unwanted inductance and failure behavior.

A heater is the same broad physics used for another purpose: resistive heating is the desired output rather than a loss.

The potentiometer belongs here too

A potentiometer commonly uses a carbon, conductive-plastic or similar resistive track. Its movable wiper touches that track.

As a variable resistor, moving the wiper changes how much resistive path is used. As a voltage divider, moving it changes the proportion of the track on either side of the wiper.

Its construction explains its limits: the track can overheat, the wiper has finite contact area, and physical contact can wear or become intermittent.

Temperature is part of resistance

Copper resistance generally rises as copper gets hotter over ordinary operating ranges. A motor winding measured cold will therefore usually have more resistance when hot.

Some materials are deliberately chosen because resistance changes strongly with temperature. In one circuit that change is an error; in another it becomes a temperature sensor.

A recurring electronics lesson is hiding here: an unwanted behavior can become a useful component when we design around it.

The electric-field family

A capacitor is fundamentally two conductive regions separated by an insulating material called a dielectric.

dielectric	insulating material between the conductive regions
capacitance	charge-separation response for a given voltage
F	farad, the unit of capacitance

Apply voltage and opposite charge accumulates on the conductors. An electric field exists in the dielectric. Energy is stored in that field.

That is why a capacitor can accept charge during a fast voltage change and release charge later. It does not store current.

Construction changes the capacitor

A multilayer ceramic capacitor stacks many thin ceramic dielectric layers with metal electrodes, producing large effective area in a small package.

A film capacitor uses thin plastic film as dielectric with conductive layers or metallization. Film construction is often useful where low loss, pulse handling or stable behavior matters.

An aluminum electrolytic capacitor uses a very thin oxide dielectric on aluminum and an electrolyte as part of the electrode system. It can provide large capacitance in a small volume.

Many electrolytics are polarized. Reverse voltage can damage the dielectric system and lead to leakage, heating, venting or catastrophic failure.

Real capacitors have losses

A real capacitor leaks a little, has some inductance and has some resistance.

ESR means equivalent series resistance. It represents resistive loss that behaves as though resistance were in series with the ideal capacitance.

Current through ESR makes heat. Aging and electrolyte loss can raise ESR in electrolytic capacitors, so a capacitor may still show capacitance yet perform badly in a power circuit.

Open, shorted, leaky and high-ESR are useful failure classes, not diagnoses made from appearance alone.

The magnetic family

Wind insulated wire into turns and pass current through it. The turns create a magnetic field.

inductor ↔ relay coil ↔ transformer winding ↔ motor winding

Their jobs differ because geometry, number of turns, core material, multiple windings, air gaps and mechanical arrangement differ. The underlying physics is shared.

Induction, revisited properly

Changing current changes the magnetic field around a conductor. A changing magnetic field linked with a conductor generates voltage in that conductor. That is electromagnetic induction.

The generated voltage acts in a direction that opposes the change that produced it.

That connects several earlier observations. An inductor fights a rapid rise in current. A collapsing relay or motor field produces flyback. A spinning motor generates back-EMF.

They are not identical events, but they belong to the same physical family.

Turns, cores and saturation

More turns can increase inductance and magnetic coupling, but they also require more wire, adding resistance, mass and size.

A magnetic core can concentrate magnetic flux, a measure of how much magnetic field passes through a chosen area, and greatly change inductance. Core material is chosen for the frequency, field strength, loss and power involved.

Many magnetic materials eventually approach saturation. Past that region, additional current produces much less additional magnetic effect than the simple model predicts.

Current can then rise rapidly while useful magnetic gain does not. More current is not always proportionally more useful magnetism.

Transformers are coupled coils

A transformer uses at least two windings linked by a common changing magnetic field.

Changing current in the first winding creates changing magnetic flux. That changing flux induces voltage in the second winding.

A steady, unchanging DC current does not continuously transfer energy through an ordinary transformer in the same way because after the initial change the field is no longer changing.

This is why transformers naturally belong to AC and switched systems.

Magnetic parts can fail electrically and magnetically

Windings have copper resistance, so they make I^2R heat. Magnetic cores also lose energy as their magnetic state changes repeatedly.

A shorted turn can create a low-resistance loop carrying large induced current, producing heat while weakening the intended magnetic behavior.

An open winding stops the intended current path. Saturation can make current rise sharply. Heat can damage insulation and create still more winding faults.

The semiconductor family

A semiconductor occupies a useful middle ground between conductor and insulator because its conductivity can be deliberately controlled.

Silicon is the most familiar semiconductor material in ordinary electronics.

Doping is deliberate contamination

Manufacturers add carefully controlled amounts of other atoms to regions of silicon to change the available charge-carrier behavior. This is called doping.

One broad region type is called N-type and another P-type. We do not need band theory here. The practical consequence is that differently doped regions can be arranged to create useful junctions and controllable conduction.

The diode grows out of a junction

A basic semiconductor diode can be made from a P-type region joined to an N-type region.

Charge rearrangement at that junction creates an internal electric-field region. Applied polarity changes how easily charge crosses it.

When voltage is applied in the diode's conducting direction, called forward bias, substantial current can flow. Applied in the blocking direction, called reverse bias, very little flows under normal rated conditions.

That is the physical basis behind Chapter 6's one-way-door model.

The door is not perfect: forward voltage, leakage, current rating, reverse-voltage rating, switching speed and temperature all matter.

The BJT uses junctions to control current

A bipolar junction transistor arranges multiple doped regions so base drive can control a larger collector-emitter current.

It is not a microscopic relay. Its behavior comes from charge movement through semiconductor junctions.

The MOSFET uses an electric field to create a path

A MOSFET's gate is insulated from the main semiconductor channel by a very thin insulating layer.

Apply the appropriate gate-to-source voltage and the electric field changes the charge distribution in the semiconductor beneath the gate. That creates or strengthens a conductive channel between source and drain.

Now Chapter 8 makes physical sense: the gate does not feed motor current. Its electric field controls the conductivity of another path.

The insulated gate also creates a vulnerability. Excessive gate voltage or a discharge of static electricity can damage that thin structure.

Hall sensors turn magnetic field into information

Moving charge carriers in a material respond to magnetic field. A Hall-effect sensor turns that small field-dependent effect into a useful electrical signal.

That lets a throttle move a magnet past a stationary sensor without a rubbing electrical wiper.

The same broad principle can report rotor position, wheel speed or current, depending on the surrounding design.

An IC is a manufactured neighborhood

An integrated circuit combines many semiconductor structures, conductive paths, resistive elements and sometimes capacitive structures on one die.

A simple regulator may contain relatively few structures. A microcontroller may contain enormous numbers of transistors plus memory and interface circuitry.

The package protects the die and provides external connections. Internally, different pins can lead to structures with very different voltage and current limits.

The electromechanical family


Some components deliberately cross between electricity and motion.

A switch moves contacts. A relay uses magnetism to move contacts. A solenoid moves a plunger. A speaker moves a cone. A brushed motor keeps a rotor turning.

Their electrical and mechanical failures overlap. Contacts wear and arc. Springs weaken. Bearings develop drag. Coils heat. Insulation breaks down. Moving parts seize.

Real components mix families

Construction creates behavior



A real motor combines all three, plus moving parts and insulation.

Figure 10. Resistance, capacitance and inductance come from physical construction. Components often combine these effects.

The family labels are thinking tools, not walls.

A motor contains conductive copper, magnetic materials, electrical insulation, bearings and often mechanical contacts.

An electrolytic capacitor combines conductive foil, dielectric oxide and electrolyte.

A MOSFET package contains a semiconductor die, metal conductors, insulating material and a thermal path to the outside world.

Understanding a component often means asking which family is responsible for the behavior you care about.

Construction predicts failure

A resistor may drift in value, burn or go open.

A capacitor may become leaky, high-ESR, open or shorted.

An inductor or transformer may develop an open winding, shorted turns, saturation-related overheating or insulation failure.

A diode may fail short, open or leaky.

A BJT or MOSFET may fail short, open, leaky or thermally damaged.

A switch, relay or motor can add contact wear, bearing problems and mechanical seizure.

You might see one symptom associated with one of these failures. That symptom is evidence, not proof. Chapter 15 will turn these categories into a formal failure vocabulary.

Open Motor M

We can now look inside the little brushed motor that has carried the first half of this book.

There is conductive copper in the windings. There are magnetic materials in the stator and rotor system. There is insulation keeping adjacent turns from becoming one shorted lump. There are bearings supporting the shaft. There are brushes and a segmented commutator making sliding electrical contact.

Every one of those pieces now belongs to a physical family we understand.

One question remains: why are the brushes and commutator there at all?

Next: the motor has to keep changing its mind

A magnetic rotor does not keep producing useful torque merely because we energize one winding forever.

The current path must change at the right rotor position so the magnetic relationship keeps pulling or pushing in the desired direction.

That switching process is commutation.

Chapter 11 shows how brushes and a copper commutator perform that timing mechanically, why the arrangement works so well, why it wears and arcs, and what problem we create when we remove it.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [10], [12], [13], [16], [17], [18], [25]. Full titles and links appear in the source list.

CHAPTER 11

11. Why brushes exist, and why we got rid of them

Motor M has been turning since Chapter 1. We have measured it, loaded it, stalled it on paper, explained its back-EMF, controlled it badly, controlled it better, and finally opened its materials.

Now we ask the question we deliberately postponed: why does a brushed motor need brushes at all?

The answer is timing. Magnetic force can start a rotor moving. Continuous rotation requires the motor to keep changing which magnetic relationship it creates as the rotor moves.

One energized coil does not make a useful continuous motor

Put a current-carrying coil in a magnetic field and forces can create torque on the rotor.

But imagine energizing one fixed magnetic relationship forever. The rotor turns toward the position favored by that field. As it approaches alignment, useful torque falls. Pass the favorable position without changing anything and the old magnetic relationship can begin fighting the rotation.

So a rotating motor needs more than magnetism. It needs magnetism at the right time.

Meet the parts that solve the timing problem

stator	the stationary magnetic structure of the motor
rotor / armature	the rotating assembly; in a simple brushed permanent-magnet (PM) motor it carries windings and commutator
armature winding	insulated copper wire on the rotor that carries current and creates magnetic field
shaft	mechanical member that carries rotor torque outward
bearings	support the rotating shaft while allowing low-friction motion
brushes	stationary conductive contacts, commonly carbon-based, that rub on the commutator
commutator	segmented conductive structure attached to the rotor that redirects winding connections as it turns

In our simple permanent-magnet brushed motor, the stator provides a fixed magnetic field. The rotor windings create a controllable magnetic field. The brushes and commutator decide which rotor winding connections receive current as the shaft turns.

The commutator is a rotating switch

Picture conductive copper segments arranged around the rotor shaft but electrically insulated from one another.

Different armature-winding connections terminate at different segments.

Stationary brushes press against the rotating segments.


As the rotor turns, a brush stops touching one segment and begins touching the next. The electrical connection to the armature therefore changes automatically with rotor position.

That is the essential trick.

rotor position changes → commutator connection changes → winding current path changes

The motor uses its own rotation to time its electrical switching.

The brushes stay still; the copper segments rotate



A and B are the two insulated commutator segments connected to the coil.

One-coil conceptual view. Real armatures commonly use many coils and segments.

Figure 11. In this simplified split-ring view, a half-turn swaps which coil terminal meets each brush. The coil connection reverses as the rotor turns, sustaining torque in the permanent-magnet field.

That switching process has a name

The deliberate redirection of winding current so torque continues in the desired direction is called commutation.

In Motor M, commutation is mechanical. The brushes and segmented commutator perform it without a microcontroller, Hall sensor or external timing signal.

The word sounds more complicated than the mechanism once you have seen the need for it.

Walk through one deliberately simplified turn

Real brushed motors may have many armature windings and many commutator segments. We will use a stripped-down conceptual motor to see the sequence.

Step 1: one brush connects source current into a commutator segment. Current passes through an armature winding and leaves through another segment and brush.

Step 2: the energized winding creates a magnetic field. Its interaction with the stator field produces torque. The rotor moves.

Step 3: because the commutator is attached to the rotor, the conductive segments move under the stationary brushes.

Step 4: near the position where the old magnetic relationship would stop helping, the brush transitions to another commutator connection.

Step 5: the armature current path changes. The rotor's magnetic relationship changes with it.

Step 6: torque continues in the same rotational direction.

Then the sequence repeats many times per revolution.

The commutator is not merely delivering electricity to something that happens to spin. It is using rotation itself to schedule the electrical events that sustain rotation.

Why the brushes stay still

The external wires cannot conveniently twist around the shaft forever. The brushes provide stationary electrical connections while the commutator rotates beneath them.

Carbon-based brush materials are useful because they conduct, tolerate sliding contact, can wear sacrificially, and can be shaped to maintain contact with the commutator.

Springs commonly push the brushes against the commutator so contact pressure remains adequate as the brushes wear.

Reverse the battery

Chapter 1 reversed Motor M by reversing the source polarity.

Now we can explain that more physically.

In a simple permanent-magnet brushed motor, the stator's permanent magnetic field stays pointed the same way. Reverse the battery and armature current reverses. The rotor's electromagnetic relationship to the stator field reverses, so torque reverses.

The commutator still performs its timing job, but the resulting rotation proceeds in the opposite direction.

This explanation applies to the simple permanent-magnet brushed motor we have been using. Other brushed-motor field arrangements can respond differently because reversing supply may reverse both field and armature current.

Back-EMF finally has hardware

Chapter 3 said a spinning motor is also acting as a generator.

Now identify the generator physically: copper armature conductors are moving through magnetic field.

That motion induces voltage in the windings.

The generated voltage opposes the applied drive voltage during normal motoring operation. Faster rotation generally produces more back-EMF for the same motor construction and field.

faster rotor → more back-EMF → less effective winding voltage at the same applied voltage

Load the shaft and it slows. Back-EMF falls. More of the applied voltage is left to drive winding current. Current rises and electromagnetic torque increases.

The behavior from Chapter 3 is no longer an abstract correction term. It is happening in the copper and magnetic field we can point to.

Brushes are brilliant and annoying

Mechanical commutation is wonderfully self-contained. Give a simple brushed permanent-magnet motor two wires and DC voltage and it handles its own timing.

That simplicity is why brushed motors remain useful.

But the timing mechanism requires sliding electrical contact.

Sliding electrical contact has a price

The brush-to-commutator interface has resistance. Current through that resistance makes heat.

The brush physically rubs the commutator, so both surfaces wear.

As a brush crosses between segments, current is being interrupted and redirected in an inductive winding. Arcing can occur.

Arcing erodes surfaces, creates electrical noise and adds heat.

Brush material wears into dust. Springs and brush holders have mechanical limits. Commutator segments can become rough, burned, contaminated or uneven.

At higher rotational speed, the available time for clean commutation becomes shorter while mechanical and electrical stresses increase.

None of this means brushes are bad engineering. It means their simplicity is purchased with wear and contact physics.

Not all motor heat comes from the winding

Chapter 3 emphasized copper I^2R loss because it explained stall heating. A real brushed motor has more loss mechanisms.

Winding resistance creates copper heat.

Brush and commutator contact create electrical and frictional loss.

Changing magnetic fields create losses in magnetic materials.

Bearings and seals create mechanical friction.

Air movement creates drag called windage.

Those losses eventually become heat.

Power loss and temperature are still different quantities. Temperature depends on how quickly heat is produced, stored and removed.

A motor can fail one segment at a time

Because the commutator connects different armature sections as the rotor turns, some electrical faults can depend on rotor position.

Suppose one armature winding goes open.

The motor may run poorly or stop at a position where the brushes connect to the damaged path. Move the shaft slightly and another winding path may allow it to start.

That creates the classic idea of a dead spot.

A dead spot does not prove an open armature winding. Brush contact, commutator damage and other faults can produce similar symptoms. But rotor-position dependence is useful evidence.

Common brushed-motor clues

Worn or sticking brush: intermittent operation, weak torque, excessive arcing or failure to start can occur.

Dirty or damaged commutator: arcing, noise, uneven operation or dead positions can occur.

Open armature winding: position-dependent no-run or weak operation can occur.

Shorted turns: excess current, heating and reduced torque can occur.

Bearing damage: drag, noise, heat and increased current can occur because the mechanical load inside the motor has increased.

These are clues, not one-to-one diagnoses. The same symptom can have several causes.

A useful mechanical test changes an electrical variable

Turn an unpowered brushed motor by hand.

If the bearings are healthy and the motor is otherwise normal, you expect a certain smooth mechanical resistance plus the characteristic feel of the magnetic structure and brushes.

A seized or rough bearing adds mechanical drag.

When powered, that extra drag becomes additional motor load. Chapter 3 tells us what additional load tends to do: speed falls, back-EMF falls and current rises.

A mechanical bearing problem can therefore present as an electrical overcurrent or heating problem.

The machine does not respect our departmental boundaries.

Why not put the windings on the stationary part?

If the windings remain stationary, we no longer need sliding contacts merely to deliver power into a rotating armature.

Put permanent magnets on the rotor and wind the electromagnetic coils on the stator.

Now the high-current windings can connect directly to stationary electronics.

That removes the brush/commutator wear mechanism.

But it also removes the automatic rotating switch that was doing commutation for free.

Remove the brushes and we lose the timing

Imagine Motor M with its brushes and commutator gone.

We still have magnetic materials. We still have copper windings. We can still create torque.

But nothing now changes the winding-current pattern automatically as the rotor moves.

An external circuit has to know, or infer, where the rotor is and energize the correct winding relationship at the correct time.

We traded mechanical commutation for electronic commutation.

That trade is the doorway to the modern scooter motor

Electronic commutation can avoid brush wear, reduce contact arcing, move heat-producing windings to the stationary structure, and give a controller much more authority over torque and speed.

It also requires more wires, more switching devices, position information or estimation, and considerably smarter electronics.

A two-wire brushed motor could hide its timing inside the can.

A brushless motor hands the timing problem to the controller.

Next: three wires are not positive, negative and ground

Chapter 12 starts with three stator winding groups and a permanent-magnet rotor.

The controller energizes those windings in sequence so the magnetic field appears to rotate.

That gives us the three fat phase wires found on a scooter hub motor.

Then we add Hall sensors or back-EMF sensing so the controller knows when to change the sequence.

By the end, six-step commutation, sine control and field-oriented control (FOC) will all be recognizable as increasingly sophisticated answers to the same question the copper commutator answered mechanically:

Where is the rotor, and which magnetic field should we create now?

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [2], [3], [10], [19]. Full titles and links appear in the source list.

CHAPTER 12

12. Three wires, three magnetic fields

Chapter 11 removed the brushes and commutator from Motor M. That eliminated sliding electrical contacts, but it also removed the mechanism that had been deciding when to redirect winding current. The timing problem did not disappear. We handed it to electronics.

A modern scooter hub motor commonly gives the controller three heavy motor wires, later called phase wires, and, on sensed systems, several smaller signal wires. The three heavy wires are not positive, negative and ground. They are three doors into a magnetic sequence.

Start with three winding groups

Imagine three groups of stator windings arranged around a permanent-magnet rotor. Call them A, B and C.

Energize a winding relationship and the resulting stator magnetic field produces torque on the rotor.

If we leave that relationship unchanged, the rotor approaches the favored alignment and useful continuous torque disappears.

So the controller changes which windings carry current and in which direction as the rotor moves.

A, B, C winding currents change in sequence → stator magnetic field rotates → rotor follows

The controller has become the commutator.

What phase means here

Phase is one of those words that becomes frightening because it is often introduced after several pages of mathematics.

Here it means position within a repeating electrical pattern.

The three winding groups are driven with electrical patterns displaced from one another. As one phase's contribution changes, the others are at different points in the sequence.

The result is a magnetic field whose useful direction moves around the stator.

The three fat wires

The three heavy motor wires are usually called phase wires.

Do not assign them permanent jobs such as positive, negative and ground.

At one instant the controller may drive current into phase A and out through phase B while phase C is not actively driven. A moment later the roles change.

Current direction in a phase lead can reverse during normal operation.

This is why randomly connecting a three-phase motor directly to a DC battery is not equivalent to powering a two-wire brushed motor. The motor needs a timed switching pattern.

A deliberately simple six-step sequence

One common basic commutation strategy is six-step commutation.

In a simplified conceptual sequence, two phases participate strongly in each step while the third is left unpowered or floating. Floating means that phase is temporarily not actively driven toward either side of the controller's main DC supply, called the DC bus.


One possible abstract sequence is:

Step 1	A driven positive, B driven negative, C floating
Step 2	A positive, C negative, B floating
Step 3	B positive, C negative, A floating
Step 4	B positive, A negative, C floating
Step 5	C positive, A negative, B floating
Step 6	C positive, B negative, A floating

Then the sequence repeats.

Positive and negative here describe which side of the DC bus the controller is connecting toward during that switching state. They are not permanent identities of the phase wires.

Real winding topology, PWM strategy and commutation implementation vary. The six rows are a teaching sequence that exposes the idea, not a universal pinout.



One electrical cycle need not equal one mechanical revolution.

Figure 12. One conceptual six-step order. Rotor position determines when to advance. The letters are phase names, not wire colors or a universal Hall code.

Why six steps make rotation

Each current combination creates a stator magnetic field pointing in a different effective direction.

Move from one step to the next and that preferred magnetic direction moves around the stator.

The permanent-magnet rotor is pulled and pushed along by that changing field.

The controller therefore does electronically what the segmented copper commutator did mechanically: change the current path at the right rotor position.

The controller needs to know when

A brushed motor used the rotor's physical commutator position to schedule switching automatically.

An electronic controller needs position information or a way to infer it.

One common solution is Hall-effect sensors.

Hall sensors report rotor position

Chapter 10 explained the Hall effect at the material level: magnetic field can be converted into an electrical signal.

Place Hall sensors in the motor where the rotor magnets pass them. As magnetic polarity and field change with rotor position, the Hall outputs change state.

A common sensed three-phase motor may use three Hall signals. Together, their changing HIGH/LOW pattern, the two logic voltage states, gives the controller coarse rotor-position information.

The controller uses that pattern to decide which commutation step belongs next.

The small Hall wires

A typical sensed motor needs more than the three Hall signal wires because the sensors themselves need power.

A common arrangement therefore includes:

Hall supply	regulated voltage supplied by the controller
Hall ground/reference	reference for Hall supply and signals
Hall A	position signal
Hall B	position signal
Hall C	position signal

That often totals five small wires.

Do not trust wire color as proof of function. Manufacturers can choose different colors, connector pinouts and supply voltages. Identify by documentation or measurement.

A Hall pattern is information, not motor power

The Hall wires do not energize the phase windings.

They report rotor position to the controller using low-power signals.

The phase wires carry the substantial winding current.

Hall signals tell → controller decides → phase switches act

This is the power/control separation from Chapters 7 through 9 appearing again inside a brushless motor system.

Sensorless control listens to the motor itself

A spinning permanent-magnet motor generates back-EMF in its windings.

During some six-step strategies, one phase is not actively driven. The controller can observe voltage behavior on that floating phase and use its back-EMF to infer rotor position.

That can eliminate Hall sensors.

There is an obvious difficulty: at zero speed there is essentially no rotational back-EMF to observe.

Sensorless controllers therefore need some startup strategy before back-EMF becomes strong enough for reliable position estimation.

This is why a motor can behave differently at startup than it does once spinning.

Battery current and phase current are not the same number

This distinction matters enormously in scooter modifications.

The battery feeds the controller's DC bus. The controller then switches that bus among the motor phases using PWM and commutation.

Battery current describes current entering the controller from the battery.

Phase current describes current flowing in the motor phase paths.

Those numbers can be very different.

At low motor speed and partial PWM duty, the controller can circulate substantial phase current while drawing a lower average current from the battery.

Energy conservation still applies. The controller is not creating free power. Voltage, current, duty cycle, stored magnetic energy and losses are being transformed through switching.

A simple power comparison

Suppose a controller draws 10 A from a 48 V battery at one operating point.

$$\text{Battery-side electrical input} = 48 \text{ V} \times 10 \text{ A} = 480 \text{ W}$$

Ignoring losses for a moment, the motor-side electrical power cannot continuously exceed that 480 W just because a phase-current measurement says 30 A.

The phase voltage waveform is not simply 48 V DC across one winding while 30 A flows continuously.

You might have expected 30 A phase current to imply $48 \text{ V} \times 30 \text{ A} = 1,440 \text{ W}$. That multiplication uses quantities that do not describe the same continuous electrical condition.

This is exactly why battery current and phase current must remain separate concepts.

Why high phase current matters at low speed

Torque in a permanent-magnet motor is strongly related to torque-producing winding current.

At low speed, back-EMF is low. The controller can use PWM to regulate substantial phase current without applying the full bus continuously.

That gives strong starting torque.

It also creates copper I^2R heating in the motor and conduction/switching losses in the controller.

A scooter can therefore be thermally stressed during a slow hill climb even though it is nowhere near top speed.

Six-step is not the only way to draw the magnetic circle

Six-step commutation changes winding states in discrete steps. It is understandable and effective, but the resulting magnetic field and torque contain more abrupt changes.

A controller can instead shape phase currents more smoothly, approximating sinusoidal waveforms.

Sine-wave control can reduce repeating variation in torque, called torque ripple, along with acoustic noise and harshness compared with crude block commutation, depending on the motor and implementation.

The goal has not changed: create the right magnetic field at the right rotor position.

FOC: stop thinking only in phase names

Field-oriented control, usually shortened to FOC, is a more sophisticated way to control the motor's magnetic field.

Instead of treating A, B and C merely as a sequence of switches, FOC uses rotor-position information and phase-current information or estimates to control the effective stator magnetic field relative to the rotor.

One useful mental model is that the controller separates current into components associated with producing torque and controlling magnetic-field relationship, then commands the three physical phases to create the desired result.

The real mathematics uses coordinate transformations and feedback control. We do not need those equations yet to understand the purpose.

know rotor position + know/control current → place magnetic field where it produces the desired torque
FOC is not a different kind of motor electricity. It is a smarter way of commanding the same windings.

Why FOC needs good information

If rotor position is wrong, current measurement is wrong, motor parameters are badly estimated, or switching timing is poor, the controller's model of the magnetic relationship can be wrong.

That can produce weak torque, roughness, noise, excess current or failure to start.

This is why configuration matters on programmable controllers. Software settings eventually become physical current in copper.

Answered diagnostic case: one Hall signal never changes

A sensed three-phase scooter motor has three heavy phase wires and five small Hall wires.

With the controller providing the expected Hall supply, slowly rotate the wheel by hand while measuring each Hall signal relative to Hall ground.

Hall A toggles between about 0 V and 5 V. Hall B also toggles. Hall C remains at 0 V through a full slow rotation.

You might immediately declare Hall sensor C dead.

The measurement does tell us the C channel is not reaching the controller as a changing position signal. It does not yet prove the semiconductor sensor itself is the failed part.

A broken Hall C wire, bad connector pin, short to ground, a missing resistor that holds the signal HIGH when released, called a pull-up, or a failed sensor could all produce a stuck-low reading.

The next test separates the sensor output from the harness and receiving input, with power removed while connections change. Some Hall outputs are open-drain: the sensor can pull the signal LOW but needs a pull-up resistor to bring it HIGH when released. A pull-up is simply a resistor to the signal supply that provides that default HIGH state. If the isolated sensor toggles correctly but the isolated harness has a short or open, wiring is implicated. A short can hold an entire connected signal line low, so merely moving the probe may not separate those possibilities.

Again: measurement narrows the map. It does not reward the fastest guess.

Phase-order problems can still make the motor move

On some systems, an incorrect relationship among phase wiring and Hall timing can produce a motor that turns but runs rough, draws excessive current, starts badly or spins in the wrong direction.

The fact that a wheel rotates is therefore not proof that commutation is correct.

Correct operation means the controller's assumed rotor position and its commanded phase-current pattern agree with the actual motor.

This becomes especially important when mixing motors and controllers that were not designed as a matched system.

We now know what the controller is trying to accomplish

The three-phase controller is no longer a mysterious box that 'sends power to three wires.'

It has a precise job:

observe or estimate rotor position → choose magnetic-field command → switch phase currents →
measure what happened → repeat

At high switching rates, this loop can happen thousands of times per second.

The remaining mystery is the hardware that survives doing it.

Next: open the controller

Chapter 13 follows battery power through protection and DC bus capacitors, then into three paired high-side/low-side MOSFET switching legs, which it will name half-bridges.

We will follow the low-voltage path separately through regulators, the microcontroller and the circuits that drive the MOSFET gates.

Then Hall inputs, current sensors, temperature sensors, brake inputs and communications join the picture.

By the end, current limiting, regenerative braking, field weakening and FOC will no longer be menu options. They will be things identifiable blocks of hardware and software are physically doing.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [19], [20]. Full titles and links appear in the source list.

CHAPTER 13

13. Inside a real scooter controller

Chapter 12 reduced a three-phase controller to a repeating job:

observe rotor → decide magnetic field → switch phase current → measure → repeat

Now open the box. Do not begin by identifying every part. Find the neighborhoods. Follow power. Follow information.

Two maps occupy the same board

Power path:

battery → protection → DC bus → power switches → motor phases

Control path:

low-voltage supply → sensors/commands → control logic → gate drivers → power switches

The maps meet at the MOSFET gates and feedback sensors. Keeping them separate in your head makes a crowded board readable.

Battery power enters first

The controller receives the battery's actual terminal voltage, not merely the nominal number printed on the scooter.

A lithium battery's voltage changes with state of charge. The controller must survive the real maximum and minimum voltages of the intended pack.

Near the input may be a fuse, reverse-polarity protection, transient suppression, a contactor, an electrically controlled switch for a high-current path, or other protection. Do not assume a feature exists because it would be sensible. Identify what is actually present.

The DC bus

The controller establishes positive and negative power rails feeding the switching stage. This shared distribution path is the DC bus.

Heavy copper traces, bus bars and large conductors often reveal it because it carries substantial current.

The large capacitors are working

Large electrolytic capacitors commonly sit across the DC bus.

The battery and its wires have resistance and inductance. The MOSFET bridge demands current in rapid pulses. Bus capacitors provide local energy during fast switching and help restrain rapid bus-voltage movement.

Their voltage rating therefore matters directly. A capacitor rated below the voltage the bus can actually reach is not protected by a lower nominal label.

Why plugging in can spark

An uncharged bus capacitor initially has little voltage across it. Connect a charged battery and current rushes in to charge it.

That is inrush current.

A connector spark on some higher-energy systems can be the result of abruptly charging controller capacitance. The spark is evidence of an energetic transient, not by itself proof of a defective controller.

Precharge makes the meeting gentler

Some systems use precharge. The bus capacitors are first charged through a resistor or another current-limited path.

first: battery → limiting path → bus capacitors

then: battery → main low-resistance path → controller

Once bus voltage approaches battery voltage, closing the main path produces much less inrush. Precharge is not universal in small scooters, but the problem grows with voltage and capacitance.


Three phases need three switching legs

Each motor phase must be controlled toward either side of the DC bus.

A half-bridge uses a high-side switch between bus positive and a phase node, and a low-side switch between that phase node and bus negative.

Three phases require three half-bridges.

Three half-bridges share one DC bus



Never turn both switches of one leg ON together. Body-diode paths omitted.

Figure 13. Each phase node sits between a high-side and low-side switch. Blocks represent controlled power switches; gate drivers, body diodes, the diodes inherent in MOSFET structure, and protection require their own design. In the drawing, bus C means the DC-bus capacitor.

High-power controllers may parallel several MOSFETs in each switch position. Functionally, they are still high-side and low-side switches.

Shoot-through: the controller shorts itself

Turn the high-side and low-side MOSFET of one leg on strongly at the same time and a low-resistance path exists directly from bus positive to bus negative.

That is shoot-through.

The motor is not the load in this path. The controller is shorting its own bus through its power devices. Current can become enormous and MOSFET failure can be nearly immediate.

Dead time

MOSFETs do not turn off instantaneously. The controller therefore inserts a short gap between turning one device off and the opposing device on.

That gap is dead time.

Too little risks shoot-through. Too much distorts current and can add loss. Software timing has become a physical survival condition.

The microcontroller cannot directly muscle the MOSFETs

A microcontroller output is a low-power logic signal. A large MOSFET gate behaves capacitively and must be charged and discharged quickly.

High-side devices are harder because their source voltage moves with the phase node.

A gate-driver circuit sits between logic and power.

microcontroller command → gate driver → MOSFET gate

A failed gate driver can leave a healthy MOSFET unable to switch correctly.

The controller contains smaller power supplies

The battery bus may be tens of volts while the microcontroller and sensors need much lower regulated rails.

A buck converter or regulator may create 12 V, 5 V and 3.3 V rails as required.

One rail may feed gate drivers, another logic, another Hall sensors or throttle circuitry.

A controller can have full battery voltage at its input and still be brain-dead because one small internal rail is missing.

The microcontroller is the traffic cop

The microcontroller or motor-control IC reads information and produces timing commands.

Inputs can include throttle, brake, Hall position, current, bus voltage, temperature, enable and communications.

It decides how the bridge should switch. It does not carry motor current.

measure → decide → command gate drivers

Current sensing closes the loop

A controller that limits current must measure or estimate current.

A shunt is a deliberately small, accurately known resistance placed in a current path. Measure its voltage drop and Ohm's law gives current.

Example: 1.0 mΩ shunt, 20 mV across it.

$$I = 0.020 \text{ V} \div 0.001 \text{ } \Omega = 20 \text{ A}$$

Other controllers use magnetic current sensors such as Hall-based devices.

Where current is measured matters

A shunt in the battery lead measures battery current. A shunt between the bridge and its local bus capacitor can instead see switching pulses used to reconstruct phase current. Phase shunts measure their own winding paths. Placement and sampling time determine what the reported number means.

Chapter 12 established that those numbers can differ substantially.

Therefore 'the controller measures 30 amps' is incomplete until we know where the measurement lives.

Temperature becomes information

A thermistor, a resistor whose resistance changes strongly with temperature, or another sensor can report controller temperature. Some motors report their own temperature.

The controller can reduce current as temperature approaches a limit. That is thermal derating. It may shut down if temperature continues rising.

The sensor does not cool anything. It gives the controller information so it can stop worsening the condition.

Throttle, brake and enable

The Hall throttle provides a low-power command.

Brake inputs can request zero drive torque, electronic braking or regenerative braking depending on design.

An enable or ignition input may wake the control electronics.

Battery voltage can therefore be present at a controller that remains intentionally inactive.

Displays and communications

A display may exchange modes, settings, faults and other data with the controller.

Universal asynchronous receiver/transmitter (UART) and Controller Area Network (CAN) are examples of communication methods. The agreed rules and message formats used for communication are called protocols; proprietary protocols are common too.

A matching connector does not establish compatibility. Identical-looking pins can carry different voltages, signals and protocols. A shared interface name still does not guarantee that two devices exchange compatible messages.

Current limiting is controlled restraint

If throttle demand would create more current than allowed, the controller measures current and changes PWM/phase commands.

The current limit is not a resistor burning the excess. It is a feedback decision that changes switching.

This is the mature descendant of Chapter 7's terrible speed controller.

Low-voltage cutoff

If battery voltage falls below a configured threshold, the controller may reduce power or stop driving.

A weak battery can measure normally at rest, sag under load, cross the cutoff threshold, then recover when load disappears.

A low-voltage fault therefore does not automatically mean the controller is wrong.

Overvoltage protection

If bus voltage exceeds an allowed range, the controller may shut down or take protective action.

Software protection only works while hardware remains alive. MOSFETs, capacitors, gate drivers and regulators still have physical voltage limits.

Soft start

The controller can ramp current or torque rather than instantly commanding the maximum.

Soft start can reduce tire spin, mechanical shock, battery sag and abrupt current demand.

Again, a software decision changes physical stress.

Electronic braking and regeneration

The controller can command motor current that produces torque opposing rotation.

During regenerative braking, the spinning motor acts as a generator and the bridge directs electrical energy toward the DC bus.

wheel motion → motor generation → controller → DC bus → battery, if accepted

The battery and BMS must permit charging current. A nearly full battery has little voltage headroom, meaning little room for the bus voltage to rise before charging limits are reached, and a BMS may refuse charge.

If generated energy cannot enter the expected sink, bus voltage can rise. Regen therefore belongs in voltage matching, not merely the feature list.

Field weakening

As speed rises, motor back-EMF approaches available bus voltage and leaves less voltage difference to force additional torque-producing current.

Field weakening deliberately changes the commanded magnetic-current relationship to extend speed beyond the ordinary back-EMF-limited region.

The cost can include more current, more loss, reduced torque margin and more heat. It is not free RPM.

FOC has a physical home

FOC lives across the control algorithm, current sensors, rotor-position information, gate drivers and MOSFET bridge.

The algorithm decides the desired magnetic relationship. Sensors report what happened. Gate drivers and MOSFETs create the phase voltages that make the commanded currents happen.

FOC is neither software alone nor a special kind of MOSFET.

Controller matching begins with maximum charged voltage

When considering a battery upgrade, first find the battery's maximum charged voltage.

Then check DC-bus capacitor ratings, MOSFET drain-to-source ratings, gate-driver architecture, low-voltage converter input limits and transient margin.

Then include regenerative braking and switching transients, which can push the bus above resting battery voltage.

A controller that survives cruising can still fail during charging, regen or a transient if margin is inadequate.

Current matching is a system too

Controller current capability depends on MOSFET conduction and switching loss, cooling, shunts, PCB copper, bus bars, phase wires, connectors and controller-software, or firmware, limits.

Motor current capability depends on winding resistance, magnetic limits, phase leads, connectors and cooling.

Battery current capability depends on cells, interconnects, BMS, wiring and connectors.

One impressive current number on one component does not upgrade the rest of the chain.

Answered diagnostic case: battery present, controller dead

A scooter battery measures 52.0 V at the controller input. The controller does not respond to throttle and the motor Hall sensors have no supply.

Expected Hall supply is 5.0 V. Measured Hall supply is 0.0 V.

You might suspect the motor because the Hall sensors are dead. But all three sensors losing supply together points toward their common 5 V source before it points toward three simultaneous sensor failures.

Measure the controller's low-voltage rail. It is also 0.0 V.

Now check the low-voltage converter input. It has 52.0 V.

That localizes the fault to the low-voltage supply section or something shorting its output rail.

The motor has not been proven healthy, but replacing it would not explain the missing controller rail.

Read a board by geography

On an unfamiliar controller, start with the large physical clues.

Large electrolytic capacitors usually sit near the DC bus. Heavy copper and MOSFET banks identify the power stage. Thick phase conductors leave toward the motor.

A smaller switching regulator and inductor may reveal the low-voltage supply. A microcontroller sits in the lower-power logic region. Thin traces run toward Hall, throttle, brake and communication connectors.

You do not need every part number to form a functional map.

identify region → predict expected voltage/signal → measure boundary → narrow fault

The whole controller is now one machine

Battery energy enters the bus. Capacitors steady the local supply. MOSFET half-bridges shape phase current. Gate drivers move the MOSFET gates. The microcontroller decides timing. Sensors report rotor position, current and temperature. Low-voltage regulators keep the decision-making layer alive.

Nothing in that sentence requires magic anymore.

Next: deliberately abuse the assumptions

Chapter 14 changes the rules.

Raise battery voltage. Raise current limit. Spin the motor faster. Ask the controller to regenerate into a full battery. Push phase current until copper heating dominates. Push RPM until magnetic, bearing or mechanical limits appear.

The question will no longer be 'can this motor take 60 volts?'

The useful question becomes:

which limit moves first when I change this variable?

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [13], [18], [19], [20], [21], [22], [27]. Full titles and links appear in the source list.

CHAPTER 14

14. What happens when you change the rules

We finally have enough pieces to talk about overvolting without reducing it to folklore.

A motor is not a voltage number. A controller is not a current number. A battery is not an endless bucket of amps. A machine is a chain of electrical, magnetic, thermal and mechanical limits.

So the useful question is not, 'Can this motor take 60 volts?'

Which limit moves first when I change this variable?

This chapter changes one thing at a time and follows the consequences.

Keep the quantities separate

Battery voltage is the electrical potential difference available at the controller input.

Battery current is current entering the controller from the battery.

Phase current is current in the motor winding paths.

RPM is rotational speed: revolutions per minute.

Torque is twisting action at the shaft or wheel.

Electrical power is the rate of electrical energy transfer.

Mechanical power is torque acting through rotational speed.

Electrical losses increase thermal energy. Heat is energy transferred because of a temperature difference. Temperature tells us how hot something became, not how much power is flowing.

Changing one can change several others, but they are not interchangeable.

Raise voltage first

For a permanent-magnet motor, back-EMF rises with speed. At a given battery voltage, speed eventually reaches a region where back-EMF leaves too little voltage difference to force much additional torque-producing current.

Raise available bus voltage and that ceiling moves upward.

If a motor reaches some no-load speed at 36 V, raising applied voltage to 60 V can make the theoretical unloaded speed tendency rise roughly in the ratio $60 \div 36$, assuming the controller applies it and no other limit intervenes.

$$60 \div 36 \approx 1.67$$

That suggests roughly 1.67 times the unloaded speed tendency, not a promise of 1.67 times road speed.

Aerodynamic drag, load, controller limits, field weakening, losses, battery sag and mechanical limits interfere with the clean ratio.

Higher voltage does not automatically mean more winding heat

Suppose a machine needs 1,000 W of electrical input at one operating point.

At 40 V, ignoring losses:

$$I_{\text{battery}} = 1,000 \text{ W} \div 40 \text{ V} = 25 \text{ A}$$

At 80 V for the same 1,000 W:

$$I_{\text{battery}} = 1,000 \text{ W} \div 80 \text{ V} = 12.5 \text{ A}$$

For the same electrical power, higher battery voltage can mean lower battery-side current. That can reduce I^2R loss in wiring and other parts carrying battery current.

But it does not prove motor phase current fell. Torque still requires winding current, and the controller transforms battery-side conditions into switched phase conditions.

Higher voltage can therefore ease one part of the system while allowing another part to work much harder.

Voltage mostly buys speed headroom; current mostly buys torque

That sentence is simplified but useful.

Raise available motor voltage and useful current can be sustained farther into higher RPM as back-EMF rises.

Raise allowed phase current and available torque rises, subject to magnetic and thermal limits.

Voltage and current are not two knobs that both merely mean 'more power.'

Raise current without raising voltage

Suppose actual phase current through a winding doubles while winding resistance is unchanged.

$$P_{\text{copper}} = I^2 \times R$$

$$(2I)^2 \times R = 4I^2R$$

Copper heating tends toward four times the original value.

Triple actual current and the I^2 term becomes nine times larger.

This is why a motor can sometimes tolerate a dramatic voltage increase at restrained current yet cook quickly when current limits are raised aggressively.

Slow hill climbing can be especially hard thermally: torque demand keeps phase current high while low speed can reduce cooling and prolong the condition.

Magnetic material has a limit

More winding current creates more magnetic field only while the magnetic circuit responds usefully.

As magnetic material approaches saturation, extra current produces progressively less additional useful magnetic effect.

The copper still gets the extra current and makes I^2R heat.

Beyond that region, the trade gets ugly: much more heat for much less additional torque.

Speed creates another family of losses

Keep current conservative but spin the motor much faster.

Copper may not complain first.

Magnetic materials experience more frequent field changes. Core losses rise with frequency and flux behavior.

Bearings and seals run faster. Windage rises. Rotor imbalance becomes more energetic. Magnets, sleeves, adhesives and rotor structures see greater centrifugal stress.

The tire, wheel and bearings are spinning faster too.

Overvolting can move the primary failure mode from copper heating to magnetic or mechanical speed stress.

The controller may die first

A motor winding does not contain a DC bus capacitor rated 63 V. A controller might.

Controller voltage limits come from actual hardware: MOSFET drain-to-source ratings, bus capacitors, gate drivers, low-voltage converter input limits, transient suppressors, PCB spacing and other structures.

Raise battery voltage and all of them see a different environment.

The controller can be the hard voltage ceiling even when the motor could tolerate the extra RPM.

Nominal voltage is not the design voltage

Lithium packs are often described by a nominal voltage lower than full charge.

For cells charged to 4.2 V, a 16-series pack, often written 16S, reaches:

$$16 \times 4.2 \text{ V} = 67.2 \text{ V}$$

That pack may be called a nominal 60 V-class pack.


Now pair it with a controller using 63 V bus capacitors.

You might think 60 V is below 63 V. The actual full-charge voltage in this example is 67.2 V, already above the capacitor rating before switching spikes or regen add anything.

Using > to mean greater than: $67.2 \text{ V} > 63 \text{ V}$.

That is not margin. It is an overvoltage condition.

Compare full charge with the component limit



Full charge already exceeds the rating. Transients need additional margin.

Figure 14. A 16S pack using 4.2 V charge-limit cells reaches 67.2 V. A 63 V capacitor is already over its voltage rating before regen or switching transients.

A rating is not an operating target

Even if maximum battery voltage sits just below a MOSFET or capacitor absolute rating, little room remains for transients, tolerance, temperature effects or regenerative bus rise.

Good design uses margin appropriate to the component, topology and transient environment.

There is no universal percentage that makes every controller safe.

Regen can raise the bus

During regenerative braking the motor generates and the controller directs energy toward the DC bus and battery.

A nearly full battery has little charging headroom. A BMS may limit or disconnect charge current. Wiring and cell resistance require voltage difference to push charge current.

Bus voltage can therefore rise during regen.

A controller that barely survives steady cruising may fail during braking because the extra voltage has nowhere comfortable to go.

The BMS is part of the modification

A higher-voltage or higher-current build has to remain compatible with cell count, maximum charge voltage, discharge current, regen/charge current, BMS switching devices, temperature limits and protection thresholds.

A BMS trip is not generic proof of 'too many amps.' The specific protection condition matters.

The connector is still in the circuit

Chapter 5 still applies.

Higher current through the same connector means more I^2R heating.

Higher voltage increases insulation and arcing demands.

Phase connectors, battery connectors, fuse holders, switches, contactors and wire remain possible bottlenecks.

A power upgrade often finds an ordinary weak connection before it finds an exotic semiconductor limit.

Torque has to react somewhere

A motor turns the wheel one way and reacts against its mounting structure the other way.

Increase phase current and available torque rises. Axle flats, dropouts, torque arms, fasteners, swingarms and frame structure must react it.

A hub motor can electrically survive a current increase and still spin its axle in the dropout.

The electrical modification moved the failure downstream into structure.

Speed moves the rest of the safety problem too

More speed is not merely a larger display number.

Braking energy rises with the square of speed. Tires see greater rotational and road stress. Bearings spin faster. Steering geometry, chassis stiffness, looseness and imbalance become more consequential.

The motor/controller question is only one part of the system.

Worked case 1: 36 V to 60 V, same phase-current ceiling

Take a hypothetical motor/controller system that operates correctly at 36 V. Replace battery and controller with hardware genuinely designed for 60 V-class operation while keeping the same phase-current ceiling.

$$60 \div 36 \approx 1.67$$

Available unloaded-speed tendency rises substantially. The motor can sustain useful torque farther into the RPM range before back-EMF consumes the voltage headroom.

At equal mechanical power, battery-side current can be lower because bus voltage is higher.

But if the rider uses the extra voltage to travel much faster, required mechanical power rises, especially from aerodynamic drag. Battery current may not fall in real use.

Even with restrained current, higher RPM raises magnetic and mechanical losses.

Same current limit controls one major stressor. It does not freeze the rest of the machine.

Worked case 2: same voltage, double phase current

Keep battery voltage unchanged. Raise actual phase current from 20 A to 40 A and suppose the motor reaches it.

$$40^2 \div 20^2 = 4$$

At unchanged winding resistance, copper heating at that current condition tends toward four times the original.

You might expect double current to mean double heat. It does not.

Torque capability can rise strongly, while saturation, controller loss, connectors, axle reaction and battery demand also move.

Worked case 3: the innocent 63 V capacitor

A nominal 60 V-class 16S lithium-ion pack reaches 67.2 V at 4.2 V per cell.

A controller contains 63 V bus capacitors.

Using > to mean greater than: 67.2 V > 63 V.

No further arithmetic is required. The capacitor rating is exceeded at full charge in this example.

A casual '60 V compatible' label cannot answer the engineering question.

Temperature turns guessing into management

When a motor is pushed beyond its original envelope, temperature is one of the most useful measurements available.

A sensor placed where it meaningfully tracks winding or stator temperature can distinguish a harmless short burst from a long hill that is accumulating heat.

The safe temperature depends on winding insulation, magnets, adhesives, bearings and construction. One internet number does not apply to every motor.

Identify material limits, measure the actual machine, and control current or duty before those limits are reached.

Mitigation is mostly refusing to be surprised

Use real voltage margin above maximum charged pack voltage and expected regen/transients.

Set battery and phase-current limits deliberately.

Monitor motor and controller temperature where practical.

Use phase wires, connectors, battery wiring, fuses and contactors suited to actual current.

Check bearing condition and speed capability. Provide cooling where required.

Make sure the BMS can supply discharge current and accept intended charge/regen current.

Strengthen axle retention when torque rises.

Make sure brakes, tires, wheels, chassis and steering belong at the new speed.

These steps do not make every extreme build safe. They turn unknown limits into named engineering questions.

Common ways the experiment ends

MOSFET fails short from voltage, current or thermal stress.

Bus capacitor vents or fails from overvoltage, ripple or heat.

Low-voltage regulator dies because its input limit was exceeded.

BMS trips on voltage, current, charge or temperature protection.

Phase connector melts from current and contact resistance.

Motor winding overheats from phase current and insufficient cooling.

Magnets weaken from heat or retention fails at RPM.

Bearing overheats or mechanically fails.

Hub axle spins in the dropout.

Tire or wheel becomes the speed limit.

These are mechanisms to investigate, not diagnoses to select by vibe.

The truth about overvolting

Overvolting a motor is not automatically abuse of the winding.

Higher voltage primarily gives more speed headroom against rising back-EMF. With current controlled, some motors tolerate far more voltage than their original system used.

But higher speed exposes magnetic and mechanical losses, and higher system voltage stresses controller, insulation, battery, BMS and regen path.

If extra voltage is also used to demand more current, copper heating and torque stress rise quickly.

So the useful answer to 'will this motor take 60 V?' is:

At what current, at what RPM, for how long, with what controller, battery, cooling and mechanical system?

Next: failures stop looking infinite

After fourteen chapters, the number of things that can break seems enormous.

The useful surprise is that their electrical symptoms collapse into a small vocabulary.

A wire opens. A MOSFET shorts. A connector develops resistance. A capacitor leaks. A reference disappears. A signal lies. Heat changes a component. A fault comes and goes.

Chapter 15 names those failure classes so that 'it doesn't work' can become a short list of electrical possibilities instead of a parts cannon.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [2], [3], [13], [18], [20], [21], [23]. Full titles and links appear in the source list.

CHAPTER 15

15. Failure has a small vocabulary

A controller can contain hundreds of components. A scooter can contain thousands of electrical connections and internal structures. That makes failure sound infinite.

Electrically, it is not.

A huge number of broken parts collapse into a surprisingly small set of conditions: a path disappeared, an unwanted path appeared, resistance changed, voltage is wrong, current is wrong, a reference moved, a signal stopped making sense, heat changed something, or the fault comes and goes.

The part names are numerous. The failure vocabulary is small.

Dead is a symptom, not a diagnosis

Suppose a scooter does absolutely nothing when the power button is pressed.

That observation does not tell us whether the battery is empty, a fuse is open, a connector is disconnected, an enable signal is missing, the controller's low-voltage supply failed, or the display stopped communicating.

All it tells us is that the expected chain of events did not complete.

The useful move is to translate the symptom into electrical conditions that could produce it.

Open circuit: the road disappeared

An open circuit is a path that should conduct but no longer does.

A broken wire is an open. A blown fuse is an intentional open created by protection. A disconnected connector is an open. A burned PCB trace can become an open. A BMS may deliberately open its electronic switching path when it detects a fault.

Expected: current can pass through the required route.

Actual: the route is interrupted.

Typical consequence: current downstream of the interruption becomes zero.

But voltage can still appear on one side of the open, which is why 'I measured battery voltage' does not prove the complete path exists.

Short circuit: a road appeared where it should not

A short circuit is an unintended low-resistance connection between points that should not be connected that way.

A damaged wire can short to the frame or another conductor. A MOSFET can fail drain-to-source short. Solder can bridge adjacent pads.

Expected: current follows the designed load path.

Actual: a much lower-resistance path bypasses some or all of that load.

The result can be excessive current, blown protection, heat, loss of control, or a component held permanently on.

A short does not always mean literal zero ohms. It means an unintended path low enough in resistance to materially change circuit behavior.

Unexpected resistance: the road is still there, but it got narrow

Chapter 5 gave us this failure in detail.

A connector can remain physically connected while corrosion, looseness, damaged plating or poor contact pressure raises its resistance.

Expected: almost no voltage appears across the closed connection under load.

Actual: substantial voltage appears across it because current is encountering resistance.

The machine may become weak while the connection becomes hot.

This is why a continuity beep can miss an important failure. A path can exist well enough for a tiny meter test current and still be terrible at carrying motor current.

Leakage: an unwanted road that is not quite a short

Leakage is unwanted current through a path that should ideally block it.

A contaminated PCB surface can leak. A damaged semiconductor junction can leak. A capacitor dielectric can leak. Wet or carbonized insulation can create a partial conductive path.

The current may be far smaller than a hard short yet still corrupt a sensor, drain a battery, heat a component or prevent a circuit designed to draw only a tiny signal current from reaching the expected voltage.

Leakage matters especially in circuits where the intended signal current is already tiny.

Missing voltage

A voltage expected at a point is absent.

Chapter 9 gave us the Hall throttle with 0 V supply instead of 5 V.

The missing throttle signal was not initially a throttle diagnosis. The common supply was missing.

Missing voltage can result from an upstream open, failed regulator, shorted rail, disabled circuit, protection state or missing reference.

The phrase 'no voltage' is incomplete until you say between which two points.

Wrong voltage

A voltage can be present and still be wrong.

A 5 V rail sitting at 3.1 V may leave logic behaving unpredictably. A battery at normal resting voltage may collapse under load. A throttle signal may sit outside the controller's valid range.

Wrong voltage can be too high, too low, unstable, noisy, or correct only when unloaded.

Always include operating condition with the number when it matters.

Voltage sag: right until asked to work

Suppose a battery measures 52.0 V at rest. Apply throttle and it falls to 38 V, then recovers to 52 V when throttle is released.

You might call the battery voltage 'good' because 52 V looked correct before the test.

The loaded measurement tells a different story. Something in the source path cannot maintain voltage at the demanded current.

Possibilities include weak cells, high internal resistance, BMS behavior, poor battery connection or another source-path resistance.

Sag is a symptom with a location to find, not automatically a verdict on the cells.

Excess current

Current is higher than expected for the operating condition.

A stalled motor can do this normally because back-EMF vanished. A dragging bearing can raise mechanical load and therefore motor current. Shorted motor turns can raise current while reducing torque. A shorted power device can create a completely different current path.

Expected current depends on what the machine is doing.

There is no meaningful universal statement that 'this motor is drawing too many amps' without voltage, speed, load, controller state and expected behavior.

Bad reference or ground

A sensor voltage is measured relative to a reference.

Suppose a Hall throttle signal is physically producing the correct difference relative to its local sensor ground, but the ground connection between throttle and controller has developed resistance.

The controller may now see a different signal voltage than the meter sees when referenced somewhere else.

A bad ground can make several signals look wrong simultaneously because their common measuring ruler moved.

When multiple unrelated signals become strange together, a shared supply or reference deserves attention.

Bad signal

A signal can be missing, stuck, noisy, outside range, mistimed or logically inconsistent.

Chapter 12 gave us one Hall channel stuck at 0 V while the other two toggled.

That narrowed the failure to the C signal path, but did not prove the Hall sensor die itself had failed.

Sensor, wire, connector, pull-up circuitry, supply, reference and controller input can all participate in the same signal path.

The correct question is: at what boundary does the signal stop being what we expect?

Thermal failure: heat can be cause or evidence

Heat can cause failure. Excess winding temperature can damage insulation. Excess semiconductor junction temperature can destroy a MOSFET. Heat can dry an electrolytic capacitor or soften a connector housing.

Heat can also be evidence of another failure. A loose connector gets hot because resistance rose. A MOSFET gets hot because gate drive is inadequate. A motor gets hot because a bearing is dragging and current rose.

Do not diagnose 'overheating' as though heat explains itself.

What produced the heat, and what did the heat change?

Component breakdown

Components work because materials behave predictably inside intended electrical and thermal limits.

Breakdown is the broad condition where an insulating layer, semiconductor junction, dielectric or other structure stops behaving as intended under stress.

A MOSFET gate oxide can be damaged. A diode junction can break down or fail. Capacitor dielectric can break down. Winding insulation can fail between turns.

The resulting electrical failure may then appear as short, leakage, open or wrong behavior.

Breakdown describes the physical loss of intended material behavior; the circuit symptom still belongs to the smaller vocabulary.

Intermittent: the failure refuses to stay failed

An intermittent fault changes state with time or conditions.

Vibration moves a broken conductor into and out of contact. Heat expands a cracked solder joint. A connector pin loses contact only when the harness moves. A semiconductor fails only when hot. A BMS trips only under a particular load.

Intermittents are difficult because a measurement taken while the fault is absent can be perfectly normal.

The operating condition becomes part of the test: temperature, vibration, load, steering position, harness movement, time since startup.

High ESR: a component can fail without becoming open or short

An electrolytic capacitor can age so its equivalent series resistance rises.

At DC with little load it may appear to have capacitance and may not be shorted.

Under ripple current, meaning current that repeatedly rises and falls, the increased ESR creates voltage drop and heat.

A controller rail can therefore become noisy or sag during switching even though the capacitor is neither obviously open nor hard-short.

This is a useful reminder that 'open or short' is not the whole failure vocabulary.

A BMS cutoff is an open with a reason

Suppose a battery powers the scooter normally until hard acceleration, then output disappears. After load is removed or the pack is reset, voltage returns.

The output path became open, but the BMS may have opened it intentionally because it detected overcurrent, undervoltage, temperature or another protection condition.

Replacing the BMS without learning why it opened can remove the messenger instead of fixing the problem.

A shorted MOSFET has a recognizable electrical class

A power MOSFET that fails drain-to-source short can leave a phase leg conducting when the controller expects it off.

Possible symptoms include wheel resistance when pushed, abnormal phase-to-bus readings, blown fuses, large current or a controller that cannot initialize safely.

Those symptoms are not unique to a shorted MOSFET.

The useful category is simpler: an unintended low-resistance path exists in the power stage. Chapter 16 will show how to test boundaries before condemning the device.

One symptom can belong to several classes

A weak motor can come from low source voltage, voltage drop in a connector, excessive mechanical load, bad phase commutation, current limiting, damaged winding or bad throttle command.

A hot controller can come from high current, poor cooling, partially driven MOSFETs, excessive switching loss or a bad connection.

A dead machine can come from an open power path, missing enable, missing low-voltage rail, bad communication or protection state.

This overlap is why parts guessing performs so badly.

Translate the symptom before naming a part

DEAD: ask about missing source, open path, missing enable, missing control supply, missing communication.

WEAK: ask about sag, voltage drop, current limiting, mechanical drag, wrong command, bad commutation.

HOT: ask where power is being dissipated and why.

INTERMITTENT: ask what condition makes the path or signal change.

TRIPS PROTECTION: ask which measured condition crossed which threshold.

WRONG SIGNAL: ask whether supply, reference, sensor, harness and receiving input agree.

EXCESS CURRENT: ask whether load, back-EMF, winding condition, switching state and current measurement are what you think they are.

The list is short enough to hold in your head.

Answered case: the burned connector

A scooter becomes weak on hills. Battery voltage at the pack remains 52.0 V under the test load. Controller input falls to 47.0 V.

You might suspect the battery because the symptom appears only under heavy load. The pack measurement rules out a five-volt collapse at the battery terminals.

Measure across the main battery connector under the same load. It has 4.7 V across it.

A healthy closed connector should have very little voltage across it.

The failure class is unexpected series resistance in the source path.

The melted housing is now supporting evidence, not the basis of the diagnosis.

Answered case: the dead Hall channel

Hall A and B toggle between expected logic levels as the wheel turns. Hall C remains low at the controller.

With power removed, isolate the Hall C signal wire. A suitably supplied test arrangement provides the pull-up required by this particular sensor. The motor-side Hall C output now toggles correctly relative to the verified sensor ground.

Testing the disconnected harness finds the C conductor shorted to ground. Replacing that harness section restores the signal at the controller.

The isolated sensor test and harness test together localize this case to wiring. A signal that is shorted on one continuous wire can be held low along the entire wire; merely moving the probe would not separate the sensor from that short.

The symptom was 'bad Hall signal.' The failed part was wiring.

Answered case: the hot motor that is not an electrical motor fault

A motor runs hot and draws more current than an identical known-good system at the same road load.

With power removed, the wheel is difficult and rough to turn by hand.

You might suspect shorted windings because heat plus current fits that failure class.

The mechanical test gives us another explanation that accounts for both: bearing drag increases mechanical load, which lowers speed/back-EMF and raises current.

The winding has not yet been proven healthy. On a brushless system, disconnect the controller from the motor before attributing drag to bearings: an electrical short can also resist rotation. Persistent roughness after electrical isolation supports a mechanical fault.

Failure signatures are filters

A symptom does not point to one part. It filters the set of possible electrical conditions.

A measurement filters it again.

Another measurement filters it again.

symptom → failure classes → measurement → fewer classes → boundary → cause

That is the transition from knowing electronics to troubleshooting electronics.

Next: stop guessing

Chapter 16 turns this vocabulary into a procedure.

Understand what the machine is supposed to do. Confirm the complaint. Establish the source. Divide the system into functional blocks. Find the boundary where expected behavior disappears. Measure. Eliminate. Move the boundary. Repair. Verify under real load.

The meter stops being a box that produces numbers.

It becomes a machine for killing wrong hypotheses.

CHAPTER 16

16. Stop guessing

Chapter 15 reduced failure to a small vocabulary. Troubleshooting is the method that turns those failure classes into a location.

The goal is not to recognize the broken part from across the room. The goal is to make the machine eliminate possibilities until only a defensible cause remains.

The whole method

1. Understand what the machine is supposed to do.
2. Confirm what it actually does.
3. Establish source integrity.
4. Divide the system into functional blocks.
5. Find a boundary where expected behavior can be tested.
6. Measure there.
7. Use the result to eliminate possibilities.
8. Move the boundary.
9. Identify the failure.
10. Repair it.
11. Verify the repair under realistic load.

Everything else in this chapter is a way of asking one of those questions.

Know what working would look like

A dead scooter is not diagnostically useful until we know the expected sequence.

Battery voltage reaches the controller. An enable condition wakes the low-voltage supply. Five volts appears for throttle and Hall sensors. The throttle produces a valid command. No brake or fault blocks operation. Gate drivers switch the power stage. Phase current produces torque.

That model does not need every transistor. It needs enough functional blocks that a measurement can tell us which side of a boundary is still behaving.


Confirm the complaint

'It won't run' can mean display dead, display alive but no motor response, one jerk, or normal unloaded operation that dies with a rider.

Reproduce the fault safely and note the exact condition: cold, hot, under load, after charging, while steering, during braking, or only from a stop.

The condition is evidence.

Choose the next boundary from the result



A reading eliminates only the faults that the test condition actually exercises.

Figure 15. Start with a boundary that separates large parts of the machine. A good unloaded reading does not substitute for a loaded test.

Establish the source

Before chasing signals, establish whether usable energy reaches the system.

Measure battery voltage at rest and, when appropriate, under the condition that produces the fault.

Then measure at the next power boundary, such as controller input.

Good battery voltage at the pack and zero at the controller makes the motor and throttle uninteresting for the moment. The fault is upstream.

Divide the machine into blocks

source → protection/connectors → controller power → control supply → command/safety → power stage → motor → mechanical load

Subdivide later. First choose boundaries that eliminate large groups of possibilities.

A useful measurement kills possibilities

Full battery voltage at controller input makes a completely open source path unlikely.

If the expected 5 V sensor rail is then 0 V, the failure boundary moves inside the controller: main power arrived, control power did not.

One measurement eliminated a large part of the machine.

Voltage asks one question

What electrical potential difference exists between these two points?

Voltage is measured across two points. Always name both points in your head. 'It has five volts' is incomplete.

Loaded voltage matters

A source or connection can look correct with almost no current and fail when the machine works. Make the fault happen, then measure under that condition when safe.

Voltage-drop testing

Put the meter directly across a connection while current flows.

A healthy closed high-current connection should have very little voltage across it.

Example: battery is 52.0 V, controller sees 47.0 V under load, and one connector measures 4.7 V across itself.

That connector has confessed to nearly the entire missing voltage.

Continuity asks whether a path exists

Continuity mode uses a tiny meter test and usually beeps below some resistance threshold.

It can find broken wires and open fuses with power removed.

It does not prove the path can carry operating current. A corroded connector can beep and still drop volts at motor current.

Resistance asks about opposition

Resistance mode uses the meter's own test signal. External power must be removed and stored capacitor energy handled safely.

In-circuit readings can be misleading because parallel paths may exist.

If the reading makes no sense, ask what else connects the probe points.

Do not measure resistance on an energized circuit

External voltage can corrupt the reading, damage the meter or create a hazard.

Power off is not always enough in power electronics. Verify stored bus voltage has discharged before treating the circuit as unpowered.

Diode mode asks about junction behavior

Diode mode applies a small test condition and reports semiconductor-junction response.

It can expose some hard-open or hard-short diode and power-stage failures.

It is not a universal semiconductor health certificate. Circuit topology and device type matter.

Current measurement is more invasive

An ammeter normally becomes part of the current path.

A voltmeter across a battery is normal. An ammeter configured for current directly across a battery creates a low-resistance path through the meter.

That can blow the meter fuse, damage equipment or create an arc.

Use the correct jack, range and rating. Large vehicle currents are often better measured with a suitable clamp meter or installed sensor.

Signals carry information

A multimeter can answer slow signal questions: does a throttle move from about 1 V to 4 V? Does a Hall line toggle while the wheel turns slowly?

It cannot show the detailed timing of fast PWM very well.

For that, we need voltage plotted against time.

The oscilloscope remembers time

An oscilloscope displays voltage vertically and time horizontally.

It can show PWM, switching transients, ripple, Hall timing and communication activity that a meter may average away.

It is not automatically a better meter. It answers a different question.

Incorrect scope grounding can create destructive shorts. Ordinary bench-scope ground clips are commonly connected together and to protective earth; they are not independent black meter probes. Use an appropriately rated differential probe or isolated-input instrument where the circuit requires it. Never defeat the scope's protective earth. Beginner physical practice here stays with protected, low-energy educational equipment; scooter and charger cases are supplied diagnostic examples.

Divide and conquer

If a signal crosses six harness sections, test near the middle.

Good there? Move downstream. Bad there but good at the sensor? Move upstream within that section.

Choose measurements that eliminate the largest useful part of the map.

Worked case: dead scooter

Complaint: display powers up, but throttle produces no motor response.

battery → controller bus → 5 V rail → throttle → logic → gate drivers → motor

Measurement 1: battery is 52.1 V at rest and 51.7 V during the attempted start. No large sag appears in this test. Because the motor has not begun drawing its normal current, this is not yet a full battery load test.

Measurement 2: controller input is 51.7 V. Main source power arrives.

Measurement 3: throttle supply is 0.0 V; expected 5.0 V.

You might replace the throttle. Do not. Its required supply is missing.

Measurement 4: controller 5 V rail is also 0.0 V.

Measurement 5: low-voltage converter input has 51.7 V.

The converter now looks suspicious. Before replacing it, power is removed, stored voltage is verified discharged, and resistance from the 5 V rail to ground is found extremely low.

That changes the hypothesis: something may be shorting the rail.

Disconnect one external 5 V branch. Rail resistance returns to a normal higher value.

A crushed sensor cable on that branch is shorting 5 V to ground.

Repair it. Reconnect. Power up. The 5 V rail and throttle supply return and the motor runs.

Then verify under load and with steering/harness movement. The repair is not complete until the original condition is gone.

The first plausible part was innocent

Zero throttle output made the throttle look guilty.

Zero throttle supply moved the fault upstream.

Zero 5 V rail moved it into the control-supply system.

Correct converter input made the converter suspicious.

Low output-rail resistance prevented blind replacement.

Disconnecting branches localized the actual short.

Every measurement killed a hypothesis.

Worked case: weak motor, wrong first guess

Complaint: scooter is weak and controller becomes warm.

Battery voltage is normal. Throttle reaches full command.

You might suspect current limiting or MOSFETs.

With power off and the motor disconnected from the controller, the wheel remains unusually hard and rough to rotate. Removing external brake or drivetrain drag localizes the roughness to a motor bearing.

Under power, that mechanical drag raises motor load, lowers speed/back-EMF and raises current.

The warm controller is carrying the consequence. It is not necessarily the cause.

After bearing replacement, verify free rotation, loaded current, speed and controller temperature.

Repair is a hypothesis too

Replacing a burned connector does not prove it was the only problem. Excess current from another fault may have helped burn it.

After repair, reproduce the original load. Recheck voltage drop, current, temperature and signal behavior.

Verification tests whether your explanation was complete.

Do not shotgun parts

Replacing several plausible components at once destroys information.

If the machine works afterward, you may not know which change mattered or whether moving a connector temporarily restored an intermittent.

Parts-cannon troubleshooting can fix machines accidentally. It does not reliably explain them.

The instrument can lie too

Probe contact adds resistance. Automatic range selection takes time. A digital meter may average PWM into a misleading number. An oscilloscope probe can load a delicate circuit. A current clamp can have a zero offset, meaning it may read slightly above or below zero when no current flows.

If a measurement contradicts everything else, inspect the model and the measurement setup.

When measurement disagrees with the model, inspect both.

The diagnostic loop

expected behavior → measure boundary → compare → eliminate → move boundary

Repeat until the remaining explanation is specific enough to repair.

Troubleshooting is not intuition versus procedure. Experience makes you choose better boundaries faster. The procedure keeps experience honest.

Next: something you have never seen before

Chapter 17 removes most of the guidance.

A dead scooter. A bad-running motor. A charger that looks alive but will not charge. A three-phase controller with no throttle response. An intermittent vehicle circuit. A simple PCB power fault.

You will still get the measurements. But first you will have to decide which measurement deserves to exist.

The final case begins with only:

It doesn't work.

By then, that sentence should feel like the beginning of a map rather than the end of one.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [5], [6], [7], [24]. Full titles and links appear in the source list.

CHAPTER 17

17. Here's something you've never seen before

For sixteen chapters, the book has usually told you what the next useful question was.

That ends here.

You are not expected to recognize these machines. Recognition is not the skill we have been building.

For each case, stop before reading the supplied measurement and decide what you would ask the machine next. Then read the answer. If your question was different, that is not failure. Compare how much of the machine each question would eliminate.

Case 1: dead scooter

Complaint: the scooter is completely dead. No display, no motor response.

You know only this: it uses a removable lithium battery feeding a controller through a main fuse and connector.

What deserves the first measurement?

Battery terminals: 54.1 V.

That establishes available resting voltage at the source. It does not prove power reaches the controller.

Next useful boundary: controller battery input.

Controller input: 0.0 V.

At this point the source path deserves attention before the throttle or motor. These voltage readings alone do not yet distinguish an open connection from protection operating in response to a downstream fault.

Measure across the main fuse: 0.0 V. Both sides of the fuse are at 54.1 V relative to battery negative.

These readings show no detectable voltage difference across the fuse in this unloaded state. They do not prove its current-carrying ability. With power removed, an isolated continuity test supports an intact fuse; the final load test will check the whole path.

Measure across the battery-to-controller connector while connected: 54.1 V.

With the negative return verified intact and a defined return through the connected load, nearly the entire source voltage across the positive connector contact points to an open or extremely resistive contact.

Inspection finds a backed-out positive terminal inside the connector.

You may have suspected the fuse because 'dead scooter' and 'fuse' naturally belong together. That was reasonable. Measuring both sides of the fuse eliminated it before a part was replaced.

Repair the terminal, secure its retention, then load-test. Controller input remains near battery voltage and the scooter operates normally.

Diagnosis: open source path at the connector.

Case 2: the brushed motor that eats current

Complaint: a small brushed drive is slow, hot and drawing much more current than it used to.

Battery voltage under load is normal.

What next?

Motor terminal voltage is also normal under load.

That makes a large upstream voltage drop unlikely.

Current is high.

You might now suspect shorted armature turns. That would fit high current, heat and weak torque.

Before opening the motor, disconnect power and turn the shaft.

The shaft is rough and requires much more torque to turn than a known-good motor.

Remove the drive load. The motor shaft itself remains rough.

One bearing is damaged.

The bearing created mechanical load. Mechanical load reduced speed, reduced back-EMF and increased current. The electrical symptoms were real; their origin was mechanical.

Replace the bearing. Free rotation returns. Under the same load, speed and current return near their previous values.

Shorted turns were a reasonable hypothesis. The mechanical test made them unnecessary to explain the evidence.

Case 3: the charger with a convincing green light

Complaint: charger indicator illuminates and its connector measures the expected voltage, but the battery does not charge.

Unloaded charger output: 58.8 V.

You might conclude the charger is good because its voltage is correct.

Connect it to the battery while measuring charger output.

Voltage collapses to 9 V and the charger cycles repeatedly.

The unloaded measurement proved the charger could establish voltage with almost no output current. It did not prove the charger could supply charging power.

Now separate charger failure from battery/BMS refusal.

Test the charger on an appropriate known-good load or battery under controlled conditions. Its output collapses there too.

The failure follows the charger.

Inside the charger, the exact failed component is not yet established. The defensible diagnosis at this boundary is that the charger cannot maintain its regulated output under load.

A green light-emitting diode (LED) was evidence that some internal circuitry was alive. It was never a load test.

Case 4: everything looks good, still no throttle

Complaint: three-phase scooter powers normally but will not drive.

Battery bus at controller: correct.

Internal 5 V rail: 5.02 V.

Motor Hall supply: 5.02 V.

All three Hall signals toggle as the wheel turns.

Throttle supply: 5.02 V.

Throttle signal: 1.0 V at rest, rising smoothly to 4.0 V at full command.

At this point, replacing the throttle would be difficult to defend.

What other condition can intentionally prevent torque?

Brake input is active continuously.

The brake input is active continuously. This example uses a switch that grounds the input to request braking.

With power removed, disconnect the lever switch from the harness. An isolated test shows that the switch opens and closes correctly.

Disconnect the harness at the controller too. The brake-signal conductor still measures shorted to ground with neither end attached. That puts the fault in the harness.

Separate an intermediate connector and test each isolated section. The short remains only in the controller-side section. Inspection finds a pinched wire.

Repair the wire, reconnect everything and verify both brake release and brake operation before a controlled drive test.

The controller was doing exactly what it was told: refuse drive while braking.

A valid throttle signal does not guarantee motor output because command is only one input to the controller's decision.

Case 5: the fault only exists when you turn left

Complaint: an auxiliary vehicle circuit cuts out intermittently, usually during left turns.

Sitting still with the wheels straight, every voltage measurement is normal.

That does not mean the fault vanished from the machine. It means the condition that exposes it is absent.

Reproduce the complaint while safely moving the steering/harness through its range.

The circuit cuts out at one steering position.

Measure source voltage while holding that position: normal.

Measure after the steering-column connector: voltage disappears.

Measure directly across that connector during the fault: nearly full source voltage appears.

Release the steering angle and the drop returns nearly to zero.

The connector path is intermittent under mechanical movement.

Inspection finds a conductor broken inside intact-looking insulation near the terminal crimp.

A continuity test with the harness relaxed would have passed. The useful measurement was made while the machine was failing.

Case 6: the dead five-volt rail on a PCB

Complaint: a small control board receives 24 V but its logic is dead.

You identify the power-input block, protection, a buck-converter section and the 5 V logic rail.

Input to the buck converter: 23.9 V.

Output: 0.0 V.

The converter IC looks guilty.

Chapter 16 taught us not to replace it yet.

Power off. Verify stored voltage is discharged. Measure resistance from the 5 V rail to ground.

Reading: 0.8 Ω .

That is suspiciously low for this board.

Disconnect the major loads on the 5 V rail one at a time. After unplugging one daughterboard, a smaller circuit board connected to the main board, resistance rises to several kilohms.

The short is on the daughterboard, not necessarily in the buck converter.

On that board, a small supply-stabilizing capacitor measures effectively shorted across the rail. Remove it and the short disappears. Confirm the capacitor is shorted out of circuit.

Replace it with an appropriate part. Reassemble. The buck converter produces 5.01 V and the logic starts.

The regulator was correctly refusing or failing to raise a rail that was being dragged almost directly to ground.

The clues are getting smaller

Notice what changed across those cases.

The first told you the source path.

The second forced electrical symptoms back into mechanics.

The third separated voltage present from power available.

The fourth made a valid command irrelevant because an interlock overruled it.

The fifth required reproducing an intermittent condition.

The sixth made the apparently dead component innocent because something downstream was loading it.

Different machines. Same method.

Final case

It doesn't work.

That is all the complaint says.

The machine is a small battery-powered pump controller. A button is supposed to start a DC pump. An indicator should illuminate while the pump runs.

You have never seen this model.

Do not begin with the pump.

First define expected behavior.

battery → protection → controller supply → button/enable → power switch → pump

The indicator may belong to the control side or load side. We do not know yet.

Confirm the complaint: pressing the button produces no pump motion and no indicator.

Establish source: battery measures 12.6 V at rest.

Press the button while measuring battery voltage: 12.5 V.

No collapse appears in this attempted-start measurement. If the load never switches on, this is not a full battery capacity or load test.

Measure controller input: 12.5 V.

Source and main path reach the controller.

Now find the control supply. A small regulator on the board is labeled by circuit context as feeding the logic section.

Its output is 5.0 V.

The controller has power and low-voltage logic supply.

Measure the button input. Released: 5.0 V. Pressed: 0.0 V.

The button signal changes cleanly.

You might now suspect the power switch or pump.

Before going there, ask whether the controller is commanding the switch.

The gate-to-source control voltage of the power MOSFET remains 0 V whether the button is pressed or released.

The power stage is not being commanded.

Why?

Look for other controller inputs. There is a two-wire level sensor intended to prevent the pump from running dry.

Trace the connector with power removed. One wire is ground. The other reaches a logic input with a pull-up. A pull-up is a resistor that gives an otherwise undriven input a defined HIGH state.

Sensor connected: the input remains 5.0 V. That is a clue, not yet proof that the sensor is faulty.

The sensor's own part marking identifies it as a float-operated dry-contact switch, meaning a plain switch contact that supplies no voltage of its own. Its specification calls for closed contacts at the full-tank position. Isolated from the board, its contacts remain open throughout the specified travel.

Connect a compatible replacement sensor with power removed, place it in its full-tank state, then power the controlled test setup. The input goes to 0 V. Press the button. The MOSFET gate command appears, the pump runs and the indicator lights.

You have not proven every detail of the controller. You have proven enough.

The controller was intentionally withholding pump drive because the level interlock never reported a safe condition.

Replace the failed level sensor with the correct part. Reassemble and verify operation at full, low and transition levels.

Then verify that the dry-run protection still works. A repair that makes the pump run but defeats the protection is not a repair.

What if your first test was different?

Perhaps you would have measured pump resistance first.

That is not stupid. If the pump were open, the measurement could reveal it.

But even a healthy pump would not explain why the indicator was also dead or why no gate command existed.

A better early test attacks a boundary shared by more possible failures.

Troubleshooting proficiency is partly learning which question buys the most information.

You are allowed to not know

There will always be a component you have never seen, a protocol you do not recognize, a controller whose firmware is undocumented, or a machine whose wiring makes no immediate sense.

The skill is not knowing every answer before touching the machine.

The skill is preserving what you know, labeling what you do not know, and asking measurements that turn unknowns into smaller unknowns.

The final change

At the beginning of this book, a battery and motor were enough to create mysteries.

Voltage, current, resistance and power had no meaning yet. A motor drawing more current under load looked like a fact to memorize.

Now an unfamiliar machine can be reduced to sources, paths, loads, control, switching, fields, signals and feedback.

A failure can be reduced to opens, shorts, resistance, leakage, wrong voltage, wrong current, bad references, bad signals, heat and intermittence.

And a diagnosis can be reduced to a loop:

expect → measure → compare → eliminate → move

That is enough to begin on almost anything.

The destination

Somebody who finishes this book is not an electrical engineer merely because they finished it.

They are not expected to design a grid inverter from scratch, repair energized mains equipment, or know every semiconductor topology.

The useful threshold is more modest and more powerful.

Put an unfamiliar low-voltage electrical machine on the bench and the reader can say:

I don't know this machine yet, but I know how to ask it questions.

That is electronics literacy.

You do not need to remember the book.

You need to understand enough of it to know what to look up.

Sources and example notes

Numerical cases are constructed teaching examples, not product test records. Use the exact component and instrument documentation for real equipment.

Written references: [5], [6], [7], [13], [21]. Full titles and links appear in the source list.

The Bench Book

The chapters were a path. This is the permanent bench reference. Look things up. Nothing here is a test of memory.

Electrical language

Quantity	Symbol / unit	Meaning
Voltage	V / volt (V)	Electrical potential difference between two points.
Current	I / ampere (A)	Rate of charge flow.
Resistance	R / ohm (Ω)	Opposition to current.
Power	P / watt (W)	Rate of energy transfer.
Energy	joule (J), watt-hour (Wh)	Amount of energy.
Capacitance	C / farad (F)	Electric-field energy/charge-storage behavior.
Inductance	L / henry (H)	Magnetic-field behavior opposing changes in current.
Frequency	f / hertz (Hz)	Cycles per second.
Time	t / second (s)	Duration.
Mark	Meaning	
=	equals	
+	add	
-	subtract	
x	multiply	
÷ /	divide	
²	square	
≈	approximately	
Δ	change in	
...	continue	
Prefix	Symbol	Multiplier
pico	p	one trillionth (0.000000000001)
nano	n	one billionth (0.000000001)
micro	μ	one millionth (0.000001)
milli	m	one thousandth (0.001)
kilo	k	one thousand (1,000)
mega	M	one million (1,000,000)
giga	G	one billion (1,000,000,000)

Prefixes change the size of a unit, not the kind of thing being measured. For example, 50 mA means 50 thousandths of an amp, or 0.050 A; 1 kΩ means 1,000 Ω. Case matters: m means milli; M means mega.

Equation library

Use $P = V \times I$ directly for steady DC. For changing waveforms, average the instantaneous voltage-current product; multiplying two separate averages can be wrong. In I^2R heating, use root-mean-square (RMS) current, the steady-current value that would produce the same resistive heating, when current varies. The Motor M equation is a settled, simplified brushed-motor model: it omits inductive transients and brush voltage drop.

Ohm's law

$$I = V \div R$$

Current equals voltage divided by resistance. Example: $6\text{ V} \div 120\ \Omega = 0.050\text{ A}$.

Voltage drop

$$V = I \times R$$

Voltage across resistance equals current times resistance.

Resistance

$$R = V \div I$$

Effective resistance equals measured voltage divided by current.

Electrical power

$$P = V \times I$$

Power equals voltage times current.

Resistive heating

$$P = I^2 \times R$$

Double current through unchanged resistance and heating becomes four times larger.

Motor with back-EMF

$$I = (V - E) \div R$$

E is opposing generated back-EMF.

Series resistance

$$R_{\text{total}} = R_1 + R_2 + \dots$$

Series resistances add.

Parallel resistance

$$1/R_{\text{total}} = 1/R_1 + 1/R_2 + \dots$$

Parallel routes reduce equivalent resistance.

Branch current

$$I_{\text{total}} = I_1 + I_2 + \dots$$

Branch currents sum at a junction.

Battery energy

$$Wh = V \times Ah$$

Approximate nominal energy from voltage and amp-hours.

PWM duty

$$\text{Duty} = \text{ON time} \div \text{cycle time}$$

Fraction of each cycle spent ON.

Shunt current

$$I = V_{\text{shunt}} \div R_{\text{shunt}}$$

Known small resistance turns voltage drop into current measurement.

Component families

Family	Members	What makes them work	Common failures
Resistive	wire, resistor, heater, pot	Material + geometry set resistance	heat, drift, open, worn contact
Capacitive	ceramic, film, electrolytic	Conductors separated by dielectric	leakage, ESR, short/open, overvoltage
Magnetic	inductor, transformer, relay/motor winding	Current and changing magnetic fields	saturation, copper/core heat, open/shorted turns
Semiconductor	diode, BJT, MOSFET, Hall, IC	Doped materials, junctions, electric fields	short, open, leakage, thermal/gate damage
Electromechanical	switch, relay, motor, solenoid	Electricity interacts with motion	wear, arcing, bearings, contacts, insulation

Component identification

Clue	Useful for	Does not prove
Part number	Finding exact datasheet	Health
Resistor bands/print	Value/tolerance	Power rating always
Capacitor print	Capacitance/voltage/polarity clues	ESR/health
Diode band	Cathode on many parts	Exact type
PCB label	Likely rail/signal	Universal meaning
Connector shape	Physical mating	Pinout/protocol/compatibility
Wire color	Documented-system clue	Universal function

Never assume transistor/MOSFET pin order. Use the exact datasheet. IC pin 1 may use a dot, notch or bevel.

Typical measurements

Thing	Useful healthy pattern
Closed power connection	Tiny loaded voltage drop
Bad connection	Unexpected loaded voltage drop, often heat
Wire/fuse	Continuity when unpowered; continuity alone does not prove current capacity
Resistor	Near marked resistance if isolated enough
Diode	Forward response one direction, blocking opposite in simple isolated test
Hall sensor	Supply/reference correct; output changes with magnetic position


Thing	Useful healthy pattern
Hall throttle	Supply correct; signal changes smoothly through expected range
Battery	Resting voltage plus loaded sag behavior
PWM node	Oscilloscope shows pulses; multimeter may show an average-like number

Failure signatures

Class	Meaning	Clues
Open	Required path disappeared	No current in the interrupted path; voltage may remain across the open
Short	Unwanted low-resistance path	High current, protection trip, stuck-on behavior
Unexpected resistance	Path got electrically narrow	Voltage drop, heat, weak load
Leakage	Unwanted partial path	Drain, corrupted signal, heating
Drift	Value moved	Wrong bias/timing/current/voltage
High ESR	Capacitor loss increased	Ripple, sag, heating
Overheated	Temperature exceeded intended condition	Derating, damage, changed value
Intermittent	Fault depends on condition	Heat, vibration, harness motion, load
Bad reference	Measuring ruler moved	Several signals wrong together
Bad signal	Information wrong/missing	Control refuses or miscommands
Breakdown	Material/junction/insulation stops behaving	Leak, arc, short or open

Schematic symbols


Schematic symbols: connections and passive parts



Common variants exist. Symbol terminals do not specify the physical package pin order.


Reference plate A. Connection and passive-component symbols. The resistor is shown in the rectangular convention.

Schematic symbols: switching and devices



Diode

K = cathode bar




LED


arrows point out




Switch




Motor



Fuse




Transformer




Relay

dashed mechanical link



NPN BJT

emitter arrow out



N-MOSFET

body diode shown

Common variants exist. Symbol terminals do not specify the physical package pin order.

Reference plate B. Common device symbols. LED means light-emitting diode. BJT means bipolar junction transistor; NPN names the N-P-N order of semiconductor regions in that BJT. N-MOSFET means N-channel MOSFET. The N-MOSFET drawing is simplified; its body diode conducts from source toward drain. Refer to the exact device schematic and pinout.

Item	Common representation / meaning
Wire	line: connection
Junction	meeting lines/dot: connected node; check drawing convention
Ground/common	reference node; not automatically earth
Battery	DC source with polarity
Resistor	zigzag (common US convention) or rectangle (common international convention)
Potentiometer	resistor plus wiper
Capacitor	two plates; polarized variants marked

Item	Common representation / meaning
Inductor	coil
Transformer	coupled coils
Diode	junction symbol with cathode bar
LED	diode plus light arrows
Switch	movable contact
Relay	coil plus linked contacts
BJT	three-terminal transistor
MOSFET	gate/channel transistor symbol
Motor	M/circle or motor symbol
Fuse	overcurrent-open device

Mechanisms

Voltage divider

Two series resistances create an intermediate voltage. In $V_{out} = V_{in} \times R_2 \div (R_1 + R_2)$, V_{in} is the input voltage, V_{out} is the output voltage, and R_1 and R_2 are the two resistance values. Best for signal/reference work, not substantial loads.

Pull-up / pull-down

A resistor gives an otherwise undriven logic/signal node a defined default state.

RC filter

A resistor and capacitor shape how quickly voltage changes; useful for filtering/noise/timing.

Flyback clamp

A diode gives DC coil current a safe recirculating path when its switch opens.

RC snubber

Resistor + capacitor tame fast switching transients and dissipate stored energy.

Buck

Switching converter that makes lower DC voltage efficiently.

Boost

Switching converter that makes higher DC voltage; energy is conserved, losses remain.

Precharge

Charge bus capacitors through a limited path before closing the main connection.

Soft start

Ramp current/torque/duty instead of applying maximum demand instantly.

PWM

Switch full voltage ON/OFF rapidly; duty cycle controls average drive, not exact RPM.

Half-bridge

High-side + low-side switch control one node toward either DC bus rail. Avoid shoot-through.

H-bridge

Four switches reverse current through a two-terminal load.

Three-phase bridge

Three half-bridges control motor phases to create a rotating magnetic field.

Hall interface

Supply/reference plus Hall signal lets magnetic position/current/speed become information.

Regen

Motor generation is switched back toward DC bus/battery if the system can accept it.

Field weakening

Changes magnetic-current relationship for extra speed at cost of current/loss/torque margin.

FOC

Uses rotor position/current information to place the effective stator field precisely.

Reference tables

American Wire Gauge (AWG) dimensions

These are conductor dimensions, not universal ampacity ratings. Ampacity means the allowable current under specified conditions.

AWG	Diameter	Area
24	0.511 mm	0.205 mm ²
22	0.644 mm	0.326 mm ²
20	0.812 mm	0.518 mm ²
18	1.024 mm	0.823 mm ²
16	1.291 mm	1.31 mm ²
14	1.628 mm	2.08 mm ²
12	2.053 mm	3.31 mm ²
10	2.588 mm	5.26 mm ²
8	3.264 mm	8.37 mm ²

Allowable current depends on insulation, ambient temperature, bundling, airflow, duty, termination and applicable standards.

Resistor color digits

Color	Digit	Multiplier
Black	0	×1
Brown	1	×10
Red	2	×100
Orange	3	×1k
Yellow	4	×10k
Green	5	×100k

Color	Digit	Multiplier
Blue	6	×1M
Violet	7	×10M
Gray	8	×100M
White	9	×1G
Gold	-	×0.1
Silver	-	×0.01

Common tolerance bands: brown $\pm 1\%$, red $\pm 2\%$, gold $\pm 5\%$, silver $\pm 10\%$. Verify ambiguous band schemes.

Four-band example: brown, black, red, gold means $10 \times 100 \Omega = 1,000 \Omega$, $\pm 5\%$. A common five-band code uses three value digits before the multiplier and tolerance. Read the part's documented scheme.

Capacitor markings

Marking	Typical meaning
100 μF 63 V	100 microfarads; 63 V rating
104	100,000 pF = 100 nF
473	47,000 pF = 47 nF
Stripe / +	Polarity clue on applicable types
105 °C	Temperature/series rating clue

Voltage rating is a limit. It is not the voltage a capacitor supplies.

Connector basics

Question	Why it matters
Contact current rating?	Pins can overheat before wire does.
Voltage/insulation rating?	Spacing/materials must tolerate real voltage and transients.
Keying/polarity?	Helps prevent some incorrect mating, not electrical incompatibility.
Correct crimp/tool?	Mechanical attachment and low electrical resistance depend on termination quality.
Pinout documented?	Best basis for assigning functions.
Protocol documented?	The same connector can carry simple voltage signals, serial data such as UART or CAN, or proprietary signals.

Connector fit is never proof of electrical compatibility.

Bench sequence

What should happen? → Where is the boundary? → What should I measure there? → What did I actually get?

Then classify the disagreement: open, short, resistance, leakage, wrong voltage/current/reference/signal, heat, intermittent, or breakdown.

measure → eliminate → move boundary → repair → verify under load

That is the whole book compressed into a bench card.

Sources and example notes

Lookup entries summarize the chapters. Component pinouts, limits and test conditions remain device-specific. American Wire Gauge (AWG) dimensions are for bare solid conductors or equivalent conductor area, not finished cable diameter.

Written references: [4], [7], [8], [9], [12], [13], [16], [18], [19], [21], [25], [26]. Full titles and links appear in the source list.

Written sources

The references below support the definitions, mechanisms and representative component behavior used in the chapters. Numerical teaching cases and original diagrams were composed for this book. A cited product is an example, not an endorsement or a universal specification.

Source numbers are shared across the book. Full titles and addresses are retained so printed copies remain usable. Online addresses can change; the publisher and document title provide a second way to find each source.

[1] **OpenStax**. University Physics Volume 2, §9.1: Electrical Current. Charge, current and the circuit model.

<https://openstax.org/books/university-physics-volume-2/pages/9-1-electrical-current>

[2] **Pololu**. Jrk G2 Motor Controller User's Guide, §4.1: Choosing the motor, power supply, and Jrk. Brushed-motor operating conditions; free-run and stall current; source capability and voltage margin.

<https://www.pololu.com/docs/0J73/4.1>

[3] **maxon Support**. Motor constants. Torque/current and speed/back-EMF relationships.

<https://support.maxongroup.com/hc/en-us/articles/360005873794-Motor-constants>

[4] **NIST**. Guide to the SI, Chapter 4: The Two Classes of SI Units and the SI Prefixes. Units and prefixes used in this book.

<https://www.nist.gov/pml/special-publication-811/nist-guide-si-chapter-4-two-classes-si-units-and-si-prefixes>

[5] **Fluke**. How to Measure DC Voltage with a Digital Multimeter. Voltage measurement, probe polarity and instrument setup.

<https://www.fluke.com/en-us/learn/blog/digital-multimeters/how-to-measure-dc-voltage-with-a-digital-multimeter>

[6] **Fluke**. How to Use a Multimeter: A Comprehensive Guide. Measurement modes, connections and limitations.

<https://www.fluke.com/en/learn/blog/maintenance-monitoring/how-to-use-a-multimeter-guide>

[7] **Fluke**. How to Measure Resistance with a Digital Multimeter. Unpowered resistance tests and parallel-path limitations.

<https://www.fluke.com/en-us/learn/blog/digital-multimeters/how-to-measure-resistance>

[8] **OpenStax**. University Physics Volume 2, §9.5: Electrical Energy and Power. Electrical power and resistive heating.

<https://openstax.org/books/university-physics-volume-2/pages/9-5-electrical-energy-and-power>

[9] **OpenStax**. University Physics Volume 2, §10.2: Resistors in Series and Parallel. Series and parallel circuit relationships.

<https://openstax.org/books/university-physics-volume-2/pages/10-2-resistors-in-series-and-parallel>

[10] **OpenStax**. University Physics Volume 2, §14.2: Self-Inductance and Inductors. Inductance and induced voltage.

<https://openstax.org/books/university-physics-volume-2/pages/14-2-self-inductance-and-inductors>

[11] **Texas Instruments**. Minimizing Ringing at the Switch Node of a Boost Converter, SLVA255. Switching transients and RC snubber behavior; application-specific example.

<https://www.ti.com/lit/an/slva255/slva255.pdf>

[12] **Bourns**. 3386 3/8-Inch Square Trim Pot Trimming Potentiometer. Representative potentiometer construction, terminal arrangement and ratings.

<https://www.bourns.com/docs/product-datasheets/3386.pdf>

[13] **Texas Instruments**. Avoid Common Mistakes When Selecting And Designing With Power MOSFETs, SLPA021. Gate-drive voltage, threshold versus low on-resistance, ratings and loss mechanisms.

<https://www.ti.com/lit/an/slpa021/slpa021.pdf>

[14] **Texas Instruments**. LM2596 SIMPLE SWITCHER Power Converter, datasheet. Representative buck-converter circuit and operating limits.

<https://www.ti.com/lit/ds/symlink/lm2596.pdf>

[15] **Texas Instruments**. DRV5055 Ratiometric Linear Hall Effect Sensor, datasheet. Hall sensing, supply dependence and analog output; not a universal throttle specification.

<https://www.ti.com/lit/ds/symlink/drv5055.pdf>

[16] **OpenStax**. University Physics Volume 2, §8.1: Capacitors and Capacitance. Capacitor structure, charge separation and capacitance.

<https://openstax.org/books/university-physics-volume-2/pages/8-1-capacitors-and-capacitance>

[17] **OpenStax**. University Physics Volume 3, §9.7: Semiconductor Devices. Doping, junctions and semiconductor devices.

<https://openstax.org/books/university-physics-volume-3/pages/9-7-semiconductor-devices>

[18] **Nichicon**. General Descriptions of Aluminum Electrolytic Capacitors. Electrolytic construction, losses, polarity and operating precautions.

<https://www.nichicon.co.jp/english/products/pdf/aluminum.pdf>

[19] **Texas Instruments**. Trapezoidal Control of BLDC Motors Using Hall Effect Sensors, SPRABQ6. Three-phase commutation, Hall sensing, power bridges and current waveforms.

<https://www.ti.com/lit/an/sprabq6/sprabq6.pdf>

[20] **Texas Instruments**. Sensorless Field Oriented Control of 3-Phase Permanent Magnet Synchronous Motors, SPRABQ3. Rotor estimation, current feedback and field-oriented control.

<https://www.ti.com/lit/an/sprabq3/sprabq3.pdf>

[21] **Texas Instruments**. DRV8301 Three-Phase Gate Driver With Dual Current Shunt Amplifiers and Buck Regulator, datasheet. Representative controller blocks, dead time, gate drivers and current sensing.

<https://www.ti.com/lit/ds/symlink/drv8301.pdf>

[22] **Texas Instruments**. Best Practices for Board Layout of Motor Drivers, SLVA959B. Grounding, thermal paths, bus capacitors, bridge layout and current-sense placement.

<https://www.ti.com/lit/an/slva959b/slva959b.pdf>

[23] **Molicel**. INR-21700-P42A Product Data Sheet, version 4. One documented 4.2 V charge-voltage cell example; chemistry-specific limits.

<https://www.molicel.com/wp-content/uploads/INR21700P42A-V4-80092.pdf>

[24] **Tektronix**. Fundamentals of Floating Measurements and Isolated Input Oscilloscopes. Floating measurements and instrument isolation.

<https://www.tek.com/en/documents/application-note/fundamentals-floating-measurements-and-isolated-input-oscilloscopes>

[25] **National Bureau of Standards**. Copper Wire Tables, Handbook 100, 1966. American Wire Gauge dimensions; Tables 1 and 10. Dimensions are not ampacity ratings.

<https://nvlpubs.nist.gov/nistpubs/Legacy/hb/nbshandbook100.pdf>

[26] **Vishay**. Color Code, document 20143. Resistor digits, multipliers and tolerance bands.

<https://www.vishay.com/docs/20143/colorcod.pdf>

[27] **Texas Instruments**. How to Drive Resistive, Inductive, Capacitive, and Lighting Loads, SLVAE30E. Inductive turn-off, demagnetization paths, capacitive inrush and thermal behavior.

<https://www.ti.com/lit/an/slvae30e/slvae30e.pdf>