Electricity

2 nd Education

Dr. Mohamed Allosh

South Valley University

2022-2023

What is Electricity?

Electricity is all around us--powering technology like our cell phones, computers, lights, soldering irons, and air conditioners. It's tough to escape it in our modern world. Even when you try to escape electricity, it's still at work throughout nature, from the lightning in a thunderstorm to the synapses inside our body. But what exactly *is* electricity? This is a very complicated question, and as you dig deeper and ask more questions, there really is not a definitive answer, only abstract representations of how electricity interacts with our surroundings.

Electricity is a natural phenomenon that occurs throughout nature and takes many different forms. In this tutorial we'll focus on current electricity: the stuff that powers our electronic gadgets. Our goal is to understand how electricity flows from a power source through wires, lighting up LEDs, spinning motors, and powering our communication devices.

Electricity is briefly defined as the flow of electric charge, but there's so much behind that simple statement. Where do the charges come from? How do we move them? Where do they move to? How does an electric charge cause mechanical motion or make things light up? So many questions! To begin to explain what electricity is we need to zoom way in, beyond the matter and molecules, to the atoms that make up everything we interact with in life.

Going Atomic

To understand the fundamentals of electricity, we need to begin by focusing in on atoms, one of the basic building blocks of life and matter. Atoms exist in over a hundred different forms as chemical elements like hydrogen, carbon, oxygen, and copper. Atoms of many types can combine to make molecules, which build the matter we can physically see and touch.

Atoms are *tiny*, stretching at a max to about 300 picometers long (that's 3×10^{-10} or 0.0000000003 meters). A copper penny (if it actually were made of 100% copper) would have 3.2×10^{22} atoms (32,000,000,000,000,000,000,000 atoms) of copper inside it.

Even the atom isn't small enough to explain the workings of electricity. We need to dive down one more level and look in on the building blocks of atoms: protons, neutrons, and electrons.

Building Blocks of Atoms

An atom is built with a combination of three distinct particles: electrons, protons, and neutrons. Each atom has a center nucleus, where the protons and neutrons are densely packed together. Surrounding the nucleus are a group of orbiting electrons.

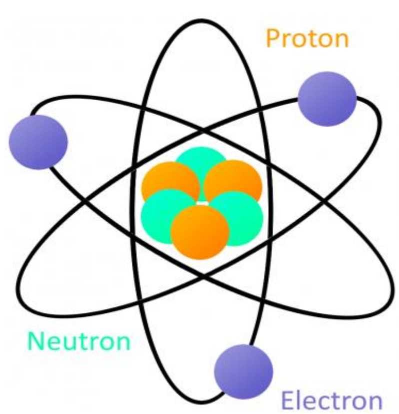

A very simple atom model. It's not to scale but helpful for understanding how an atom is built. A core nucleus of protons and neutrons is surrounded by orbiting electrons.

Every atom must have at least one proton in it. The number of protons in an atom is important, because it defines what chemical element the atom represents. For example, an atom with just one proton is hydrogen, an atom with 29 protons is copper, and an atom with 94 protons is plutonium. This count of protons is called the atom's atomic number.

The proton's nucleus-partner, neutrons, serves an important purpose; they keep the protons in the nucleus and determine the isotope of an atom. They're not critical to our understanding of electricity, so let's not worry about them for this tutorial.

Electrons are critical to the workings of electricity (notice a common theme in their names?) In its most stable, balanced state, an atom will have the same number of electrons as protons. As in the Bohr atom model below, a nucleus with 29 protons (making it a copper atom) is surrounded by an equal number of electrons.

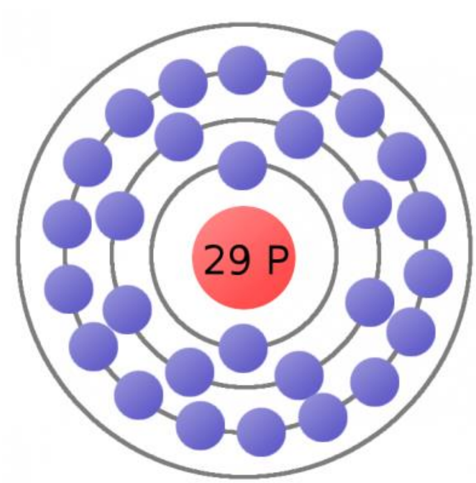

As our understanding of atoms has evolved, so too has our method for modeling them. The Bohr model is a very useful atom model as we explore electricity.

The atom's electrons aren't all forever bound to the atom. The electrons on the outer orbit of the atom are called valence electrons. With enough outside force, a valence electron can escape orbit of the atom and become free. Free electrons allow us to move charge, which is what electricity is all about. Speaking of charge...

Flowing Charges

As we mentioned at the beginning of this tutorial, electricity is defined as the flow of electric charge. Charge is a property of matter--just like mass, volume, or density. It is measurable. Just as you can quantify how much mass something has, you can measure how much charge it has. The key concept with charge is that it can come in two types: positive (+) or negative (-).

In order to move charge we need charge carriers, and that's where our knowledge of atomic particles--specifically electrons and protons--comes in handy. Electrons always carry a negative charge, while protons are always positively charged. Neutrons (true to their name) are neutral, they have no charge. Both electrons and protons carry the same amount of charge, just a different type.

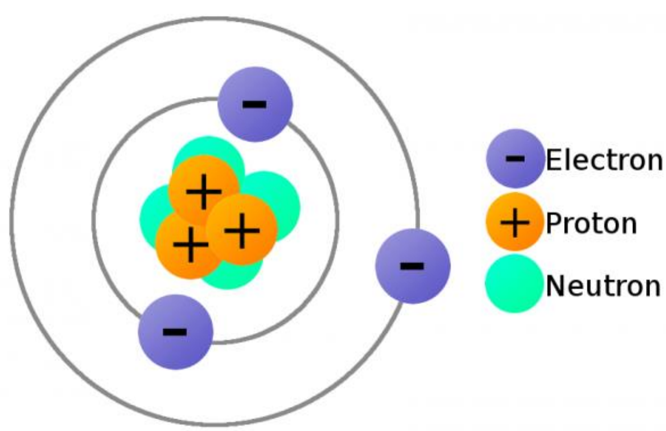

A lithium atom (3 protons) model with the charges labeled.

The charge of electrons and protons is important, because it provides us the means to exert a force on them. Electrostatic force!

Electrostatic Force

Electrostatic force (also called Coulomb's law) is a force that operates between charges. It states that charges of the same type repel each other, while charges of opposite types are attracted together. Opposites attract, and likes repel.

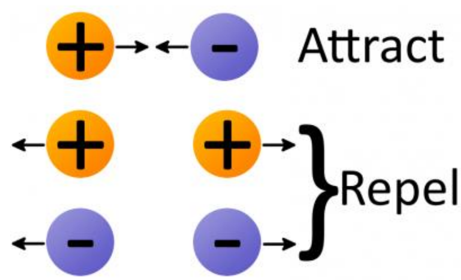

The amount of force acting on two charges depends on how far they are from each other. The closer two charges get, the greater the force (either pushing together, or pulling away) becomes.

Thanks to electrostatic force, electrons will push away other electrons and be attracted to protons. This force is part of the "glue" that holds atoms together, but it's also the tool we need to make electrons (and charges) flow!

Making Charges Flow

We now have all the tools to make charges flow. Electrons in atoms can act as our charge carrier, because every electron carries a negative charge. If we can free an electron from an atom and force it to move, we can create electricity.

Consider the atomic model of a copper atom, one of the preferred elemental sources for charge flow. In its balanced state, copper has 29 protons in its nucleus and an equal number of electrons orbiting around it. Electrons orbit at varying distances from the nucleus of the atom. Electrons closer to the nucleus feel a much stronger attraction to the center than those in distant orbits. The outermost electrons of an atom are called the valence electrons; these require the least amount of force to be freed from an atom.

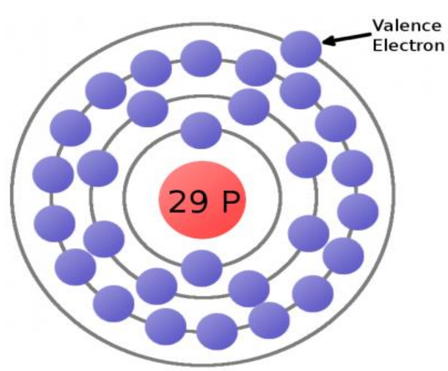

This is a copper atom diagram: 29 protons in the nucleus, surrounded by bands of circling electrons. Electrons closer to the nucleus are hard to remove while the valence (outer ring) electron requires relatively little energy to be ejected from the atom.

Using enough electrostatic force on the valence electron--either pushing it with another negative charge or attracting it with a positive charge--we can eject the electron from orbit around the atom creating a free electron.

Now consider a copper wire: matter filled with countless copper atoms. As our free electron is floating in a space between atoms, it's pulled and prodded by surrounding charges in that space. In this chaos the free electron eventually finds a new atom to latch on to; in doing so, the negative charge of that electron ejects another valence electron from the atom. Now a new electron is drifting through free space looking to do the same thing. This chain effect can continue on and on to create a flow of electrons called electric current.

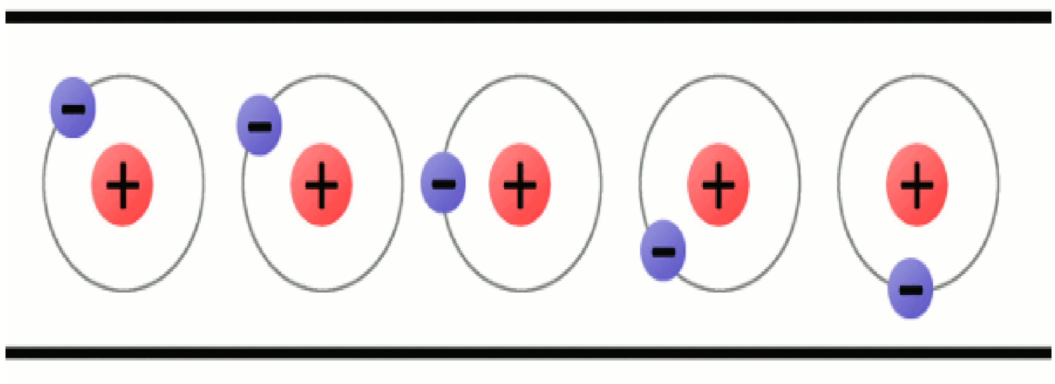

A very simplified model of charges flowing through atoms to make current

Conductivity

Some elemental types of atoms are better than others at releasing their electrons. To get the best possible electron flow we want to use atoms which don't hold very tightly to their valence electrons. An element's conductivity measures how tightly bound an electron is to an atom.

Elements with high conductivity, which have very mobile electrons, are called conductors. These are the types of materials we want to use to make wires and other components which aid in electron flow. Metals like copper, silver, and gold are usually our top choices for good conductors.

Elements with low conductivity are called insulators. Insulators serve a very important purpose: they prevent the flow of electrons. Popular insulators include glass, rubber, plastic, and air.

Static or Current Electricity

Before we get much further, let's discuss the two forms electricity can take: static or current. In working with electronics, current electricity will be much more common, but static electricity is important to understand as well.

Static Electricity

Static electricity exists when there is a build-up of opposite charges on objects separated by an insulator. Static (as in "at rest") electricity exists until the two groups of opposite charges can find a path between each other to balance the system out.

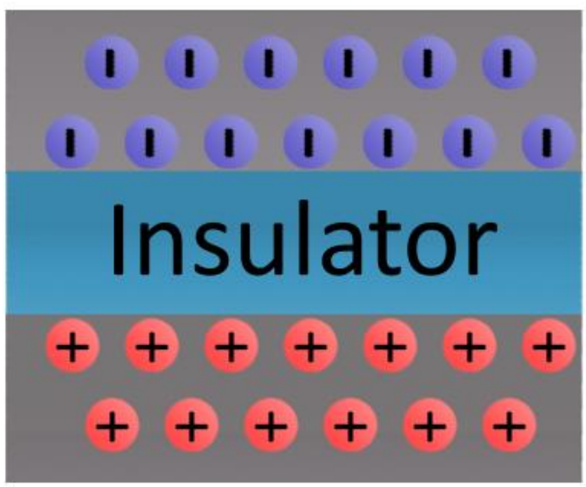

When the charges do find a means of equalizing, a static discharge occurs. The attraction of the charges becomes so great that they can flow through even the best of insulators (air, glass, plastic, rubber, etc.). Static discharges can be harmful depending on what medium the charges travel through and to what surfaces the charges are transferring. Charges equalizing through an air gap can result in a visible shock as the traveling electrons collide with electrons in the air, which become excited and release energy in the form of light.

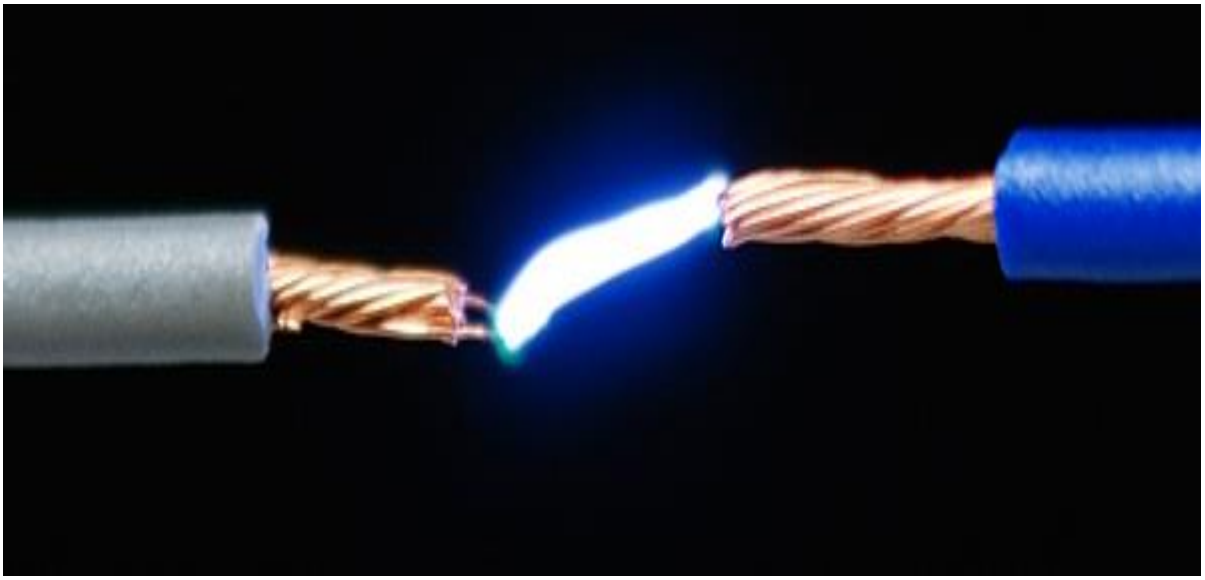

Spark gap igniters are used to create a controlled static discharge. Opposite charges build up on each of the conductors until their attraction is so great charges can flow through the air

One of the most dramatic examples of static discharge is lightning. When a cloud system gathers enough charge relative to either another group of clouds or the earth's ground, the charges will try to equalize. As the cloud discharges, massive quantities of positive (or sometimes negative) charges run through the air from ground to cloud causing the visible effect we're all familiar with.

Static electricity also familiarly exists when we rub balloons on our head to make our hair stand up, or when we shuffle on the floor with fuzzy slippers and shock the family cat (accidentally, of course). In each case, friction from rubbing different types of materials transfers electrons. The object losing electrons becomes positively charged, while the object gaining electrons becomes negatively charged. The two objects become attracted to each other until they can find a way to equalize.

Working with electronics, we generally don't have to deal with static electricity. When we do, we're usually trying to protect our sensitive electronic components from being subjected to a static discharge. Preventative measures against static electricity include wearing ESD (electrostatic discharge) wrist straps, or adding special components in circuits to protect against very high spikes of charge.

Current Electricity

Current electricity is the form of electricity which makes all of our electronic gizmos possible. This form of electricity exists when charges are able to constantly flow. As opposed to static electricity where charges gather and remain at rest, current electricity is dynamic; charges are always on the move. We'll be focusing on this form of electricity throughout the rest of the tutorial.

Circuits

In order to flow, current electricity requires a circuit: a closed, never-ending loop of conductive material. A circuit could be as simple as a conductive wire connected end-to-end, but useful circuits usually contain a mix of wire and other components which control the flow of electricity. The only rule when it comes to making circuits is they can't have any insulating gaps in them.

If you have a wire full of copper atoms and want to induce a flow of electrons through it, all free electrons need somewhere to flow in the same general direction. Copper is a great conductor, perfect for making charges flow. If a circuit of copper wire is broken, the charges can't flow through the air, which will also prevent any of the charges toward the middle from going anywhere.

On the other hand, if the wire were connected end-to-end, the electrons all have a neighboring atom and can all flow in the same general direction.

We now understand how electrons can flow, but how do we get them flowing in the first place? Then, once the electrons are flowing, how do they produce the energy required to illuminate light bulbs or spin motors? For that, we need to understand electric fields.

Electric Fields

We have a handle on how electrons flow through matter to create electricity. That's all there is to electricity. Well, almost all. Now we need a source to induce the flow of electrons. Most often that source of electron flow will come from an electric field.

What's a Field?

A field is a tool we use to model physical interactions which don't involve any observable contact. Fields can't be seen as they don't have a physical appearance, but the effect they have is very real.

We're all subconsciously familiar with one field in particular: Earth's gravitational field, the effect of a massive body attracting other bodies. Earth's gravitational field can be modeled with a set of vectors all pointing into the center of the planet; regardless of where you are on the surface, you'll feel the force pushing you towards it.

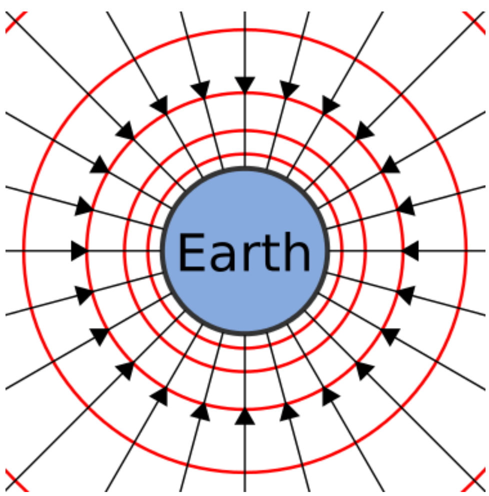

The strength or intensity of fields isn't uniform at all points in the field. The further you are from the source of the field the less effect the field has. The magnitude of Earth's gravitational field decreases as you get further away from the center of the planet.

As we go on to explore electric fields in particular remember how Earth's gravitational field works, both fields share many similarities. Gravitational fields exert a force on objects of mass, and electric fields exert a force on objects of charge.

Electric Fields

Electric fields (e-fields) are an important tool in understanding how electricity begins and continues to flow. Electric fields describe the pulling or pushing force in a space between charges. Compared to Earth's gravitational field, electric fields have one major difference: while Earth's field generally only attracts other objects of mass (since everything is so significantly less massive), electric fields push charges away just as often as they attract them.

The direction of electric fields is always defined as the direction a positive test charge would move if it was dropped in the field. The test charge has to be infinitely small, to keep its charge from influencing the field.

We can begin by constructing electric fields for solitary positive and negative charges. If you dropped a positive test charge near a negative charge, the test charge would be attracted towards the negative charge. So, for a single, negative charge we draw our electric field arrows pointing inward at all directions. That same test charge dropped near another positive charge would result in an outward repulsion, which means we draw arrows going out of the positive charge.

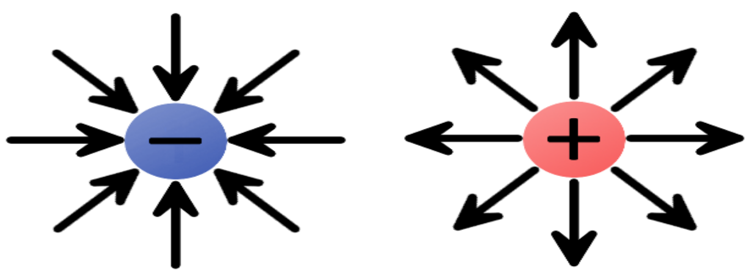

The electric fields of single charges. *A negative charge has an inward electric field because it attracts positive charges. The positive charge has an outward electric field, pushing away like charges.*

Groups of electric charges can be combined to make more complete electric fields.

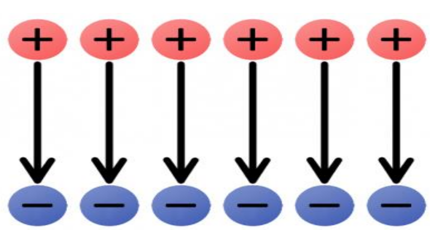

The uniform e-field above points away from the positive charges, towards the negatives. Imagine a tiny positive test charge dropped in the e-field; it should follow the direction of the arrows. As we've seen, electricity usually involves the flow of electrons--negative charges--which flow against electric fields.

Electric fields provide us with the pushing force we need to induce current flow. An electric field in a circuit is like an electron pump: a large source of negative charges that can propel electrons, which will flow through the circuit towards the positive lump of charges.

Electric Potential (Energy)

When we harness electricity to power our circuits, gizmos, and gadgets, we're really transforming energy. Electronic circuits must be able to store energy and transfer it to other forms like heat, light, or motion. The stored energy of a circuit is called electric potential energy.

Energy? Potential Energy?

To understand potential energy we need to understand energy in general. Energy is defined as the ability of an object to do work on another object, which means moving that object some distance. Energy comes in many forms, some we can see (like mechanical) and others we can't (like chemical or electrical). Regardless of what form it's in, energy exists in one of two states: kinetic or potential.

An object has kinetic energy when it's in motion. The amount of kinetic energy an object has depends on its mass and speed. Potential energy, on the other hand, is a stored energy when an object is at rest. It describes how much work the object could do if set into motion. It's an energy we can generally control. When an object is set into motion, its potential energy transforms into kinetic energy.

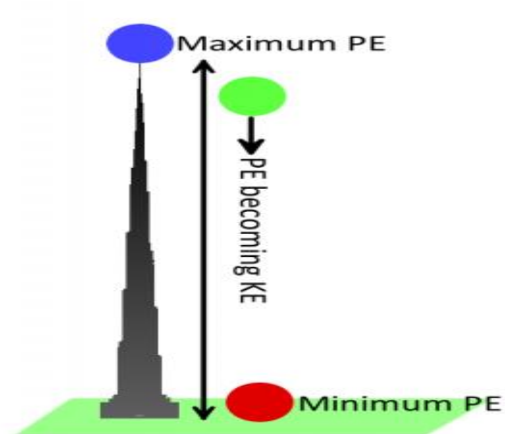

Let's go back to using gravity as an example. A bowling ball sitting motionless at the top of Khalifa tower has a lot of potential (stored) energy. Once dropped, the ball--pulled by the gravitational field--accelerates towards the ground. As the ball accelerates, potential energy is converted into kinetic energy (the energy from motion). Eventually all of the ball's energy is converted from potential to kinetic, and then passed on to whatever it hits. When the ball is on the ground, it has a very low potential energy.

Electric Potential Energy

Just like mass in a gravitational field has gravitational potential energy, charges in an electric field have an electric potential energy. A charge's electric potential energy describes how much stored energy it has, when set into motion by an electrostatic force, that energy can become kinetic, and the charge can do work.

Like a bowling ball sitting at the top of a tower, a positive charge in close proximity to another positive charge has a high potential energy; left free to move, the charge would be repelled away from the like charge. A positive test charge placed near a negative charge would have low potential energy, analogous to the bowling ball on the ground.

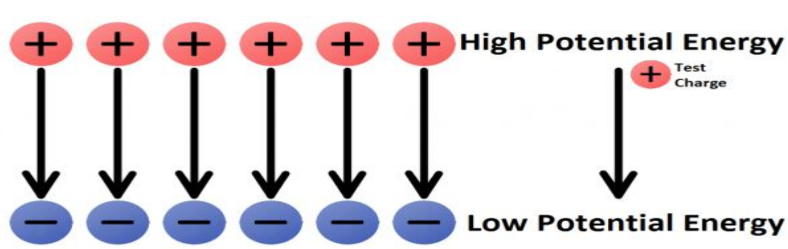

To instill anything with potential energy, we have to do work by moving it over a distance. In the case of the bowling ball, the work comes from carrying it up 163 floors, against the field of gravity. Similarly, work must be done to push a positive charge against the arrows of an electric field (either towards another positive charge, or away from a negative charge). The further up the field the charge goes, the more work you have to do. Likewise, if you try to pull a negative charge away from a positive charge--against an electric field--you have to do work.

For any charge located in an electric field its electric potential energy depends on the type (positive or negative), amount of charge, and its position in the field. Electric potential energy is measured in units of joules (J).

Electric Potential

Electric potential builds upon electric potential energy to help define how much energy is stored in electric fields. It's another concept which helps us model the behavior of electric fields. Electric potential is not the same thing as electric potential energy!

At any point in an electric field the electric potential is the amount of electric potential energy divided by the amount of charge at that point. It takes the charge quantity out of the equation and leaves us with an idea of how much potential energy specific areas of the electric field may provide. Electric potential comes in units of joules per coulomb (J/C), which we define as a volt (V).

In any electric field there are two points of electric potential that are of significant interest to us. There's a point of high potential, where a positive charge would have the highest possible potential energy, and there's a point of low potential, where a charge would have the lowest possible potential energy.

One of the most common terms we discuss in evaluating electricity is voltage. A voltage is the difference in potential between two points in an electric field. Voltage gives us an idea of just how much pushing force an electric field has.

With potential and potential energy under our belt we have all of the ingredients necessary to make current electricity. Let's do it!

Electricity in Action!

After studying particle physics, field theory, and potential energy, we now know enough to make electricity flow. Let's make a circuit!

First we will review the ingredients we need to make electricity:

The definition of electricity is the flow of charge. Usually our charges will be carried by free-flowing electrons.

Negatively-charged electrons are loosely held to atoms of conductive materials. With a little push we can free electrons from atoms and get them to flow in a generally uniform direction.

A closed circuit of conductive material provides a path for electrons to continuously flow.

The charges are propelled by an electric field. We need a source of electric potential (voltage), which pushes electrons from a point of low potential energy to higher potential energy.

A Short Circuit

Batteries are common energy sources which convert chemical energy to electrical energy. They have two terminals, which connect to the rest of the circuit. On one terminal there are an excess of negative charges, while all of the positive charges coalesce on the other. This is an electric potential difference just waiting to act!

If we connected our wire full of conductive copper atoms to the battery, that electric field will influence the negatively-charged free electrons in the copper atoms. Simultaneously pushed by the negative terminal and pulled by the positive terminal, the electrons in the copper will move from atom to atom creating the flow of charge we know as electricity.

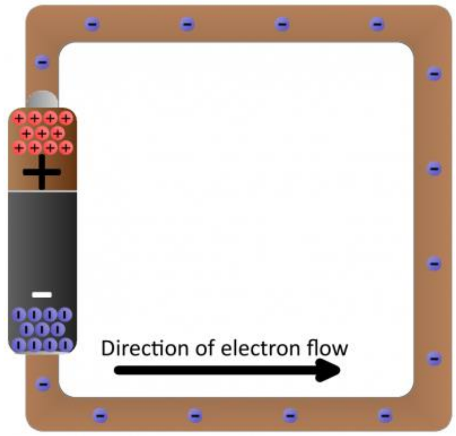

After a second of the current flow, the electrons have actually moved very little--fractions of a centimeter. However, the energy produced by the current flow is huge, especially since there's nothing in this circuit to slow down the flow or consume the energy. Connecting a pure conductor directly across an energy source is a bad idea. Energy moves very quickly through the system and is transformed into heat in the wire, which may quickly turn into melting wire or fire.

Illuminating a Light Bulb

Instead of wasting all that energy, not to mention destroying the battery and wire, let's build a circuit that does something useful! Generally an electric circuit will transfer electric energy into some other form--light, heat, motion, etc. If we connect a light bulb to the battery with wires in between, we have a simple, functional circuit.

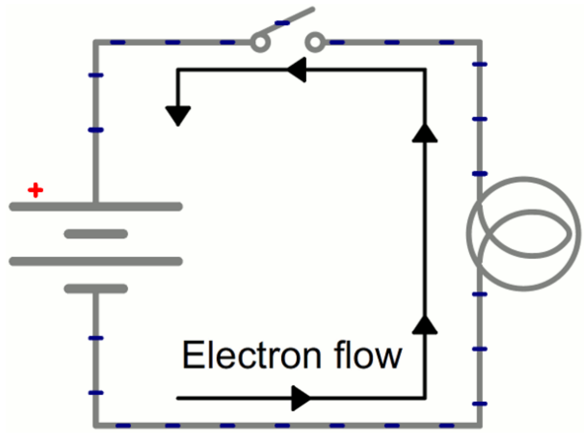

Schematic: A battery (left) connecting to a lightbulb (right), the circuit is completed when the switch (top) closes. With the circuit closed, electrons can flow, pushed from the negative terminal of the battery through the lightbulb, to the positive terminal.

While the electrons move at a snail's pace, the electric field affects the entire circuit almost instantly (we're talking speed of light fast). Electrons throughout the circuit, whether at the lowest potential, highest potential, or right next to the light bulb, are influenced by the electric field. When the switch closes and the electrons are subjected to the electric field, all electrons in the circuit start flowing at seemingly the same time. Those charges nearest the light bulb will take one step through the circuit and start transforming energy from electrical to light (or heat).

Electric Power

Why do we care about power? Power is the measurement of energy transfer over time, and energy costs money. Batteries aren't free, and neither is that stuff coming out of your electrical outlet. So, power measures how fast the pennies are draining out of your wallet!

Also, energy is...energy. It comes in many, potentially harmful, forms -- heat, radiation, sound, nuclear, etc. -- ,and more power means more energy. So, it's important to have an idea of what kind of power you're working with when playing with electronics. Fortunately, in playing with Arduinos, lighting up LEDs, and spinning small motors, losing track of how much power you're using only means smoking a resistor or melting an IC. Nevertheless, Uncle Ben's advice doesn't just apply to superheros.

What is Electric Power?

There are many types of power -- physical, social, super, odor blocking, love -- but in this tutorial, we'll be focusing on electric power. So what is electric power?

In general physics terms, *power* is defined as the **rate at which energy is transferred (or transformed)**.

So, first, what is energy and how is it transferred? It's hard to state simply, but energy is basically the ability of *something* to *move* something else. There are many forms of energy: mechanical, electrical, chemical, electromagnetic, thermal, and many others.

Energy can never be created or destroyed, only transferred to another form. A lot of what we're doing in electronics is converting different forms of energy to and from **electric energy**. Lighting LEDs turns electric energy into electromagnetic energy. Spinning motors turns electric energy into mechanical energy. Buzzing buzzers makes sound energy. Powering a circuit off a 9V alkaline battery turn chemical energy into electrical energy. All of these are forms of **energy transfers**.

Energy type converted	Converted by
Mechanical	Electric Motor
Electromagnetic	LED
Heat	Resistor
Chemical	Battery
Wind	Windmill

Example electric components, which transfer electric energy to another form

Electric energy in particular, begins as electric *potential* energy -- what we lovingly refer to as voltage. When electrons flow through that potential energy, it turns into electric energy. In most useful circuits, that electric energy transforms into some other form of energy. Electric power is measured by combining both **how much** electric energy is transferred, and **how fast** that transfer happens.

Producers and Consumers

Each component in a circuit either **consumes** or **produces** electric energy. A consumer transforms electric energy into another form. For example, when an LED lights up, electric energy is transformed into electromagnetic. In this case, the light bulb **consumes** power. Electric power is **produced** when energy is transferred *to* electric from some other form. A battery supplying power to a circuit is an example of a **power producer**.

Wattage

Energy is measured in terms of joules (J). Since power is a measure of energy over a set amount of time, we can measure it in **joules per second**. The SI unit for joules per second is the **watt** abbreviated as *W*.

$$watt = W = \frac{joule}{second} = \frac{J}{s}$$

It's very common to see "watts" preceded by one of the standard SI prefixes: microwatts (µW), miliwatt (mW), kilowatt (kW), megawatt (MW), and gigawatts (GW), are all common depending on the situation.

Prefix Name	Prefix Abbreviation	Weight
Nanowatt	nW	10^{-9}
Microwatt	µW	10^{-6}
Milliwatt	mW	10^{-3}
Watt	W	10^{0}
Kilowatt	kW	10^{3}
Megawatt	MW	10^{6}
Gigawatt	GW	10^{9}

Microcontrollers, like the Arduino will usually operate in the the µW or mW range. Laptop and desktop computers operate in the standard watt power range. Energy consumption of a house is usually in the kilowatt range. Large stadiums might operate at the megawatt scale. And gigawatts come into play for large-scale power stations and time machines.

Calculating Power

Electric power is the rate at which energy is transferred. It's measured in terms of joules per second (J/s) -- a watt (W). Given the few basic electricity terms we know, how could we calculate power in a circuit? Well, we've got a very standard measurement involving potential energy -- volts (V) -- which are defined in terms of joules per unit of charge (coulomb) (J/C). Current, another of our favorite electricity terms, measures charge flow over time in terms of the ampere (A) -- coulombs per second (C/s). Put the two together and what do we get?! Power!

$$power = volts \times amperes = \frac{joules}{\cancel{coulomb}} \times \frac{\cancel{coulomb}}{second} = watt$$

To calculate the power of any particular component in a circuit, multiply the voltage drop across it by the current running through it.

$$P = VI$$

For Example

Below is a simple (though not all that functional) circuit: a 9V battery connected across a 10Ω resistor.

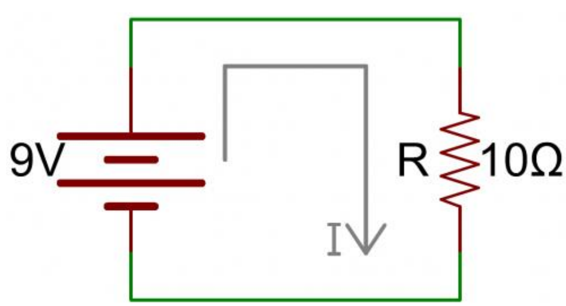

How do we calculate the power across the resistor? First we have to find the current running through it. Easy enough...Ohm's law!

$$I = \frac{V}{R} = \frac{9\,V}{10\,\Omega} = 0.9\,A = 900\,mA$$

Alright, 900mA (0.9A) running through the resistor, and 9V across it. What kind of power is being applied to the resistor then?

$$P = I \times V = 9\,V \times 0.9\,A = 8.1\,W$$

A resistor transforms electric energy into heat. So this circuit transforms 8.1 joules of electric energy to heat every second.

Calculating Power in Resistive Circuits

When it comes to calculating power in a purely resistive circuit, knowing two of three values (voltage, current, and/or resistance) is all you really need.

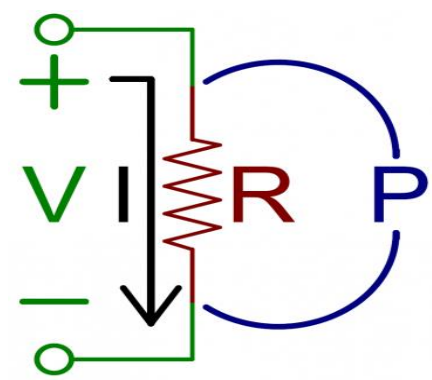

By plugging Ohm's law (V=IR or I=V/R) into our traditional power equation we can create two new equations. The first, purely in terms of voltage and resistance:

$$P = \frac{V^2}{R}$$

So, in our previous example, 9V2/10Ω (V2/R) is 8.1W, and we never have to calculate the current running through the resistor.

A second power equation can be formed solely in terms of current and resistance:

$$P = I^2 \times R$$

Why do we care about the power dropped on a resistor? Or any other component for that matter. Remember that power is the transfer of energy from one type to another. When that electrical energy running from the power source hits the resistor, the energy transforms into heat. Possibly more heat than the resistor can handle. Which leads us to...power ratings.

Power Ratings

All electronic components transfer energy from one type to another. Some energy transfers are desired: LEDs emitting light, motors spinning, batteries charging. Other energy transfers are undesirable, but also unavoidable. These unwanted energy transfers are **power losses**, which usually show up in the form of heat. Too much power loss -- too much heat on a component -- can become *very* undesirable.

Even when energy transfers are the main goal of a component, there'll still be losses to other forms of energy. LEDs and motors, for example, will still produce heat as a byproduct of their other energy transfers.

Most components have a rating for maximum power they can dissipate, and it's important to keep them operating under that value. This'll help you avoid what we lovingly refer to as "letting the magic smoke out".

Resistor Power Ratings

Resistors are some of the more notorious culprits of power loss. When you drop some voltage across a resistor, you're also going to induce current flow across it. More voltage, means more current, means more power.

Remember back to our first power-calculation example, where we found that if 9V were dropped across a 10Ω resistor, that resistor would dissipate 8.1W. 8.1 is a *lot* of watts for most resistors. Most resistors are rated for anywhere from ⅛W (0.125W) to ½W (0.5W). If you drop 8W across a standard ½W resistor, ready a fire extinguisher.

If you've seen resistors before, you've probably seen these. Top is a ½W resistor and below that a ¼W. These aren't built to dissipate very much power.

There are resistors built to handle large power drops. These are specifically called out as **power resistors**.

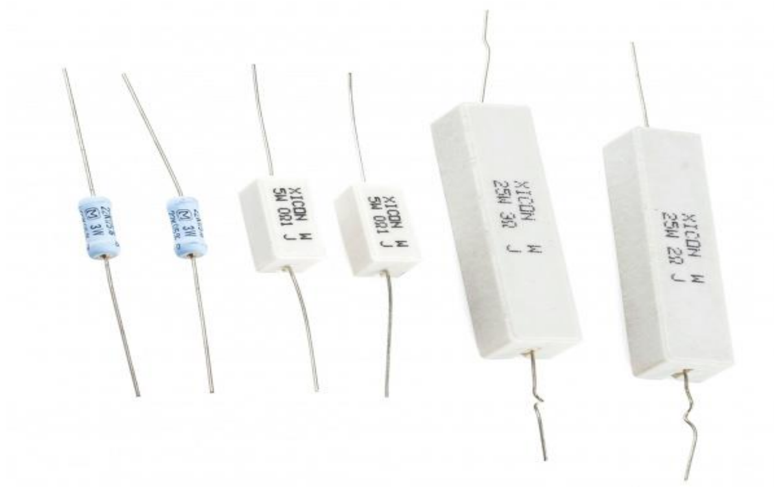

These large resistors are built to dissipate lots of power. From left to right: two 3W 22kΩ resistors, two 5W 0.1Ω resistors, and 25W 3Ω and 2Ω resistors.

If you ever find yourself picking out a resistor value. Keep it’s power rating in mind as well. And, unless your goal is to heat something up (heating elements are basically really high-power resistors), try to minimize power loss in a resistor.

For Example

Resistor power ratings can come into play when you’re trying to decide on a value for an LED current-limiting resistor. Say, for example, you want to light up a 10mm super-bright red LED at maximum brightness, using a 9V battery.

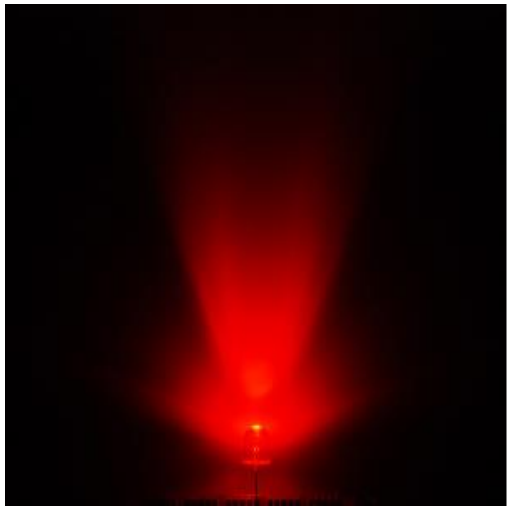

That LED has a maximum forward current of 80mA, and a forward voltage of about 2.2V. So to deliver 80mA to the LED, you'd need an 85Ω resistor to do so.

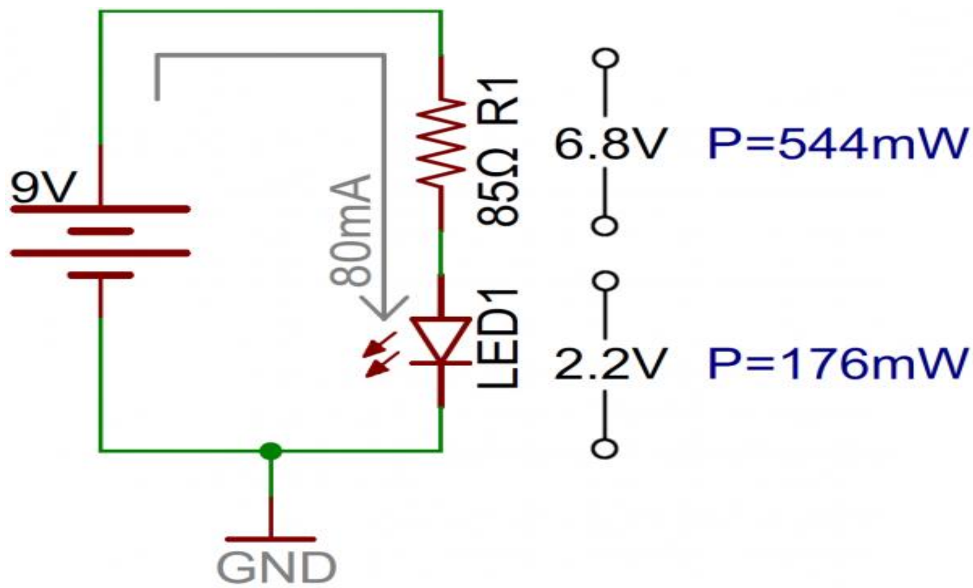

6.8V dropped on the resistor, and 80mA running through it means 0.544W (6.8V*0.08A) of power lost on it. A half-watt resistor isn't going to like that very much! It probably won't melt, but it'll get hot. Play it safe and move up to a 1W resistor (or save power and use a dedicated LED driver).

Resistors certainly aren't the only components where maximum power ratings must be considered. Any component with a resistive property to it is going to produce thermal power losses. Working with components that are commonly subjected to high power -- voltage regulators, diodes, amplifiers, and motor drivers, for example -- means paying extra special attention to power loss and thermal stress.

Electricity Basics

When beginning to explore the world of electricity and electronics, it is vital to start by understanding the basics of voltage, current, and resistance. These are the three basic building blocks required to manipulate and utilize electricity. At first, these concepts can be difficult to understand because we cannot "see" them. One cannot see with the naked eye the energy flowing through a wire or the voltage of a battery sitting on a table. Even the lightning in the sky, while visible, is not truly the energy exchange happening from the clouds to the earth, but a reaction in the air to the energy passing through it. In order to detect this energy transfer, we must use measurement tools such as multimeters, spectrum analyzers, and oscilloscopes to visualize what is happening with the charge in a system. Fear not, however, this tutorial will give you the basic understanding of voltage, current, and resistance and how the three relate to each other.

Georg Ohm

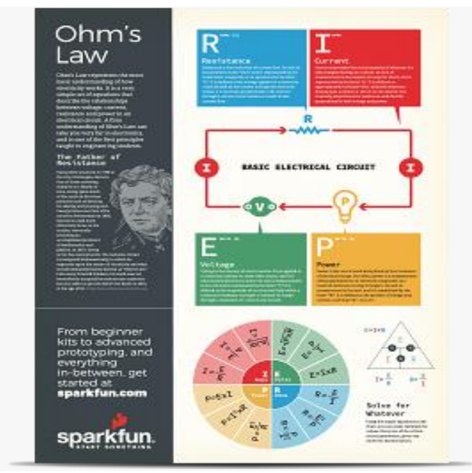

Electrical Charge

Electricity is the movement of electrons. Electrons create charge, which we can harness to do work. Your light bulb, your stereo, your phone, etc., are all harnessing the movement of the electrons in order to do work. They all operate using the same basic power source: the movement of electrons.

The three basic principles for this tutorial can be explained using electrons, or more specifically, the charge they create:

- Voltage is the difference in charge between two points.
- Current is the rate at which charge is flowing.
- Resistance is a material's tendency to resist the flow of charge (current).

So, when we talk about these values, we're really describing the movement of charge, and thus, the behavior of electrons. A circuit is a closed loop that allows charge to move from one place to another. Components in the circuit allow us to control this charge and use it to do work.

Georg Ohm was a Bavarian scientist who studied electricity. Ohm starts by describing a unit of resistance that is defined by current and voltage. So, let's start with voltage and go from there.

Voltage

We define voltage as the amount of potential energy between two points on a circuit. One point has more charge than another. This difference in charge between the two points is called voltage. It is measured in volts, which, technically, is the potential energy difference between two points that will impart one joule of energy per coulomb of charge that passes through it (don't panic if this makes no sense, all will be explained). The unit "volt" is named after the Italian physicist Alessandro Volta who invented what is considered the first chemical battery. Voltage is represented in equations and schematics by the letter "V".

When describing voltage, current, and resistance, a common analogy is a water tank. In this analogy, charge is represented by the water amount, voltage is represented by the water pressure, and current is represented by the water flow. So for this analogy, remember:

- Water = Charge
- Pressure = Voltage
- Flow = Current

Consider a water tank at a certain height above the ground. At the bottom of this tank there is a hose.

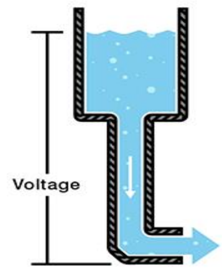

The pressure at the end of the hose can represent voltage. The water in the tank represents charge. The more water in the tank, the higher the charge, the more pressure is measured at the end of the hose.

We can think of this tank as a battery, a place where we store a certain amount of energy and then release it. If we drain our tank a certain amount, the pressure created at the end of the hose goes down. We can think of this as decreasing voltage, like when a flashlight gets dimmer as the batteries run down. There is also a decrease in the amount of water that will flow through the hose. Less pressure means less water is flowing, which brings us to current.

Current

We can think of the amount of water flowing through the hose from the tank as current. The higher the pressure, the higher the flow, and vice-versa. With water, we would measure the volume of the water flowing through the hose over a certain period of time. With electricity, we measure the amount of charge flowing through the circuit over a period of time. Current is measured in Amperes (usually just referred to as "Amps"). An ampere is defined as 6.241*10^18 electrons (1 Coulomb) per second passing through a point in a circuit. Amps are represented in equations by the letter "I".

Let's say now that we have two tanks, each with a hose coming from the bottom. Each tank has the exact same amount of water, but the hose on one tank is narrower than the hose on the other.

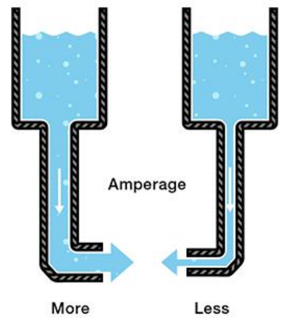

We measure the same amount of pressure at the end of either hose, but when the water begins to flow, the flow rate of the water in the tank with the narrower hose will be less than the flow rate of the water in the tank with the wider hose. In electrical terms, the current through the narrower hose is less than the current through the wider hose. If we want the flow to be the same through both hoses, we have to increase the amount of water (charge) in the tank with the narrower hose.

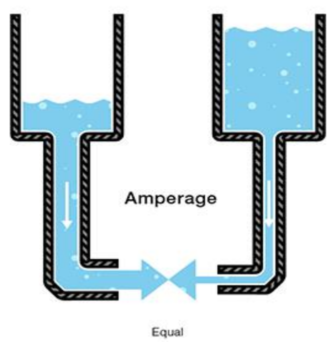

This increases the pressure (voltage) at the end of the narrower hose, pushing more water through the tank. This is analogous to an increase in voltage that causes an increase in current.

Now we're starting to see the relationship between voltage and current. But there is a third factor to be considered here: the width of the hose. In this analogy, the width of the hose is the resistance. This means we need to add another term to our model:

- Water = Charge (measured in Coulombs)

- Pressure = Voltage (measured in Volts)
- Flow = Current (measured in Amperes, or "Amps" for short)
- Hose Width = Resistance

Resistance

Consider again our two water tanks, one with a narrow pipe and one with a wide pipe.

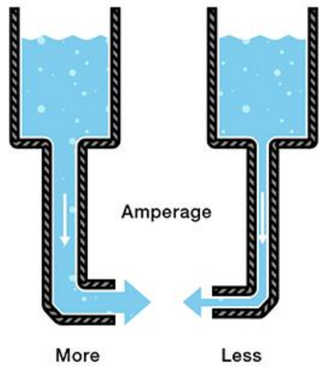

It stands to reason that we can't fit as much volume through a narrow pipe than a wider one at the same pressure. This is resistance. The narrow pipe "resists" the flow of water through it even though the water is at the same pressure as the tank with the wider pipe.

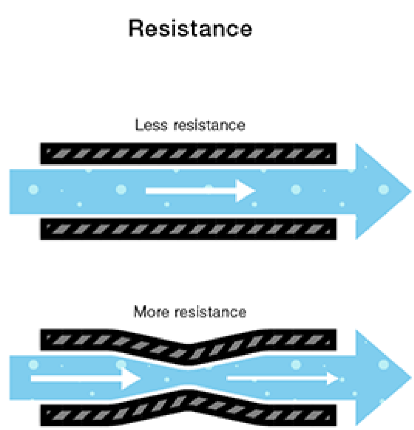

In electrical terms, this is represented by two circuits with equal voltages and different resistances. The circuit with the higher resistance will allow less charge to flow, meaning the circuit with higher resistance has less current flowing through it.

This brings us back to Georg Ohm. Ohm defines the unit of resistance of "1 Ohm" as the resistance between two points in a conductor where the application of 1 volt will push 1 ampere, or 6.241×10^18 electrons. This value is usually represented in schematics with the greek letter "Ω", which is called omega, and pronounced "ohm".

Ohm's Law

Combining the elements of voltage, current, and resistance, Ohm developed the formula:

$$V = I \cdot R$$

Where

- V = Voltage in volts
- I = Current in amps
- R = Resistance in ohms

This is called Ohm's law. Let's say, for example, that we have a circuit with the potential of 1 volt, a current of 1 amp, and resistance of 1 ohm. Using Ohm's Law we can say:

$$1V = 1A \cdot 1\Omega$$

Let's say this represents our tank with a wide hose. The amount of water in the tank is defined as 1 volt and the "narrowness" (resistance to flow) of the hose is defined as 1 ohm. Using Ohms Law, this gives us a flow (current) of 1 amp.

Using this analogy, let's now look at the tank with the narrow hose. Because the hose is narrower, its resistance to flow is higher. Let's define this resistance as 2 ohms. The amount of water in the tank is the same as the other tank, so, using Ohm's Law, our equation for the tank with the narrow hose is

$$1V = ?A \cdot 2\Omega$$

But what is the current? Because the resistance is greater, and the voltage is the same, this gives us a current value of 0.5 amps:

$$1V = 0.5A \cdot 2\Omega$$

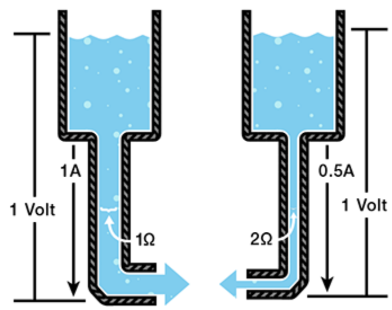

So, the current is lower in the tank with higher resistance. Now we can see that if we know two of the values for Ohm's law, we can solve for the third. Let's demonstrate this with an experiment.

An Ohm's Law Experiment

For this experiment, we want to use a 9 volt battery to power an LED. LEDs are fragile and can only have a certain amount of current flowing through them before they burn out. In the documentation for an LED, there will always be a "current rating". This is the maximum amount of current that can flow through the particular LED before it burns out.

Materials Required

In order to perform the experiments listed at the end of the tutorial, you will need:

- A multimeter
- A 9-Volt battery
- A 560-Ohm resistor(or the next closest value)
- An LED

NOTE: LEDs are what's known as a "non-ohmic" devices. This means that the equation for the current flowing through the LED itself is not as simple as V=IR.

The LED introduces something called a "voltage drop" into the circuit, thus changing the amount of current running through it. However, in this experiment we are simply trying to protect the LED from over-current, so we will neglect the current characteristics of the LED and choose the resistor value using Ohm's Law in order to be sure that the current through the LED is safely under 20mA.

For this example, we have a 9 volt battery and a red LED with a current rating of 20 milliamps, or 0.020 amps. To be safe, we'd rather not drive the LED at its maximum current but rather its suggested current, which is listed on its datasheet as 18mA, or 0.018 amps. If we simply connect the LED directly to the battery, the values for Ohm's law look like this:

$$V = I \cdot R$$

therefore:

$$I = \frac{V}{R}$$

and since we have no resistance yet:

$$I = \frac{9V}{0R}$$

Dividing by zero gives us infinite current! Well, not infinite in practice, but as much current as the battery can deliver. Since we do NOT want that much current flowing through our LED, we're going to need a resistor. Our circuit should look like this:

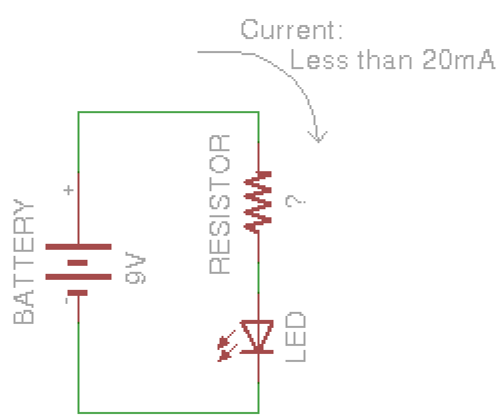

We can use Ohm's Law in the exact same way to determine the resistor value that will give us the desired current value:

$$V = I \cdot R$$

Therefore:

$$R = \frac{V}{I}$$

plugging in our values:

$$R = \frac{9V}{0.018A}$$

solving for resistance:

$$R = 500\Omega$$

So, we need a resistor value of around 500 ohms to keep the current through the LED under the maximum current rating.

500 ohms is not a common value for off-the-shelf resistors, so this device uses a 560 ohm resistor in its place. Here's what our device looks like all put together.

Success! We've chosen a resistor value that is high enough to keep the current through the LED below its maximum rating, but low enough that the current is sufficient to keep the LED nice and bright.

This LED/current-limiting resistor example is a common occurrence in hobby electronics. You'll often need to use Ohm's Law to change the amount of current flowing through the circuit. Another example of this implementation is seen in the LilyPad LED boards.

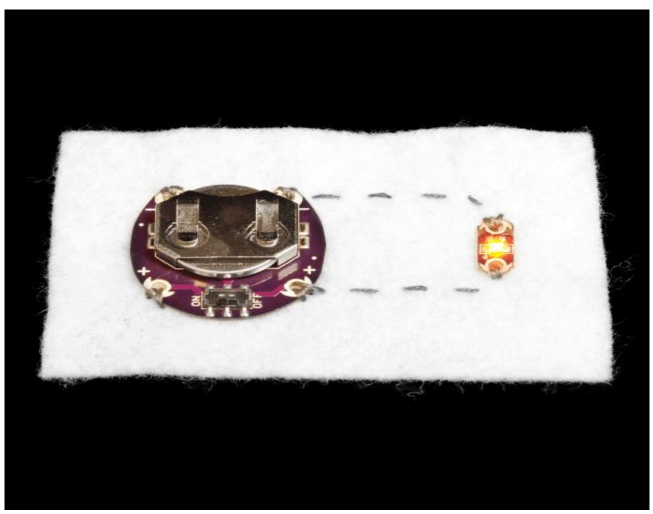

With this setup, instead of having to choose the resistor for the LED, the resistor is already on-board with the LED so the current-limiting is accomplished without having to add a resistor by hand.

Current Limiting Before or After the LED?

To make things a little more complicated, you can place the current limiting resistor on either side of the LED, and it will work just the same!

Many folks learning electronics for the first time struggle with the idea that a current limiting resistor can live on either side of the LED and the circuit will still function as usual.

Imagine a river in a continuous loop, an infinite, circular, flowing river. If we were to place a dam in it, the entire river would stop flowing, not just one side. Now imagine we place a water wheel in the river which slows the flow of the river. It wouldn't matter where in the circle the water wheel is placed, it will still slow the flow on the entire river.

This is an oversimplification, as the current limiting resistor cannot be placed anywhere in the circuit; it can be placed on either side of the LED to perform its function.

For a more scientific answer, we turn to Kirchoff's Voltage Law. It is because of this law that the current limiting resistor can go on either side of the LED and still have the same effect.

What is a Circuit?

.

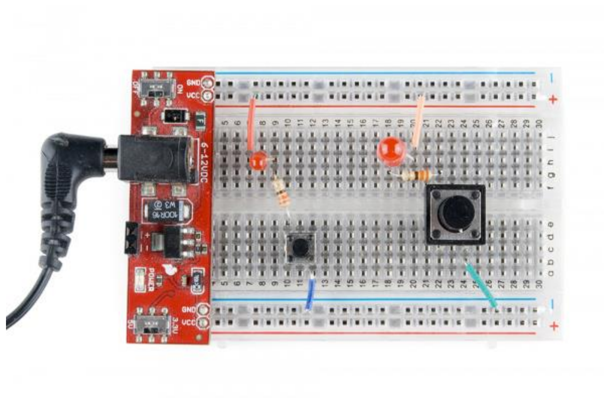

A simple circuit, involving a button, an LED, and a resistor, built two different ways.

Circuit Basics

Voltage and How it Works

You've probably heard that a battery or a wall outlet has a certain number of **volts**. This is a measurement of the electrical **potential** produced by the battery, or the utility grid connected to the wall outlet.

All those volts are sitting there waiting for you to use them, but there's a catch: **in order for electricity to do any work, it needs to be able to move**. It's kind of like a blown-up balloon; if you pinch it off, there is air in there that *could* do something if it's released, but it won't actually do anything until you let it out.

Unlike air coming out of a balloon, electricity can only flow through materials that can conduct electricity, such as copper wire. If you connect a wire to a battery or wall outlet (**WARNING:** the voltage in a wall outlet is dangerous, don't do this!), you will be giving the electricity a path to follow. But if the wire isn't connected to anything else, the electricity won't have anywhere to go and still won't move.

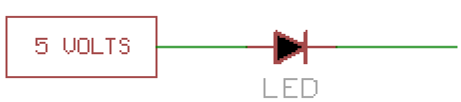

What makes electricity move? **Electricity wants to flow from a higher voltage to a lower voltage.** This is exactly like the balloon: the pressurized air in the balloon wants to flow from inside the balloon (higher pressure) to outside the balloon (lower pressure). If you create a conductive path between a higher voltage and a lower voltage, electricity will flow along that path. And if you insert something useful into that path like an LED, the flowing electricity will do some work for you, like lighting up that LED. Huzzah!

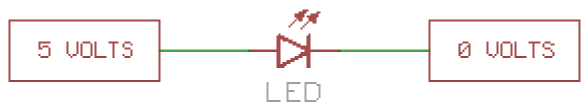

So, where do you find a higher voltage and a lower voltage? Here's something really useful to know: **every source of electricity has two sides**. You can see this on batteries, which have metal caps on both ends, or your wall outlet that has two (or more) holes. In batteries and other DC (Direct Current) voltage sources, these sides (often called **terminals**) are named **positive** (or "+"), and **negative** (or "-").

Why does every source of electricity have two sides? This goes back to the idea of “potential”, and that you need a voltage difference in order to get electricity to flow. It sounds silly, but you can’t have a difference without two things to be different. In any power supply, the positive side will have a higher voltage than the negative side, which is exactly what we want. In fact, when we measure voltage, we usually say that the negative side is 0 volts, and the positive side is however many volts the supply can provide.

Electrical sources are like pumps. Pumps always have two sides, an outlet that blows something out, and an inlet that sucks something in. Batteries and generators and solar panels work the same way. Something inside them is hard at work moving electricity towards the outlet (the positive side), but all that electricity leaving the device creates a void, which means that the negative side needs to pull electricity in to replace it.*

What have we learned so far?

- Voltage is potential, but electricity needs to flow to do anything useful.
- Electricity needs a path to flow through, which must be an electrical conductor such as copper wire.
- Electricity will flow from a higher voltage to a lower voltage.
- DC voltage sources always have two sides, called positive and negative, with the positive side a higher voltage than the negative side.

The Simplest Circuit

We're finally ready to make electricity work for us! If we connect the positive side of a voltage source, through something that does some work such as a Light Emitting Diode (LED), and back to the negative side of the voltage source; electricity, or **current**, will flow. And we can put things in the path that do useful things when current flows through them, like LEDs that light up.

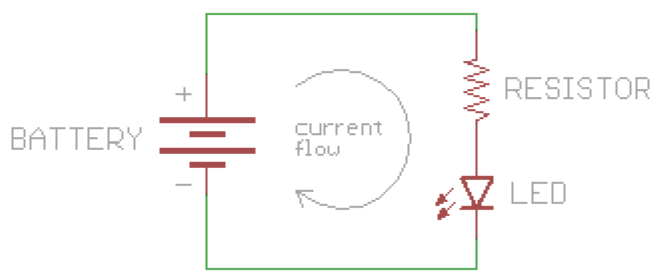

This circular path, which is always required to get electricity to flow and do something useful, is called a circuit. A circuit is a path that starts and stops at the same place, which is exactly what we're doing.

Short and Open Circuits

What is a “Load”?

The reason we want to build circuits is to make electricity do useful things for us. The way we do that is by putting things in the circuit that use the current flow to light up, make noise, run programs, etc.

These things are called **loads**, because they “load down” the power supply, just like you’re “loaded down” when you’re carrying something. The same way you could be loaded down with too much weight, it’s possible to load down a power supply too much, which will slow down the current flow. But unlike you, it’s also possible to load down a circuit too little - this may let too much current flow (imagine running too fast if you weren’t carrying any weight), which can burn out your parts or even the power supply.

You’ll learn all about voltage, current, and loads in the next tutorial: Voltage, Current, Resistance, and Ohm's Law. But for now, let’s learn about two special cases of circuit: **short circuit**, and **open circuit**. Knowing about these will help tremendously when you're troubleshooting your own circuits.

Short Circuit

DON'T DO THIS, but if you connect a wire directly from the positive to the negative side of a power supply, you'll create what is called a **short circuit**. This is a very bad idea.

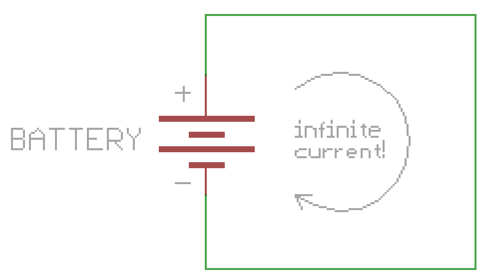

This seems like the best possible circuit, so why is it a bad idea? Remember that electrical current wants to flow from a higher voltage to a lower voltage, and if you put a load into the current, you can do something useful like light up an LED.

If you DO have a load in the current, the current flow through your circuit will be limited to that which your device consumes, which is usually a very small amount. However, if you DON'T put anything in to restrict the current flow, there won't be anything to slow down the current, and it will try to be infinite!

Your power supply can't provide infinite current, but it will provide as much as it can, which may be a lot. This could cause your wire to burn up, damage the power supply, drain your battery, or other exciting things. Most of the time your power supply will have some sort of safety mechanism built into it to limit the maximum current in the event of a short circuit, but not always. This is the reason all homes and buildings have circuit breakers, to prevent fires from starting in the event of a short circuit somewhere in the wiring.

A closely related problem is accidentally letting too much current flow through part of your circuit, causing a part to burn up. This isn't quite a short circuit, but it's close. This most often happens when you use the incorrect **resistor** value, which lets too much current flow through another component such as an LED.

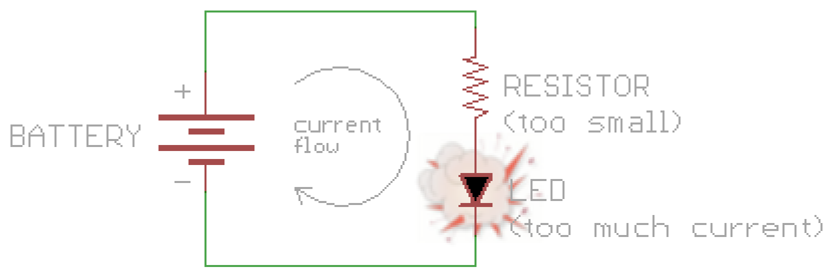

The bottom line: **if you notice that things are suddenly becoming hot or a part suddenly burns out, immediately turn off the power and look for possible short circuits.**

Open Circuit

The opposite of a short circuit is an **open circuit**. This is a circuit where the loop isn't fully connected (and therefore this isn't really a circuit at all).

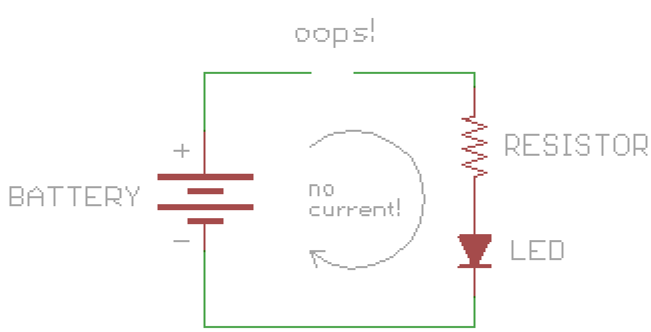

Unlike the short circuit above, nothing will get hurt by this "circuit", but your circuit won't work either. If you're new at circuits, it can often be hard to find where the break is, especially if you're using breadboards where all the conductors are hidden.

If your circuit doesn't work, the most likely cause is an open circuit. This is usually due to a broken connection or a loose wire. (Short circuits can steal all the power from the rest of your circuit, so be sure to look for those as well.)

TIP: if you can't easily find where your circuit is open, a multimeter can be very useful tool. If you set it to measure volts, you can use it to check the voltage at various points in your powered circuit, and eventually find the point where voltage isn't getting through.

Series and Parallel Circuits

Series and Parallel Circuits

Simple circuits (ones with only a few components) are usually fairly straightforward for beginners to understand. But, things can get sticky when other components come to the party. Where's the current going? What's the voltage doing? Can this be simplified for easier understanding? Fear not, intrepid reader. Valuable information follows.

In this tutorial, we'll first discuss the difference between series circuits and parallel circuits, using circuits containing the most basic of components -- resistors and batteries -- to show the difference between the two configurations. We'll then explore what happens in series and parallel circuits when you combine different types of components, such as capacitors and inductors.

Series Circuits

Nodes and Current Flow

Before we get too deep into this, we need to mention what a **node** is. It's nothing fancy, just representation of an electrical junction between two or more components. When a circuit is modeled on a schematic, these nodes represent the wires between components.

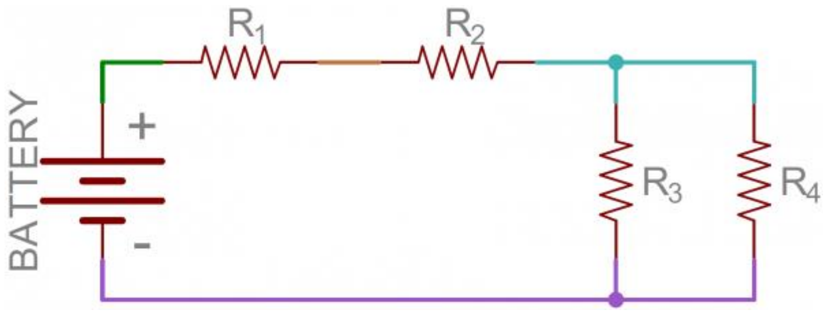

Example schematic with four uniquely colored nodes.

That's half the battle towards understanding the difference between series and parallel. We also need to understand **how current flows** through a circuit. Current flows from a high voltage to a lower voltage in a circuit. Some amount of current will flow through every path it can take to get to the point of lowest voltage (usually called ground). Using the above circuit as an example, here's how current would flow as it runs from the battery's positive terminal to the negative:

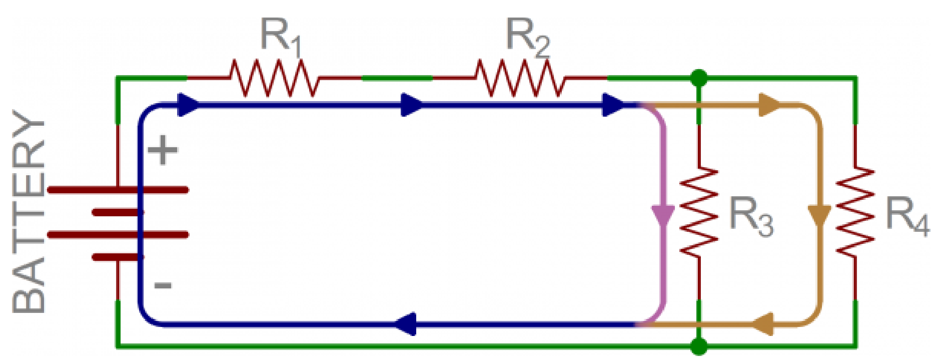

Current (indicated by the blue, orange, and pink lines) flowing through the same example circuit as above. Different currents are indicated by different colors.

Notice that in some nodes (like between R1 and R2) the current is the same going in as at is coming out. At other nodes (specifically the three-way junction between R2, R3, and R4) the main (blue) current splits into two different ones. *That's* the key difference between series and parallel!

Series Circuits Defined

Two components are in series if they share a common node and if the **same current** flows through them. Here's an example circuit with three series resistors:

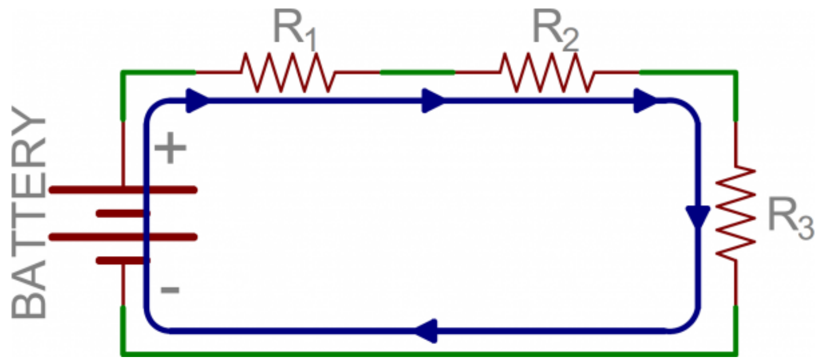

There's only one way for the current to flow in the above circuit. Starting from the positive terminal of the battery, current flow will first encounter R1. From there the current will flow straight to R2, then to R3, and finally back to the negative terminal of the battery. Note that there is only one path for current to follow. These components are in series.

Parallel Circuits

Parallel Circuits Defined

If components share *two* common nodes, they are in parallel. Here's an example schematic of three resistors in parallel with a battery:

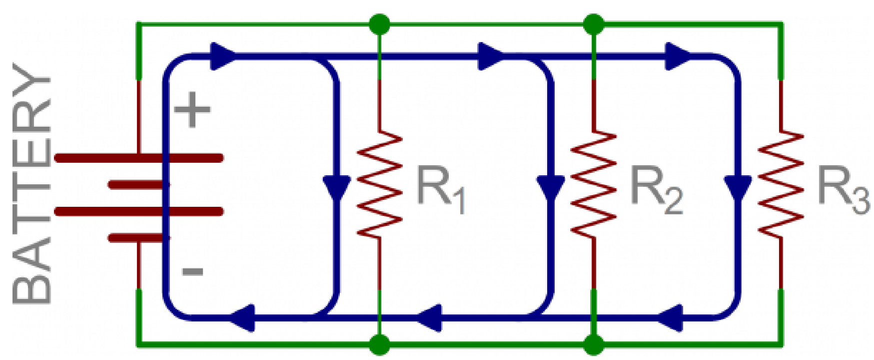

From the positive battery terminal, current flows to R1... and R2, and R3. The node that connects the battery to R1 is also connected to the other resistors. The other ends of these resistors are similarly tied together, and then tied back to the negative terminal of the battery. There are three distinct paths that current can take before returning to the battery, and the associated resistors are said to be in parallel.

Where series components all have equal currents running through them, parallel components all have the same voltage drop across them -- series:current::parallel:voltage.

Series and Parallel Circuits Working Together

From there we can mix and match. In the next picture, we again see three resistors and a battery. From the positive battery terminal, current first encounters R1. But, at the other side of R1 the node splits, and current can go to both R2 and R3. The current paths through R2 and R3 are then tied together again, and current goes back to the negative terminal of the battery.

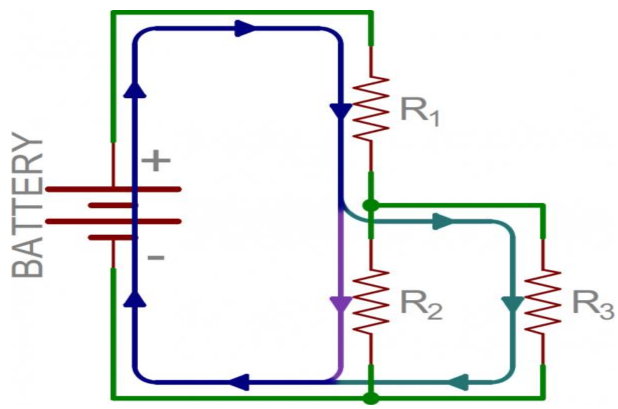

In this example, R2 and R3 are in parallel with each other, and R1 is in series with the parallel combination of R2 and R3.

Calculating Equivalent Resistances in Series Circuits

Here's some information that may be of some more practical use to you. When we put resistors together like this, in series and parallel, we change the way current flows through them. For example, if we have a 10V supply across a 10kΩ resistor, Ohm's law says we've got 1mA of current flowing.

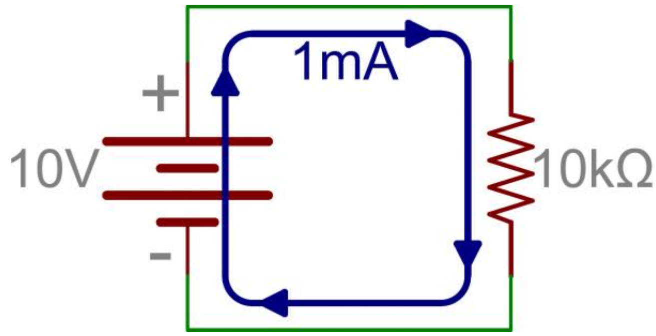

If we then put another 10kΩ resistor in series with the first and leave the supply unchanged, we've cut the current in half because the resistance is doubled.

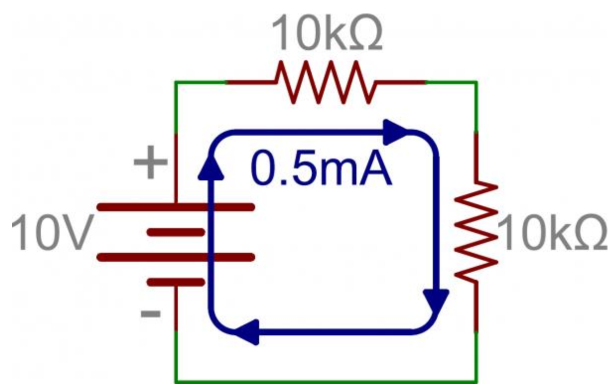

In other words, there's still only one path for current to take and we just made it even harder for current to flow. How much harder? 10kΩ + 10kΩ = 20kΩ. And, that’s how we calculate resistors in series -- just **add their values**.

To put this equation more generally: the total resistance of *N* -- some arbitrary number of -- resistors is their total sum.

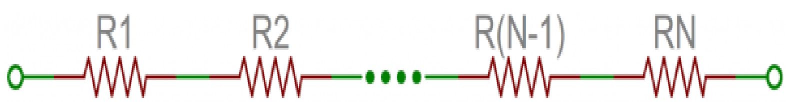

$$R_{tot} = R_1 + R_2 + \ldots + R_{N-1} + R_N$$

Calculating Equivalent Resistances in Parallel Circuits

What about parallel resistors? That’s a bit more complicated, but not by much. Consider the last example where we started with a 10V supply and a 10kΩ resistor, but this time we add another 10kΩ in parallel instead of series. Now there are two paths for current to take. Since the supply voltage didn’t change, Ohm's Law says the first resistor is still going to draw 1mA. But, so is the second resistor, and we now have a total of 2mA coming from the supply, doubling the original 1mA. This implies that we've cut the total resistance in half.

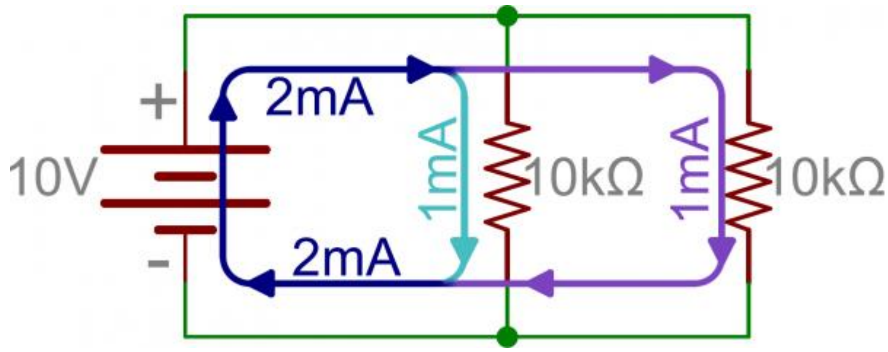

While we can say that 10kΩ || 10kΩ = 5kΩ ("||" roughly translates to "in parallel with"), we're not always going to have 2 identical resistors. What then?

The equation for adding an arbitrary number of resistors in parallel is:

$$\frac{1}{R_{tot}} = \frac{1}{R_1} + \frac{1}{R_2} + \ldots + \frac{1}{R_{N-1}} + \frac{1}{R_N}$$

If reciprocals aren't your thing, we can also use a method called "product over sum" when we have two resistors in parallel:

$$R_{tot} = \frac{R_1 \cdot R_2}{R_1 + R_2}$$

However, this method is only good for two resistors in one calculation. We can combine more than 2 resistors with this method by taking the result of R1 || R2 and calculating that value in parallel with a third resistor (again as product over sum), but the reciprocal method may be less work.

Experiment Time - Part 1

What you'll need:

- A handful of 10kΩ resistors
- A multimeter
- A breadboard

Let's try a simple experiment just to prove that these things work the way we're saying they do.

First, we're going to hook up some 10kΩ resistors in series and watch them add in a most un-mysterious way. Using a breadboard, place one 10kΩ resistor as indicated in the figure and measure with a multimeter. Yes, we already know it's going to say it's 10kΩ, but this is what we in the biz call a "sanity check". Once we've convinced ourselves that the world hasn't changed significantly since we last looked at it, place another one in similar fashion but with a lead from each resistor connecting electrically through the breadboard and measure again. The meter should now say something close to 20kΩ.

You may notice that the resistance you measure might not be exactly what the resistor says it should be. Resistors have a certain amount of **tolerance**, which means they can be off by a certain percentage in either direction. Thus, you may read 9.99kΩ or 10.01kΩ. As long as it's close to the correct value, everything should work fine.

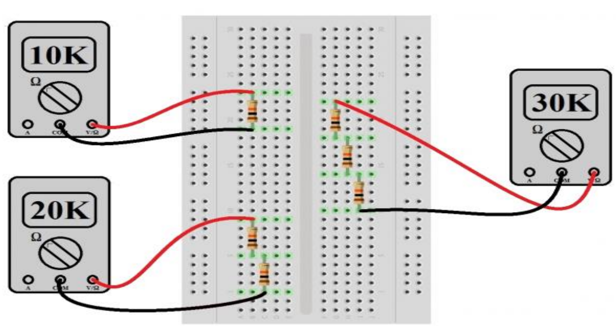

The reader should continue this exercise until convincing themselves that they know what the outcome will be before doing it again, or they run out of resistors to stick in the breadboard, whichever comes first.

Experiment Time - Part 2

Now let's try it with resistors in a **parallel** configuration. Place one 10kΩ resistor in the breadboard as before (we'll trust that the reader already believes that a single 10kΩ resistor is going to measure something close to 10kΩ on the multimeter). Now place a second 10kΩ resistor next to the first, taking care that the leads of each resistor are in electrically connected rows. But before measuring the combination, calculate by either product-over-sum or reciprocal methods what the new value should be (hint: it's going to be 5kΩ). Then measure. Is it something close to 5kΩ? If it's not, double check the holes into which the resistors are plugged.

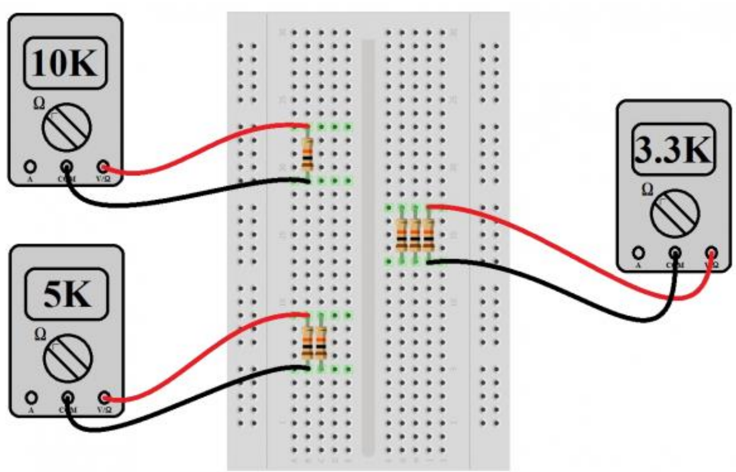

Repeat the exercise now with 3, 4 and 5 resistors. The calculated/measured values should be 3.33kΩ, 2.5kΩ and 2kΩ, respectively. Did everything come out as planned? If not, go back and check your connections. If it did, EXCELSIOR! Go have a milkshake before we continue. You've earned it.

Rules of Thumb for Series and Parallel Resistors

There are a few situations that may call for some creative resistor combinations. For example, if we're trying to set up a very specific reference voltage you'll almost always need a very specific ratio of resistors whose values are unlikely to be "standard" values. And while we can get a very high degree of precision in resistor values, we may not want to wait the X number of days it takes to ship something, or pay the price for non-stocked, non-standard values. So in a pinch, we can always build our own resistor values.

Tip #1: Equal Resistors in Parallel

Adding *N* like-valued resistors *R* in parallel gives us *R/N* ohms. Let's say we need a 2.5kΩ resistor, but all we've got is a drawer full of 10kΩ's. Combining four of them in parallel gives us 10kΩ/4 = 2.5kΩ.

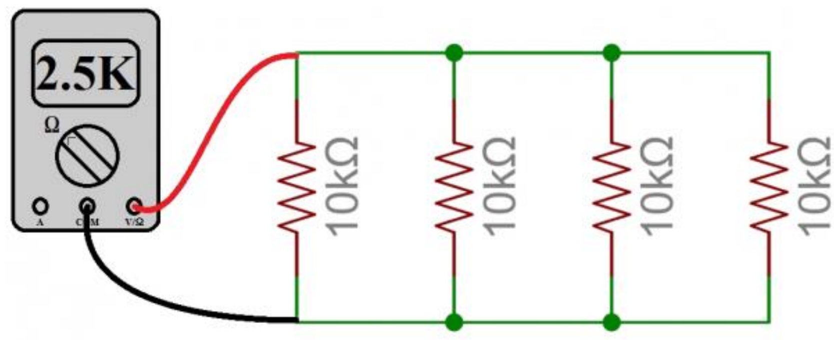

Total Resistance = Value of One Resistor / Number of Resistors
2.5K = 10K / 4

Tip #2: Tolerance

Know what kind of tolerance you can tolerate. For example, if you needed a 3.2kΩ resistor, you could put 3 10kΩ resistors in parallel. That would give you 3.3kΩ, which is about a 4% tolerance from the value you need. But, if the circuit you're building needs to be closer than 4% tolerance, we can measure our stash of 10kΩ's to see which are lowest values because they have a tolerance, too. In theory, if the stash of 10kΩ resistors are all 1% tolerance, we can only get to 3.3kΩ. But part manufacturers are known to make just these sorts of mistakes, so it pays to poke around a bit.

Tip #3: Power Ratings in Series/Parallel

This sort of series and parallel combination of resistors works for power ratings, too. Let's say that we need a 100Ω resistor rated for 2 watts (W), but all we've got is a bunch of 1kΩ quarter-watt (¼W) resistors (and it's 3am, all the Mountain Dew is gone, and the coffee's cold). You can combine 10 of the 1kΩ's to get 100Ω

(1kΩ/10 = 100Ω), and the power rating will be 10x0.25W, or 2.5W. Not pretty, but it will get us through a final project, and might even get us extra points for being able to think on our feet.

We need to be a little more careful when we combine resistors of dissimilar values in parallel where total equivalent resistance and power ratings are concerned. It should be completely obvious to the reader, but...

Tip #4: Different Resistors in Parallel

The combined resistance of two resistors of different values is always less than the smallest value resistor. The reader would be amazed at how many times someone combines values in their head and arrives at a value that's halfway between the two resistors (1kΩ || 10kΩ does NOT equal anything around 5kΩ!). The total parallel resistance will always be dragged closer to the lowest value resistor. Do yourself a favor and read tip #4 10 times over.

Tip #5: Power Dissipation in Parallel

The power dissipated in a parallel combination of dissimilar resistor values is not split evenly between the resistors because the currents are not equal. Using the previous example of (1kΩ || 10kΩ), we can see that the 1kΩ will be drawing 10X the current of the 10kΩ. Since Ohm's Law says power = voltage x current, it follows that the 1kΩ resistor will dissipate 10X the power of the 10kΩ.

Ultimately, the lessons of tips 4 and 5 are that we have to pay closer attention to what we're doing when combining resistors of dissimilar values in parallel. But tips 1 and 3 offer some handy shortcuts when the values are the same.

Series and Parallel Capacitors

Combining capacitors is just like combining resistors...only the opposite. As odd as that sounds, it's absolutely true. Why would this be?

A capacitor is just two plates spaced very close together, and it's basic function is to hold a whole bunch of electrons. The greater the value of capacitance, the more electrons it can hold. If the size of the plates is increased, the capacitance goes up because there's physically more space for electrons to hang out. And if the plates are moved farther apart, the capacitance goes down, because the electric field strength between them goes down as the distance goes up.

Now let’s say we've got two 10µF capacitors wired together in series, and let’s say they’re both charged up and ready discharge into the friend sitting next to you.

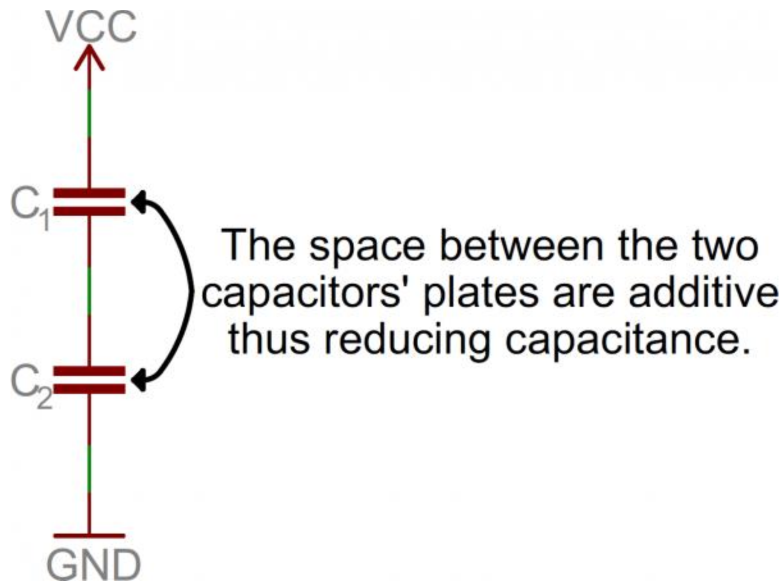

Remember that in a series circuit there's only one path for current to flow. It follows that the number of electrons that are discharging from the cap on the bottom is going to be the same number of electrons coming out of the cap on the top. So the capacitance hasn't increased, has it?

In fact, it’s even worse than that. By placing the capacitors in series, we've effectively spaced the plates farther apart because the spacing between the plates of the two capacitors adds together. So we don't have 20µF, or even 10µF. We’ve got 5µF. The upshot of this is that we add series capacitor values the same way we add parallel resistor values. Both the product-over-sum and reciprocal methods are valid for adding capacitors in series.

$$\frac{1}{C_{Tot}} = \frac{1}{C_1} + \frac{1}{C_2} + \ldots + \frac{1}{C_{N-1}} + \frac{1}{C_N}$$

It may seem that there's no point to adding capacitors in series. But it should be pointed out that one thing we did get is twice as much voltage (or voltage ratings). Just like batteries, when we put capacitors together in series the voltages add up.

Adding capacitors in parallel is like adding resistors in series: the values just add up, no tricks. Why is this? Putting them in parallel effectively increases the size of the plates without increasing the distance between them. More area equals more capacitance. Simple.

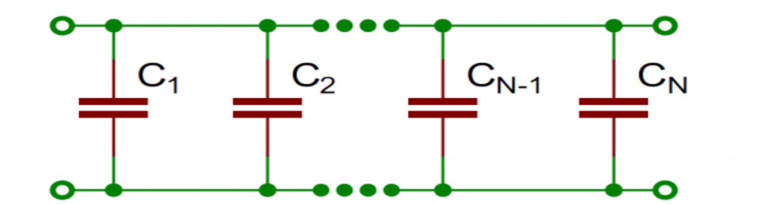

$$C_{Tot} = C_1 + C_2 + \ldots + C_{N-1} + C_N$$

Experiment Time - Part 3

What you'll need:

- One 10kΩ resistor
- Three 100µF caps
- A 3-cell AA battery holder
- Three AA cells
- A breadboard
- A multimeter
- Clip-leads

Let's see some series and parallel connected capacitors in action. This will be a little trickier than the resistor examples, because it's harder to measure capacitance directly with a multimeter.

Let's first talk about what happens when a capacitor charges up from zero volts. When current starts to go in one of the leads, an equal amount of current comes out the other. And if there's no resistance in series with the capacitor, it can be quite a lot of current. In any case, the current flows until the capacitor starts to charge up to the value of the applied voltage, more slowly trickling off until the voltages are equal, when the current flow stops entirely.

As stated above, the current draw can be quite large if there's no resistance in series with the capacitor, and the time to charge can be very short (like milliseconds or less). For this experiment, we want to be able to watch a capacitor charge up, so we're going to use a 10kΩ resistor in series to slow the action down to a point where we can see it easily. But first we need to talk about what an RC time constant is.

$$\tau = R * C$$

What the above equation says is that one time constant in seconds (called tau) is equal to the resistance in ohms times the capacitance in farads. Simple? No? We shall demonstrate on the next page.

Experiment Time - Part 3, Continued...

For the first part of this experiment, we're going to use one 10K resistor and one 100µF (which equals 0.0001 farads). These two parts create a time constant of 1 second:

$$\tau = (10,000\ \Omega)*(0.0001\ F) = 1\ second$$

When charging our 100µF capacitor through a 10kΩ resistor, we can expect the voltage on the cap to rise to about 63% of the supply voltage in 1 time constant, which is 1 second. After 5 time constants (5 seconds in this case) the cap is about 99% charged up to the supply voltage, and it will follow a charge curve something like the plot below.

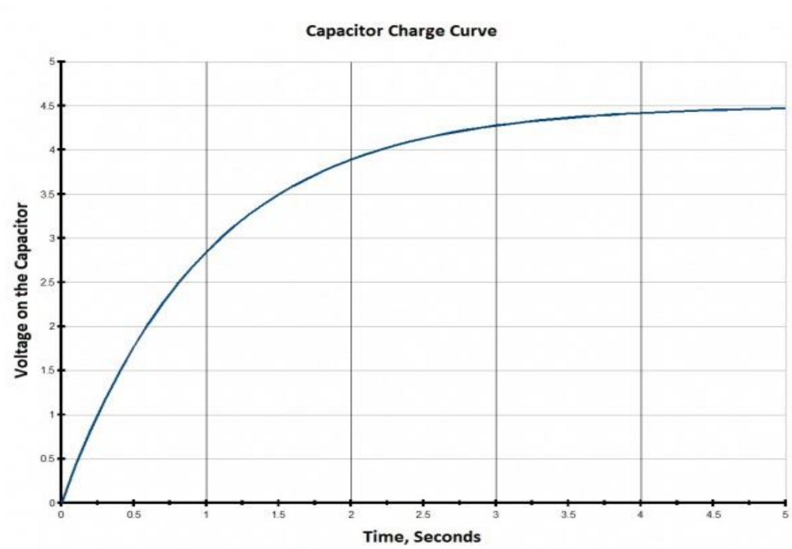

Now that we know that stuff, we're going to connect the circuit in the diagram (make sure to get the polarity right on that capacitor!).

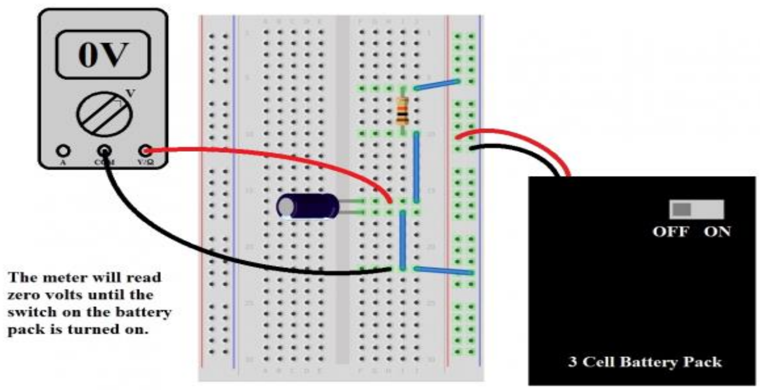

With our multimeter set to measure volts, check the output voltage of the pack with the switch turned on. That’s our supply voltage, and it should be something around 4.5V (it'll be a bit more if the batteries are new). Now connect the circuit, taking care that the switch on the battery pack is in the “OFF” position before plugging it into the breadboard. Also, take care that the red and black leads are going to the right places. If it’s more convenient, you can use alligator clips to attach the meter probes to the legs of the capacitor for measurement (you can also spread those legs out a bit to make it easier).

Once we’re satisfied that the circuit looks right and our meter’s on and set to read volts, flip the switch on the battery pack to “ON”. After about 5 seconds, the meter should read pretty close to the battery pack voltage, which demonstrates that the equation is right and we know what we’re doing. Now turn the switch off. It’s still holding that voltage pretty well, isn't it? That’s because there's no path for current to discharge the capacitor; we've got an open circuit. To discharge the cap, you can use another 10K resistor in parallel. After about 5 seconds, it will be back to pretty close to zero.

Experiment Time - Part 3, Even More...

Now we’re on to the interesting parts, starting with connecting two capacitors in series. Remember that we said the result of which would be similar to connecting two resistors in parallel. If this is true, we can expect (using product-over-sum)

$$C = \frac{100\ \mu F \cdot 100\ \mu F}{100\ \mu F + 100\ \mu F} = 50\ \mu F$$

What's that going to do to our time constant?

$$\tau = (10,000\ \Omega){*}(0.00005\ F) = 0.5\ seconds$$

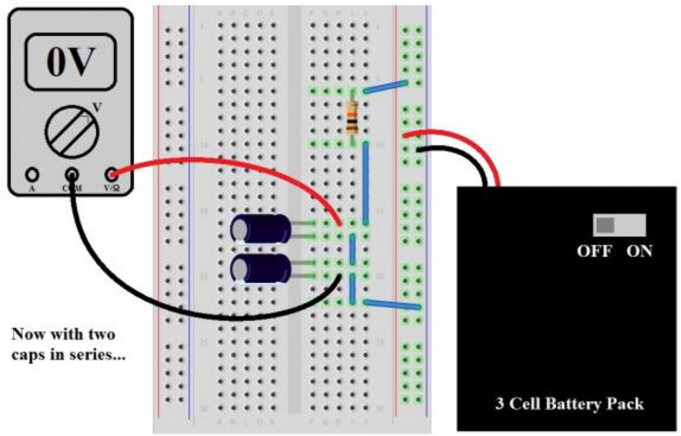

Now with two caps in series...

With that in mind, plug in another capacitor in series with the first, make sure the meter is reading zero volts (or there-abouts) and flip the switch to "ON". Did it take about half as much time to charge up to the battery pack voltage? That's because there's half as much capacitance. The electron gas tank got smaller, so it takes less time to charge it up. A third capacitor is suggested for this experiment just to prove the point, but we're betting the reader can see the writing on the wall.

Now we'll try capacitors in parallel, remembering that we said earlier that this would be like adding resistors in series. If that's true, then we can expect 200µF, right? Then our time constant becomes

$$\tau = (10,000\ \Omega){*}(0.0002\ F) = 2\ second$$

This means that it will now take about 10 seconds to see the parallel capacitors charge up to the supply voltage of 4.5V.

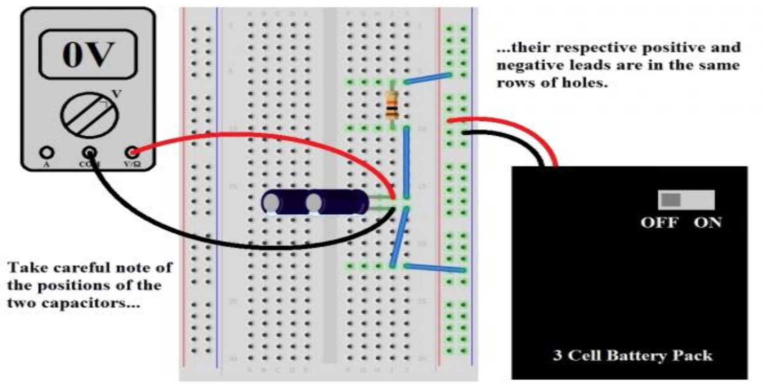

For the proof, start with our original circuit of one 10kΩ resistor and one 100µF capacitor in series, as hooked up in the first diagram for this experiment. We already know that the capacitor is going to charge up in about 5 seconds. Now add a second capacitor in parallel. Make sure the meter is reading close to zero volts (discharge through a resistor if it isn't reading zero), and flip the switch on the battery pack to “ON”. Takes a long time, doesn't it? Sure enough, we made the electron gas tank bigger and now it takes longer to fill it up. To prove it to yourself, try adding the third 100µF capacitor, and watch it charge for a good, long time.

Series and Parallel Inductors

Series and Parallel Inductors

Cases where inductors need to be added either in series or in parallel are rather rare, but not unheard of. In any case, let’s address them just to be complete.

In a nutshell they add just like resistors do, which is to say they add with a plus sign when in series, and with product-over-sum when in parallel. The tricky part comes when they are placed close together so as to have interacting magnetic fields, whether intentionally or not. For this reason, it is preferable to have a single component rather than two or more, though most inductors are shielded to prevent interacting magnetic fields.

In any case, suffice it to say that they add like resistors do. More information than that regarding inductors is well beyond the scope of this tutorial.

Capacitors

Introduction

A capacitor is a two-terminal, electrical component. Along with resistors and inductors, they are one of the most fundamental **passive** components we use. You would have to look very hard to find a circuit which *didn't* have a capacitor in it.

What makes capacitors special is their ability to **store energy**; they're like a fully charged electric battery. *Caps*, as we usually refer to them, have all sorts of critical applications in circuits. Common applications include local energy storage, voltage spike suppression, and complex signal filtering.

Symbols and Units

Circuit Symbols

There are two common ways to draw a capacitor in a schematic. They always have two terminals, which go on to connect to the rest of the circuit. The capacitors symbol consists of two parallel lines, which are either flat or curved; both lines should be parallel to each other, close, but not touching (this is actually representative of how the capacitor is made. Hard to describe, easier to just show:

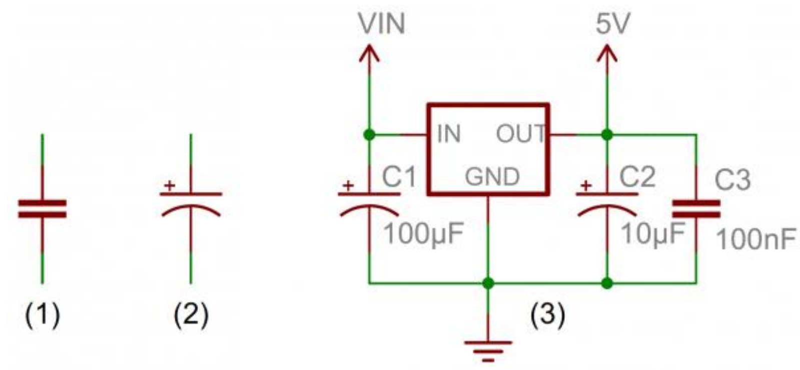

(1) and (2) are standard capacitor circuit symbols. (3) is an example of capacitors symbols in action in a voltage regulator circuit.

The symbol with the curved line (#2 in the photo above) indicates that the capacitor is polarized, meaning it's probably an electrolytic capacitor. More on that in the types of capacitors section of this tutorial.

Each capacitor should be accompanied by a name -- C1, C2, etc.. -- and a value. The value should indicate the capacitance of the capacitor; how many farads it has. Speaking of farads...

Capacitance Units

Not all capacitors are created equal. Each capacitor is built to have a specific amount of capacitance. The capacitance of a capacitor tells you **how much charge it can store**, more capacitance means more capacity to store charge. The standard unit of capacitance is called the **farad**, which is abbreviated *F*.

It turns out that a farad is a *lot* of capacitance, even 0.001F (1 milifarad -- 1mF) is a big capacitor. Usually you'll see capacitors rated in the pico- (10-12) to microfarad (10-6) range.

Prefix Name	Abbreviation	Weight	Equivalent Farads
Picofarad	pF	10^{-12}	0.000000000001 F
Nanofarad	nF	10^{-9}	0.000000001 F
Microfarad	µF	10^{-6}	0.000001 F
Milifarad	mF	10^{-3}	0.001 F
Kilofarad	kF	10^{3}	1000 F

When you get into the farad to kilofarad range of capacitance, you start talking about special caps called *super* or *ultra*-capacitors.

Capacitor Theory

Note: The stuff on this page isn't completely critical for electronics beginners to understand...and it gets a little complicated towards the end. We recommend reading the *How a Capacitor is Made* section, the others could probably be skipped if they give you a headache.

How a Capacitor Is Made

The schematic symbol for a capacitor actually closely resembles how it's made. A capacitor is created out of two metal plates and an insulating material called a **dielectric**. The metal plates are placed very close to each other, in parallel, but the dielectric sits between them to make sure they don't touch.

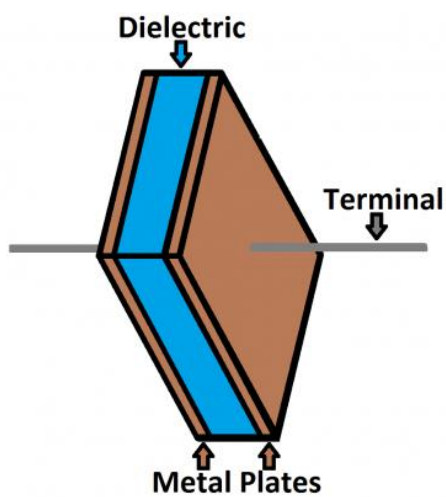

Your standard capacitor sandwich: two metal plates separated by an insulating dielectric.

The dielectric can be made out of all sorts of insulating materials: paper, glass, rubber, ceramic, plastic, or anything that will impede the flow of current.

The plates are made of a conductive material: aluminum, tantalum, silver, or other metals. They're each connected to a terminal wire, which is what eventually connects to the rest of the circuit.

The capacitance of a capacitor -- how many farads it has -- depends on how it's constructed. More capacitance requires a larger capacitor. Plates with more overlapping surface area provide more capacitance, while more distance between the plates means less capacitance. The material of the dielectric even has an effect on how many farads a cap has. The total capacitance of a capacitor can be calculated with the equation:

$$C = \varepsilon r \frac{A}{4\pi d}$$

Where εr is the dielectric's relative permittivity (a constant value determined by the dielectric material), *A* is the amount of area the plates overlap each other, and *d* is the distance between the plates.

How a Capacitor Works

Electric current is the flow of electric charge, which is what electrical components harness to light up, or spin, or do whatever they do. When current flows into a capacitor, the charges get "stuck" on the plates because they can't get past the insulating dielectric. Electrons -- negatively charged particles -- are sucked into one of the plates, and it becomes overall negatively charged. The large mass of negative charges on one plate pushes away like charges on the other plate, making it positively charged.

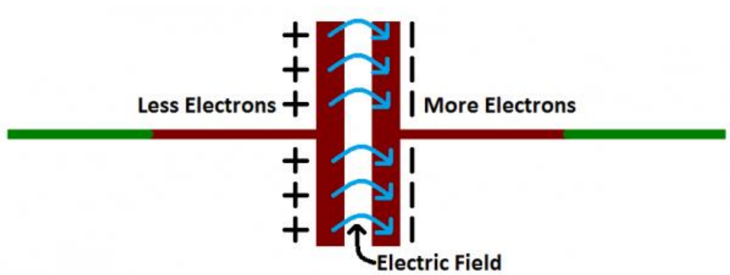

The positive and negative charges on each of these plates attract each other, because that's what opposite charges do. But, with the dielectric sitting between them, as much as they want to come together, the charges will forever be stuck on the plate (until they have somewhere else to go). The stationary charges on these plates create an electric field, which influence electric potential energy and voltage. When charges group together on a capacitor like this, the cap is storing electric energy just as a battery might store chemical energy.

Charging and Discharging

When positive and negative charges coalesce on the capacitor plates, the capacitor becomes **charged**. A capacitor can retain its electric field -- hold its charge -- because the positive and negative charges on each of the plates attract each other but never reach each other.

At some point the capacitor plates will be so full of charges that they just can't accept any more. There are enough negative charges on one plate that they can repel any others that try to join. This is where the **capacitance** (farads) of a capacitor comes into play, which tells you the maximum amount of charge the cap can store.

If a path in the circuit is created, which allows the charges to find another path to each other, they'll leave the capacitor, and it will **discharge**.

For example, in the circuit below, a battery can be used to induce an electric potential across the capacitor. This will cause equal but opposite charges to build up on each of the plates, until they're so full they repel any more current from flowing. An LED placed in series with the cap could provide a path for the current, and the energy stored in the capacitor could be used to briefly illuminate the LED.

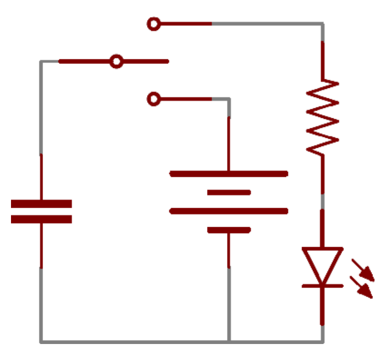

Calculating Charge, Voltage, and Current

A capacitor's capacitance -- how many farads it has -- tells you how much charge it can store. How much charge a capacitor is *currently* storing depends on the potential difference (voltage) between its plates. This relationship between charge, capacitance, and voltage can be modeled with this equation:

$$Q = CV$$

Charge (Q) stored in a capacitor is the product of its capacitance (C) and the voltage (V) applied to it.

The capacitance of a capacitor should always be a constant, known value. So we can adjust voltage to increase or decrease the cap's charge. More voltage means more charge, less voltage...less charge.

That equation also gives us a good way to define the value of one farad. One farad (F) is the capacity to store one unit of energy (coulombs) per every one volt.

Calculating Current

We can take the charge/voltage/capacitance equation a step further to find out how capacitance and voltage affect current, because current is the *rate* of flow of charge. The gist of a capacitor's relationship to voltage and current is this: the amount of **current through a capacitor** depends on both the capacitance and how quickly the **voltage is rising or falling**. If the voltage across a capacitor swiftly rises, a large positive current will be induced through the capacitor. A slower rise in voltage across a capacitor equates to a smaller current through it. If the voltage across a capacitor is steady and unchanging, no current will go through it.

(This is ugly, and gets into calculus. It's not all that necessary until you get into time-domain analysis, filter-design, and other gnarly stuff, so skip ahead to the next page if you're not comfortable with this equation.) The equation for calculating current through a capacitor is:

$$i = C\frac{dv}{dt}$$

The *dV/dt* part of that equation is a derivative (a fancy way of saying *instantaneous rate*) of voltage over time, it's equivalent to saying "how fast is voltage going up or down at this very moment". The big takeaway from this equation is that if **voltage is steady**, the derivative is zero, which means **current is**

also zero. This is why current cannot flow through a capacitor holding a steady, DC voltage.

Types of Capacitors

There are all sorts of capacitor types out there, each with certain features and drawbacks which make it better for some applications than others.

When deciding on capacitor types there are a handful of factors to consider:

- **Size** - Size both in terms of physical volume and capacitance. It's not uncommon for a capacitor to be the largest component in a circuit. They can also be very tiny. More capacitance typically requires a larger capacitor.
- **Maximum voltage** - Each capacitor is rated for a maximum voltage that can be dropped across it. Some capacitors might be rated for 1.5V, others might be rated for 100V. Exceeding the maximum voltage will usually result in destroying the capacitor.
- **Leakage current** - Capacitors aren't perfect. Every cap is prone to leaking some tiny amount of current through the dielectric, from one terminal to the other. This tiny current loss (usually nanoamps or less) is called leakage. Leakage causes energy stored in the capacitor to slowly, but surely drain away.
- **Equivalent series resistance (ESR)** - The terminals of a capacitor aren't 100% conductive, they'll always have a tiny amount of resistance (usually less than 0.01Ω) to them. This resistance becomes a problem when a lot of current runs through the cap, producing heat and power loss.
- **Tolerance** - Capacitors also can't be made to have an exact, precise capacitance. Each cap will be rated for their nominal capacitance, but, depending on the type, the exact value might vary anywhere from ±1% to ±20% of the desired value.

Ceramic Capacitors

The most commonly used and produced capacitor out there is the ceramic capacitor. The name comes from the material from which their dielectric is made.

Ceramic capacitors are usually both physically and capacitance-wise **small**. It's hard to find a ceramic capacitor much larger than 10µF. A surface-mount ceramic cap is commonly found in a tiny 0402 (0.4mm x 0.2mm), 0603 (0.6mm x 0.3mm) or 0805 package. Through-hole ceramic caps usually look like small (commonly yellow or red) bulbs, with two protruding terminals.

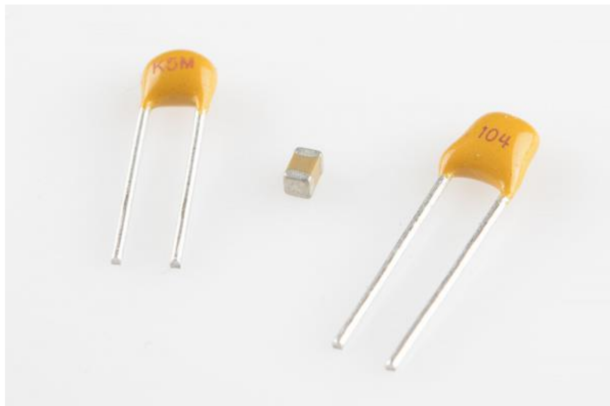

Two caps in a through-hole, radial package; a 22pF cap on the left, and a 0.1µF on the right. In the middle, a tiny 0.1µF 0603 surface-mount cap.

Compared to the equally popular electrolytic caps, ceramics are a more near-ideal capacitor (much lower ESR and leakage currents), but their small capacitance can be limiting. They are usually the least expensive option too. These caps are well-suited for high-frequency coupling and decoupling applications.

Aluminum and Tantalum Electrolytic

Electrolytics are great because they can pack *a lot* of capacitance into a relatively small volume. If you need a capacitor in the range of 1µF-1mF, you're most likely to find it in an electrolytic form. They're especially well suited to high-voltage applications because of their relatively high maximum voltage ratings.

Aluminum electrolytic capacitors, the most popular of the electrolytic family, usually look like little tin cans, with both leads extending from the bottom.

An assortment of through-hole and surface-mount electrolytic capacitors. Notice each has some method for marking the cathode (negative lead).

Unfortunately, electrolytic caps are usually **polarized**. They have a positive pin -- the anode -- and a negative pin called the cathode. When voltage is applied to an electrolytic cap, the anode must be at a higher voltage than the cathode. The cathode of an electrolytic capacitor is usually identified with a '-' marking, and a colored strip on the case. The leg of the anode might also be slightly longer as another indication. If voltage is applied in reverse on an electrolytic cap, they'll fail spectacularly (making a *pop* and bursting open), and permanently. After popping an electrolytic will behave like a short circuit.

These caps also notorious for **leakage** -- allowing small amounts of current (on the order of nA) to run through the dielectric from one terminal to the other. This makes electrolytic caps less-than-ideal for energy storage, which is unfortunate given their high capacity and voltage rating.

Supercapacitors

If you're looking for a capacitor made to store energy, look no further than supercapacitors. These caps are uniquely designed to have *very* high capacitances, in the range of farads.

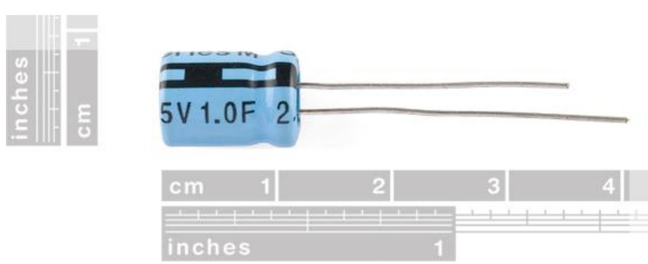

A 1F (!) supercapacitor. High capacitance, but only rated for 2.5V. Notice these are also polarized.

While they can store a huge amount of charge, supercaps can't deal with very high voltages. This 10F supercap is only rated for 2.5V max. Any more than that will destroy it. Super caps are commonly placed in series to achieve a higher voltage rating (while reducing total capacitance).

The main application for supercapacitors is in **storing and releasing energy**, like batteries, which are their main competition. While supercaps can't hold as much energy as an equally sized battery, they can release it much faster, and they usually have a much longer lifespan.

Others

Electrolytic and ceramic caps cover about 80% of the capacitor types out there (and supercaps only about 2%, but they're super!). Another common capacitor type is the **film capacitor**, which features very low parasitic losses (ESR), making them great for dealing with very high currents.

There's plenty of other less common capacitors. **Variable capacitors** can produce a range of capacitances, which makes them a good alternative to variable resistors in tuning circuits. Twisted wires or PCBs can create capacitance (sometimes undesired) because each consists of two conductors separated by an insulator. Leyden Jars -- a glass jar filled with and surrounded by conductors -- are the O.G. of the capacitor family. Finally, of course, flux capacitors (a strange combination of inductor and capacitor) are critical if you ever plan on traveling back to the glory days.

Capacitors in Series/Parallel

Much like resistors, multiple capacitors can be combined in series or parallel to create a combined equivalent capacitance. Capacitors, however, add together in a way that's **completely the opposite** of resistors.

Capacitors in Parallel

When capacitors are placed in parallel with one another the total capacitance is simply the **sum of all capacitances**. This is analogous to the way resistors add when in series.

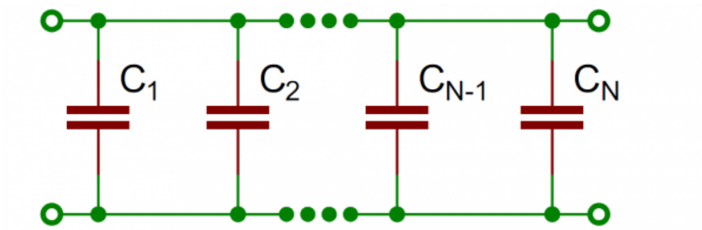

$$C_{Tot} = C_1 + C_2 + \ldots + C_{N-1} + C_N$$

So, for example, if you had three capacitors of values 10µF, 1µF, and 0.1µF in parallel, the total capacitance would be 11.1µF (10+1+0.1).

Capacitors in Series

Much like resistors are a pain to add in parallel, capacitors get funky when placed in *series*. The total capacitance of *N* capacitors in series is the inverse of the sum of all inverse capacitances.

$$\frac{1}{C_{Tot}} = \frac{1}{C_1} + \frac{1}{C_2} + \ldots + \frac{1}{C_{N-1}} + \frac{1}{C_N}$$

If you only have **two** capacitors in series, you can use the "product-over-sum" method to calculate the total capacitance:

$$C_{Tot} = \frac{C_1 C_2}{C_1 + C_2}$$

Taking that equation even further, if you have **two equal-valued capacitors in series**, the total capacitance is half of their value. For example two 10F supercapacitors in series will produce a total capacitance of 5F (it'll also have the benefit of doubling the voltage rating of the total capacitor, from 2.5V to 5V).

Application Examples

There are tons of applications for this nifty little (actually they're usually pretty large) passive component. To give you an idea of their wide range of uses, here are a few examples:

Decoupling (Bypass) Capacitors

A lot of the capacitors you see in circuits, especially those featuring an integrated circuit, are decoupling. A decoupling capacitor's job is to supress high-frequency noise in power supply signals. They take tiny voltage ripples, which could otherwise be harmful to delicate ICs, out of the voltage supply.

In a way, decoupling capacitors act as a very small, local power supply for ICs (almost like an uninterruptible power supply is to computers). If the power supply very temporarily drops its voltage (which is actually pretty common, especially when the circuit it's powering is constantly switching its load requirements), a decoupling capacitor can briefly supply power at the correct voltage. This is why these capacitors are also called **bypass** caps; they can temporarily act as a power source, *bypassing* the power supply.

Decoupling capacitors connect between the power source (5V, 3.3V, etc.) and ground. It's not uncommon to use two or more different-valued, even different types of capacitors to bypass the power supply, because some capacitor values will be better than others at filtering out certain frequencies of noise.

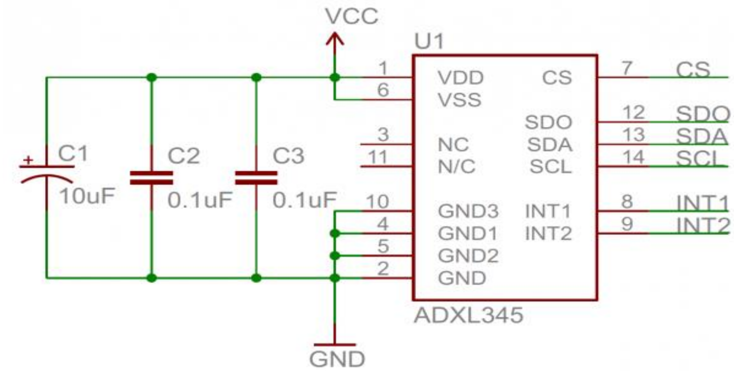

In [this schematic](), three decoupling capacitors are used to help reduce the noise in an accelerometer's voltage supply. Two ceramic 0.1µF and one tantalum electrolytic 10µF split decoupling duties.

While it seems like this might create a short from power to ground, only high-frequency signals can run through the capacitor to ground. The DC signal will go to the IC, just as desired. Another reason these are called bypass capacitors is because the high frequencies (in the kHz-MHz range) bypass the IC, instead running through the capacitor to get to ground.

When physically placing decoupling capacitors, they should always be located as close as possible to an IC. The further away they are, they less effective they'll be.

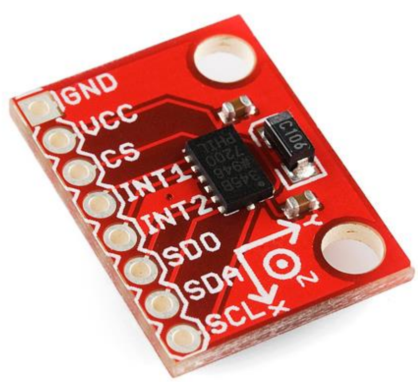

Here's the physical circuit layout from the schematic above. The tiny, black IC is surrounded by two 0.1µF capacitors (the brown caps) and one 10µF electrolytic tantalum capacitor (the tall, black/grey rectangular cap).

To follow good engineering practice, always add at least one decoupling capacitor to every IC. Usually 0.1µF is a good choice, or even add some 1µF or 10µF caps. They're a cheap addition, and they help make sure the chip isn't subjected to big dips or spikes in voltage.

Power Supply Filtering

Diode rectifiers can be used to turn the AC voltage coming out of your wall into the DC voltage required by most electronics. But diodes alone can't turn an AC signal into a clean DC signal, they need the help of capacitors! By adding a **parallel capacitor** to a bridge rectifier, a rectified signal like this:

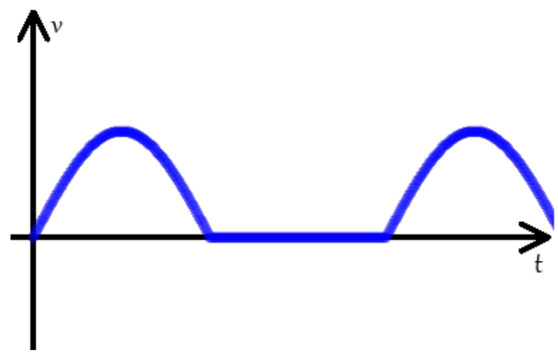

Can be turned into a near-level DC signal like this:

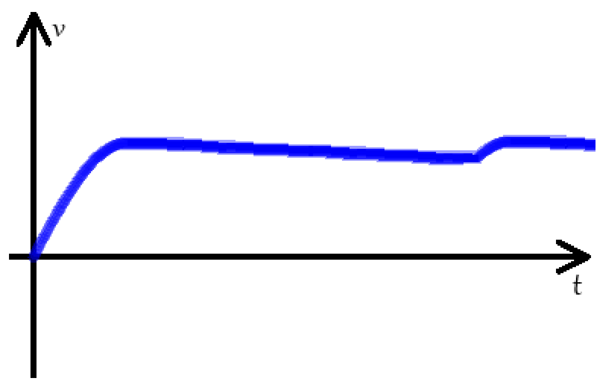

Capacitors are stubborn components, they'll always try to resist sudden changes in voltage. The filter capacitor will charge up as the rectified voltage increases. When the rectified voltage coming into the cap starts its rapid decline, the capacitor will access its bank of stored energy, and it'll discharge very slowly, supplying energy to the load. The capacitor shouldn't fully discharge before the input rectified signal starts to increase again, recharging the cap. This dance plays out many times a second, over-and-over as long as the power supply is in use.

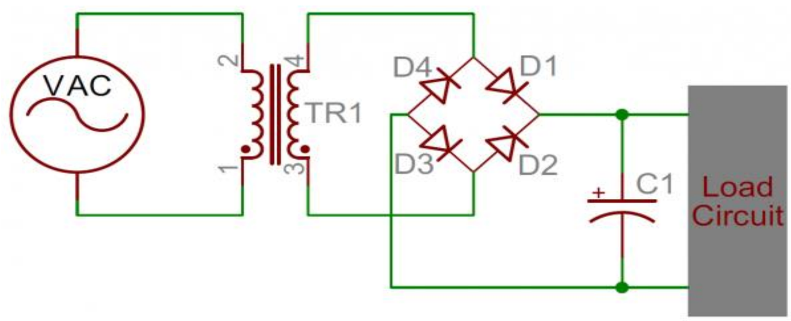

An AC-to-DC power supply circuit. The filter cap (C1) is critical in smoothing out the DC signal sent to the load circuit.

If you tear apart any AC-to-DC power supply, you're bound to find at least one rather large capacitor. Below are the guts of a 9V DC wall adapter. Notice any capacitors in there?

There might be more capacitors than you think! There are four electrolytic, tin-can-looking caps ranging from 47µF to 1000µF. The big, yellow rectangle in the foreground is a high-voltage 0.1µF polypropylene film cap. The blue disc-shaped cap and the little green one in the middle are both ceramics.

Energy Storage and Supply

It seems obvious that if a capacitor stores energy, one of it's many applications would be supplying that energy to a circuit, just like a battery. The problem is capacitors have a much lower **energy density** than batteries; they just can't pack as much energy as an equally sized chemical battery (but that gap is narrowing!).

The upside of capacitors is they usually lead longer lives than batteries, which makes them a better choice environmentally. They're also capable of delivering energy much faster than a battery, which makes them good for applications which need a short, but high burst of power. A camera flash might get its power from a capacitor (which, in turn, was probably charged by a battery).

Battery or Capacitor?

	Battery	Capacitor
Capacity	✓	
Energy Density	✓	
Charge/Discharge Rate		✓
Life Span		✓

Signal Filtering

Capacitors have a unique response to signals of varying frequencies. They can block out low-frequency or DC signal-components while allowing higher frequencies to pass right through. They're like a bouncer at a very exclusive club for high frequencies only.

Filtering signals can be useful in all sorts of signal processing applications. Radio receivers might use a capacitor (among other components) to tune out undesired frequencies.

Another example of capacitor signal filtering is passive **crossover** circuits inside speakers, which separate a single audio signal into many. A series capacitor will block out low frequencies, so the remaining high-frequency parts of the signal can go to the speaker's tweeter. In the low-frequency passing, subwoofer circuit, high-frequencies can mostly be shunted to ground through the parallel capacitor.

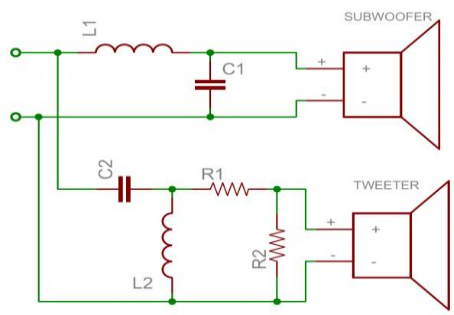

A very simple example of an audio crossover circuit. The capacitor will block out low frequencies, while the inductor blocks out high frequencies. Each can be used to deliver the proper signal to tuned audio drivers.