Modular Electronics Learning (ModEL) project

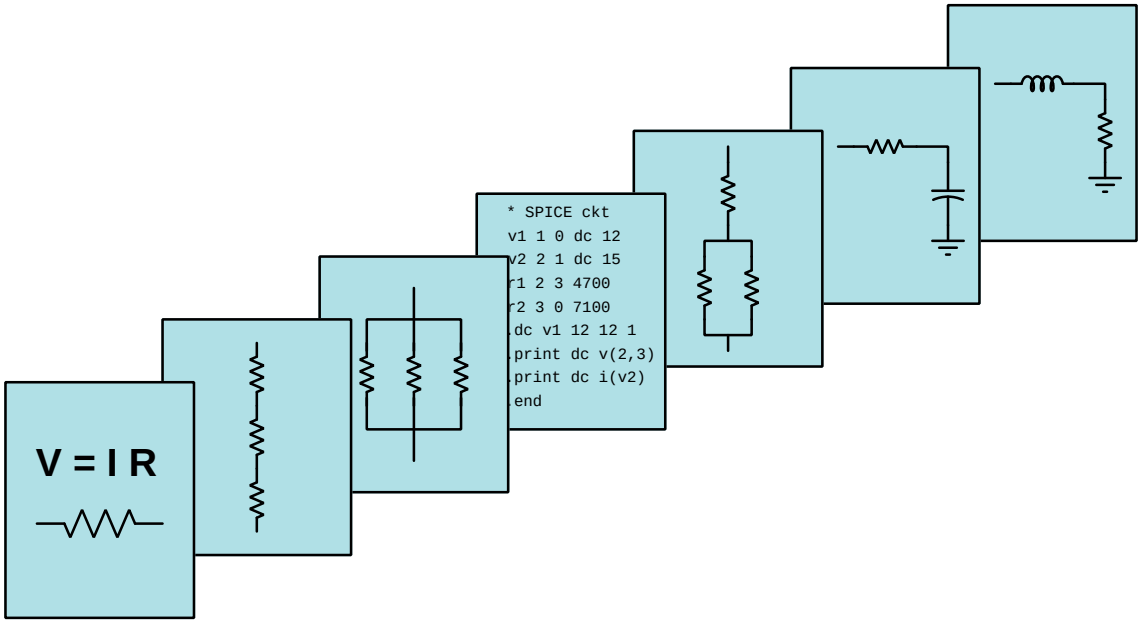

DC-DC Power Conversion

© 2019-2026 by Tony R. Kuphaldt – under the terms and conditions of the Creative Commons Attribution 4.0 International Public License

Last update = 11 October 2026

This is a copyrighted work, but licensed under the Creative Commons Attribution 4.0 International Public License. A copy of this license is found in the last Appendix of this document. Alternatively, you may visit `http://creativecommons.org/licenses/by/4.0/` or send a letter to Creative Commons: 171 Second Street, Suite 300, San Francisco, California, 94105, USA. The terms and conditions of this license allow for free copying, distribution, and/or modification of all licensed works by the general public.

Contents

1 Introduction **3**
- 1.1 Recommendations for students . . . 3
- 1.2 Challenging concepts related to DC-DC power conversion . . . 5
- 1.3 Recommendations for instructors . . . 6

2 Case Tutorial **7**
- 2.1 Example: simple buck converter . . . 8
- 2.2 Example: simple boost converter . . . 11
- 2.3 Example: measuring signal rates of change . . . 15

3 Tutorial **19**
- 3.1 Matching source and load voltages . . . 20
- 3.2 Capacitor and inductor review . . . 23
- 3.3 BJT and FET review . . . 26
- 3.4 Simple step-down converter . . . 31
- 3.5 Buck converter . . . 35
- 3.6 Boost converter . . . 39
- 3.7 Inverting converter . . . 41
- 3.8 Ćuk converter . . . 43
- 3.9 Flyback converter . . . 47
- 3.10 Forward converter . . . 49
- 3.11 Continuous and discontinuous converter operation . . . 51
- 3.12 Charge pumps . . . 57
- 3.13 Hybrid switching-linear converters . . . 62
- 3.14 Integrated circuit power converters . . . 65

4 Historical References **67**
- 4.1 Battery substations . . . 68
- 4.2 Ćuk converter patent . . . 75

5 Derivations and Technical References **79**
- 5.1 Capacitor parameters and dielectric types . . . 80
- 5.2 Inductor parameters and core types . . . 84
- 5.3 Capacitor parasitics . . . 88

5.3.1 Model of a real capacitor . . . 88
5.3.2 Parasitic resistance in capacitors . . . 89
5.3.3 Parasitic inductance in capacitors . . . 89
5.3.4 Other parasitic effects in capacitors . . . 90
5.4 Inductor parasitics . . . 91
5.4.1 Model of a real inductor . . . 91
5.4.2 Parasitic resistance in inductors . . . 92
5.4.3 Parasitic capacitance in inductors . . . 93
5.4.4 Other parasitic effects in inductors . . . 93
5.4.5 Inductor testing . . . 94

6 Animations **95**
6.1 VFD transistor switching sequence . . . 96
6.2 Animation of an NPN BJT turning on . . . 123
6.3 Animation of an E-type MOSFET turning on . . . 138

7 Questions **151**
7.1 Conceptual reasoning . . . 155
7.1.1 Reading outline and reflections . . . 156
7.1.2 Foundational concepts . . . 157
7.1.3 Buck converter regulation . . . 160
7.1.4 24 Volt spotlight . . . 161
7.1.5 Linear vs. switch-mode . . . 161
7.1.6 Poor converter efficiency . . . 162
7.1.7 Source/load functions in Ćuk converter . . . 163
7.1.8 Transistor biasing . . . 164
7.2 Quantitative reasoning . . . 165
7.2.1 Miscellaneous physical constants . . . 166
7.2.2 Introduction to spreadsheets . . . 167
7.2.3 Converter efficiency calculations . . . 170
7.2.4 Solving for duty cycle . . . 171
7.2.5 Dying battery source . . . 172
7.2.6 Instantaneous analyses of DC-DC converter circuits . . . 173
7.3 Diagnostic reasoning . . . 175
7.3.1 Faults in a buck converter . . . 176
7.3.2 Faults in a boost converter . . . 177
7.3.3 Identifying possible root-cause faults . . . 178

A Problem-Solving Strategies **179**

B Instructional philosophy **181**
B.1 First principles of learning . . . 182
B.2 Proven strategies for instructors . . . 183
B.3 Proven strategies for students . . . 185
B.4 Design of these learning modules . . . 186

C Tools used **189**

D Creative Commons License **193**

E References **201**

F Version history **203**

Index **205**

Chapter 1

Introduction

1.1 Recommendations for students

A *DC-DC converter* circuit does for DC what a transformer does for AC, namely allow the transformation of voltage and current to new levels with very little power dissipation. Several different types of converter circuits exist, each one with its own advantages and disadvantages.

Important concepts related to DC-DC power conversion include **sources** versus **loads**, **transformers**, the **Conservation of Energy**, energy **efficiency**, the **Second Law of Thermodynamics**, **Joule's Law**, **capacitance**, **inductance**, **transistors**, **pulse-width modulation**, voltage **regulation**, **buck** versus **boost**, **high-side** versus **low-side** switching, **rates of change**, **duty cycle**, **mutual inductance**, electrical **isolation**, **flying capacitors**, **charge pumps**, **Faraday's Law of Electromagnetic Induction**, magnetic **saturation**, and **ground loops**.

A useful problem-solving strategy employed in the Tutorial is that of *limiting cases*, where we imagine the function of a system as some variable achieves extreme values. This works to simplify the problem, and is usually a good method for discerning the *direction* that a variable's influence will take in a system.

Here are some good questions to ask of yourself while studying this subject:

- How might an experiment be designed and conducted to measure the duty cycle of a pulse waveform? What hypothesis (i.e. prediction) might you pose for that experiment, and what result(s) would either support or disprove that hypothesis?
- How might an experiment be designed and conducted to measure the energy conversion efficiency of a DC-to-DC converter circuit? What hypothesis (i.e. prediction) might you pose for that experiment, and what result(s) would either support or disprove that hypothesis?
- What is a disadvantage of linear voltage regulation?
- Why can't transformers be used with direct current (DC) power?

- Typically, how efficient are power transformers at converting energy?
- How exactly does capacitance store energy?
- How exactly does inductance store energy?
- What parameter is most stable – voltage or current – where there exists capacitance, and why?
- What parameter is most stable – voltage or current – where there exists inductance, and why?
- What are some of the complementary properties of capacitance versus inductance?
- What role does the *duty cycle* of pulsing play in a power converter circuit?
- What relevance does Fourier's Theorem have in switch-mode power conversion circuits?
- Why are switch-mode power supply circuits more energy-efficient than linear regulator circuits?
- What role does the diode play in any switch-mode converter design?
- How may voltage regulation be added to a switch-mode converter circuit?
- How do *synchronous* and *asynchronous* converter designs differ, and what are their relative advantages and disadvantages?
- How does mutual inductance differ from self-inductance?
- Where might electrical isolation be an advantageous feature in a power converter circuit?
- What are some of the advantages of the Ćuk converter over the standard inverting converter circuit design?
- What is the purpose of the third coil in the "forward" converter design?
- Which components are acting as sources and which as loads in Figure 5 of the Ćuk patent shown in the Tutorial?
- How do charge pumps differ in design and operation from inductor/capacitor switch-mode converter circuits?
- What is a *flying capacitor* and what it is useful for?
- Why are "forward" converter circuits more appropriate when used with certain types of power sources than others?
- What are the respective advantages and disadvantages of the different types of switch-mode power converter circuits?

1.2 Challenging concepts related to DC-DC power conversion

The following list cites concepts related to this module's topic that are easily misunderstood, along with suggestions for properly understanding them:

- **Necessary conditions for BJT operation** – bipolar junction transistors require a base current in the proper direction to inject minority carriers into the base layer (i.e. PNP transistors require hole injection from emitter to base, and so the emitter must be made positive and the base made negative, relative to each other; NPN transistors require electron injection from emitter to base, and so the emitter must be negative and the base positive.), and the collector-to-emitter voltage must be of the correct polarity (i.e. so that collector terminal current goes in the same direction as base terminal current, either *out* or *in*). Both currents join at the emitter terminal, making the emitter current the sum of the base and collector currents.

- **Necessary conditions for FET operation** – field effect transistors work on the principle of a transverse electric field affecting the conductivity of a homogenously-doped channel. For insulated-gate FETs (i.e. MOSFETs), the electric field exists across a thin layer of dielectric between the gate terminal and the modulated channel. When that applied gate-channel voltage is of such a polarity as to attract majority charge-carriers into the channel (e.g. electrons into an N-type channel, holes into a P-type channel) then the channel becomes more conductive.

- **Creative application of foundational concepts** – many students enter college-level study of electronics with an educational background stressing rote memorization at the expense of logical reasoning from trusted principles, and as such tend to find circuit analysis daunting where there is no single procedure or single formula always yielding the correct answer(s). The educational solution to this is for the instructor to *constantly* tie all new circuit concepts back to previously-learned principles and show the logical relationships between them, and then with every new circuit-analysis example challenge students to reason from those concepts on their own. When students construct a mental "toolbox" they may apply in a plurality of ways to any new challenge they encounter, they will be as free as any skilled mechanic who knows how and why each of their hand tools functions to freely select and use whichever of those tools is most appropriate at any given time. A word of caution for new instructors: *be very sparing in your direction on how to use those tools, instead engaging with students Socratically when they are "stuck" and cannot see a path to a solution!*

1.3 Recommendations for instructors

This section lists realistic student learning outcomes supported by the content of the module as well as suggested means of assessing (measuring) student learning. The outcomes state what learners should be able to do, and the assessments are specific challenges to prove students have learned.

- **Outcome** – Demonstrate effective technical reading and writing

  Assessment – Students present their outlines of this module's instructional chapters (e.g. Case Tutorial, Tutorial, Historical References, etc.) ideally as an entry to a larger Journal document chronicling their learning. These outlines should exhibit good-faith effort at summarizing major concepts explained in the text.

  Assessment – Students show how quantitative results were obtained by the author in the Tutorial chapter's examples.

  Assessment – Students present their "Conceptual Toolbox" Theory Journal entries listing important concepts applied in the Tutorial text(s). A "Conceptual Toolbox" serves as inspiration when solving problems, each concept suggesting a potential problem-solving action. Challenge students to collaboratively generate a Conceptual Toolbox for the day's topic during the class session and then use it when solving problems together, thus fostering the habit of reasoning from principles rather than recalling procedures!

- **Outcome** – Apply the concept of energy efficiency to DC-DC power conversion circuits

  Assessment – Calculate either input voltage, input current, or efficiency given the other input and output parameters; e.g. pose problems in the form of the "Converter efficiency calculations" Quantitative Reasoning question.

- **Outcome** – Apply the foundational concepts to the solution of instantaneous voltages and currents in a DC-DC converter circuit

  Assessment – e.g. pose problems in the form of the "Instantaneous analyses of DC-DC converter circuits" Quantitative Reasoning question.

- **Outcome** – Independent research

  Assessment – Locate DC-DC converter datasheets and properly interpret some of the information contained in those documents including maximum continuous power output, maximum peak power output, DC voltage input, efficiency, converter topology, etc.

Chapter 2

Case Tutorial

The idea behind a *Case Tutorial* is to explore new concepts by way of example. In this chapter you will read less presentation of theory compared to other Tutorial chapters, but by close observation and comparison of the given examples be able to discern patterns and principles much the same way as a scientific experimenter. Hopefully you will find these cases illuminating, and a good supplement to text-based tutorials.

These examples also serve well as challenges following your reading of the other Tutorial(s) in this module – can you explain *why* the circuits behave as they do?

2.1 Example: simple buck converter

Switch-mode power conversion circuits typically operate at fairly high frequency so as to allow small-value reactive components. This conserves space on a printed circuit board, as well as weight and cost for the design. However, high frequencies can be problematic for student-constructed "breadboard" circuits where conductor length is poorly controlled and parasitic inductance and capacitance values tend to run high, and also where student-grade test equipment (especially multimeters!) may not have enough bandwidth to perform reliable AC measurements. The following circuit is designed to work at relatively low frequency (below 10 kHz) and use commonly-available components (e.g. all diodes are 1N4001):

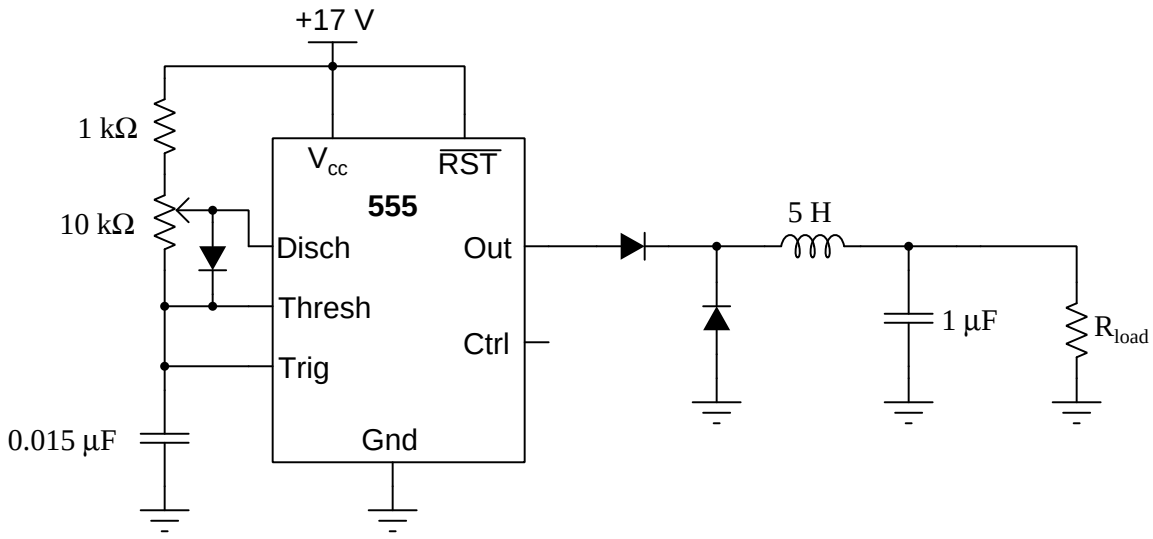

A suitable inductor for this converter circuit is the primary winding of a small step-down power transformer, such as the type commonly used to energize residential doorbell circuits. Such transformers are inexpensive, readily available at hardware stores, and with multiple windings a single transformer offers a range of inductance values for experimentation. In this particular case I used a Heath-Zenith model EM57583H transformer with primary winding rated at 120 Volts AC and secondary windings rated at 8 Volts, 16 Volts, and 24 Volts. The primary winding measured approximately 5 Henrys, and the 24 Volt secondary winding approximately 220 milliHenrys. Such an inductor provides a relatively high amount of inductance (necessary to effectively convert at low switching frequencies) and a high current rating and saturation limit so that the inductor will behave well when passing substantial DC current.

The 555 timer generates a variable duty cycle square wave output, and instead of using that signal to drive a transistor to handle the power switching, the 555's output itself is the switched power source[1].

[1] It is worth noting that the 555 timer is equipped with robust output transistors, capable of sourcing and sinking in excess of 200 milliAmperes!

Proof that the buck converter is indeed functioning consists of measuring the load voltage to see that it is a direct function of pulse duty cycle, and also that output (load) current exceeds input current (measured at the Out terminal of the 555).

Waveforms measured at the load and at the 555 timer's Out terminal show the converter outputting a DC voltage (yellow trace) proportional to duty cycle (violet trace):

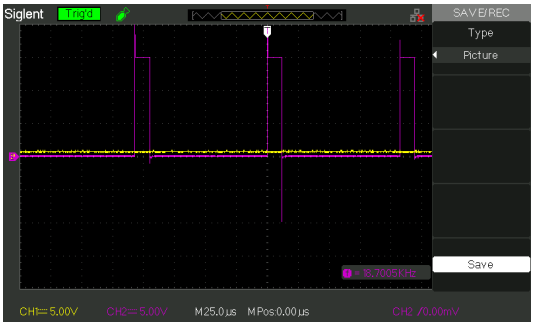

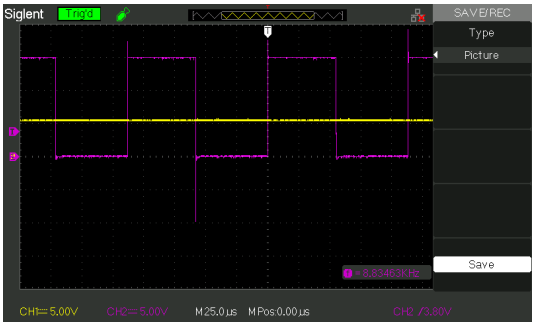

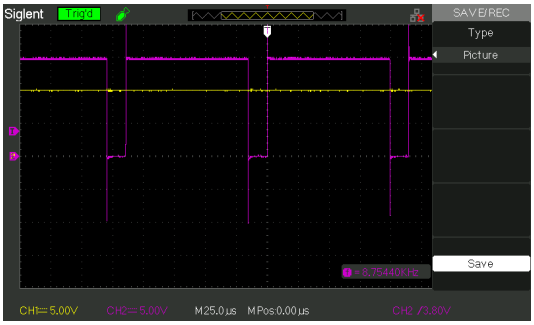

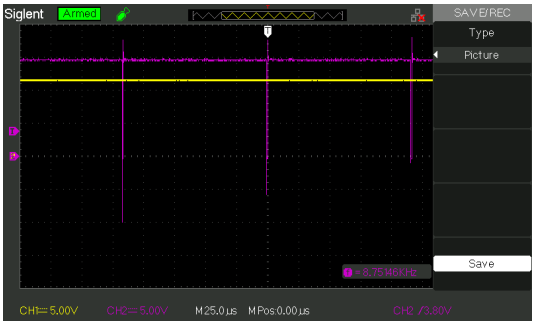

If we modify the 555 timer circuit to output a much lower frequency, we may see the output voltage rise and fall significantly, and this gives us some insight into the reactive components' absorption and release of energy:

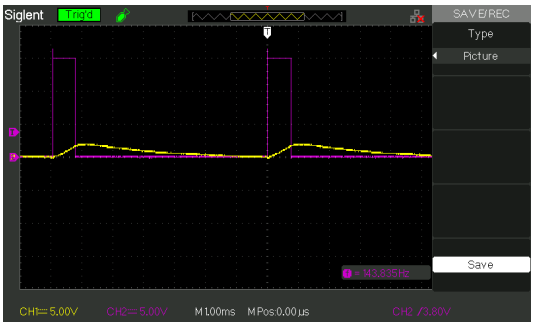

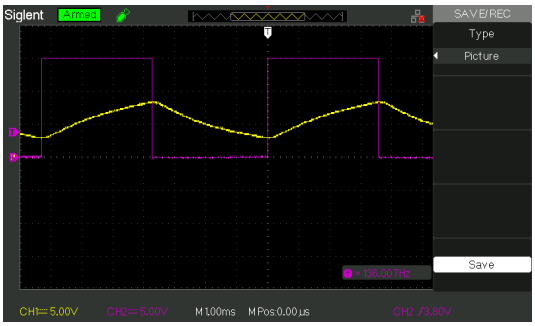

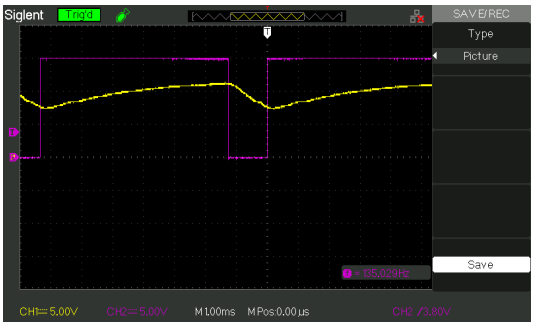

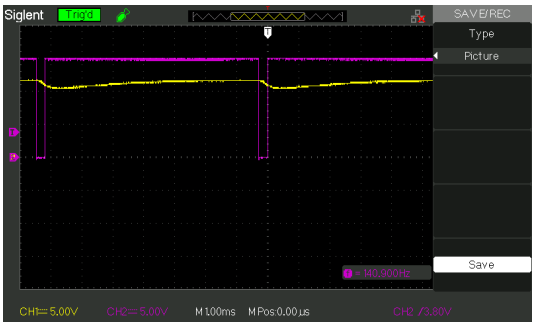

An interesting advantage of using a *transformer* for the inductor is that the unused windings provide a means of connecting an oscilloscope to monitor what is happening in the inductor's magnetic core, since by Faraday's Law of Electromagnetic Induction $V = N\frac{d\Phi}{dt}$.

2.2 Example: simple boost converter

The following is a simple "boost" form of switch-mode power conversion circuit:

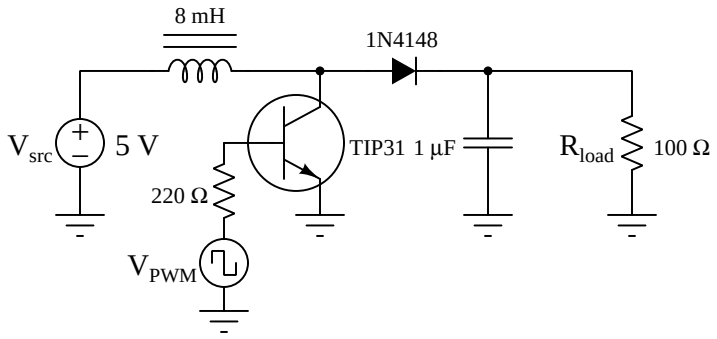

With a much smaller inductance value than the previous (buck) converter design, this circuit must be operated at much higher PWM frequency in order to provide a relatively smooth output voltage to the load. Note that the choice to use an 8 mH inductor rather than a 5 H inductor has nothing whatsoever to do with the distinction between boost and buck conversion – it's just arbitrary. What is definitely *not* arbitrary about our choice of inductors are parameters such as DC resistance, maximum rated current, and saturation current. We need an inductor that is able to handle the expected amount of current without dropping excessive voltage due to its own wire resistance ($V = IR$), overheat the coil due to its own resistance ($P = I^2R$), or fail to store enough energy because it saturates during its energization period.

The following oscilloscope screenshots show the circuit operating at a PWM switching frequency of 40 kHz, with different duty cycle values resulting in 7 Volts, 9 Volts, and 10 Volts, respectively, output to the 100 Ω load. The red trace is measured at the output of the PWM signal source while the blue trace senses load voltage:

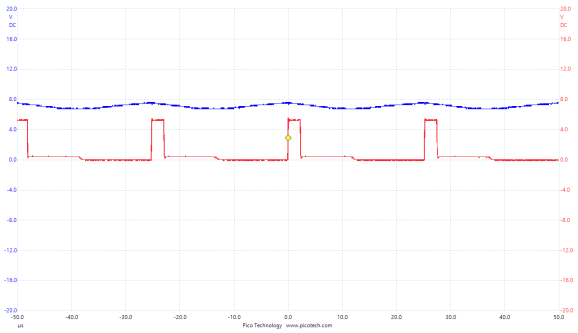

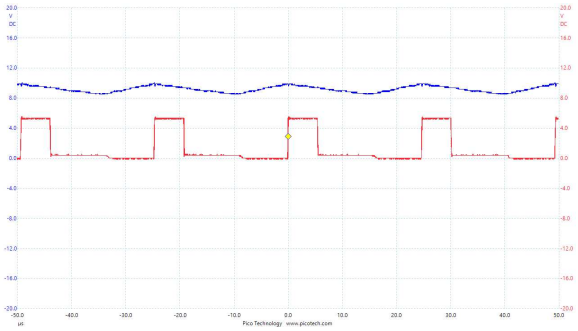

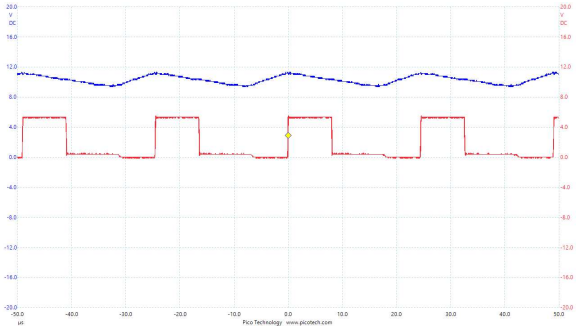

If we slow the PWM switching frequency down to 7.?? kHz from 40 kHz, we see much greater voltage ripple at the output (load). Such ripple could be minimized at this lower frequency by using higher-valued reactive components:

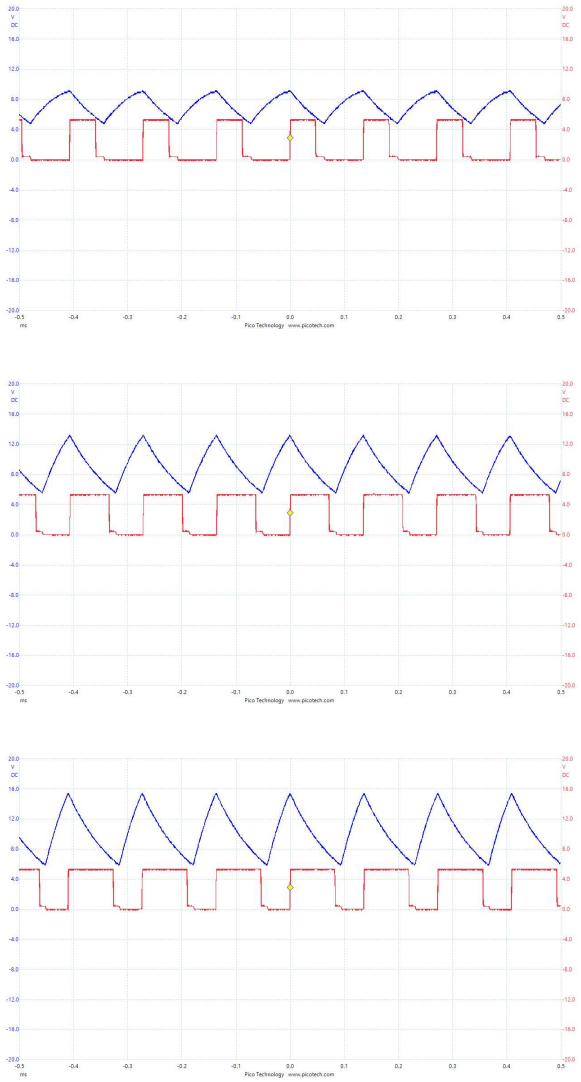

Here we see the same circuit operating with a PWM frequency that is much too low for the components:

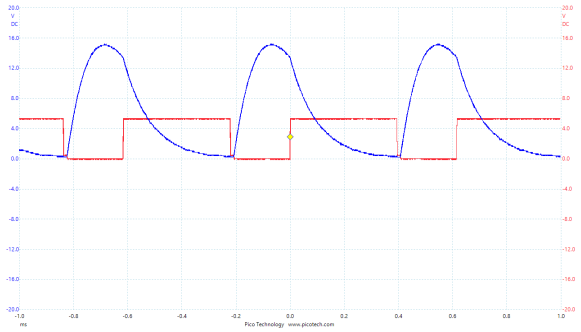

The capacitor loses nearly all of its charge by the end of each PWM cycle, resulting in an enormous peak-to-peak variation[2] at the load.

The fact that we obtain an output voltage greater than our DC supply voltage (5 Volts) proves this converter circuit is indeed "boosting" voltage. More evidence of its voltage/current "transforming" action would come from measurements of supply versus output current: with greater load voltage than supply voltage, we would expect to see greater supply current than load current.

[2] Is it even fair to call it "ripple" now?

2.3 Example: measuring signal rates of change

Capacitors and inductors relate voltage to current by rates of change. For a capacitor, the amount of current is proportional to how quickly voltage across that capacitor either rises or falls over time ($I = C\frac{dV}{dt}$). For an inductor, the amount of voltage is proportional to how quickly current through that inductor rises or falls over time ($V = L\frac{dI}{dt}$).

For example, a 330 microFarad capacitor experiencing a voltage increasing at a rate of 45 Volts per second will pass 14.85 milliAmperes. If the voltage happens to decrease at an equivalent rate (i.e. $\frac{dV}{dt} = -45$ Volts per second) then the 14.85 milliAmpere current will reverse direction through the capacitor compared to how it flowed with the increasing voltage.

Similarly, a 100 milliHenry inductor experiencing a current increasing at a rate of 5 Amperes per second will induce a voltage of 500 milliVolts. If the current happens to decrease by an equivalent rate (i.e. $\frac{dI}{dt} = -5$ Amperes per second) then the 500 milliVolt voltage induced across the inductor will reverse polarity from what it was during the period of increasing current.

Not only are rates-of-change important in determining how energy-storing devices such as capacitors and inductors will respond in circuits, but rates-of-change are also important for determining how *parasitic* capacitances and inductances will affect intended circuit behavior. Parasitic capacitance exists between any two conducting surfaces separated by an electrically insulating medium, and parasitic inductance exists along any length of conductor. This means any rate-of-change of voltage over time between two separated conductors will cause some amount of current to "pass"[3] between them, and that any rate-of-change of current over time through any single conductor will cause some amount of voltage to drop across its length. In many circuits these parasitic effects are negligible, but in circuits experiencing extremely fast rates of change for voltage and/or current the effects can be significant or even severe.

Oscilloscopes are ideal for performing empirical measurements of voltage rates-of-change, and of current rates-of-change given the proper accessories[4]. Some skill is required to do this, though, and here we will explore practical examples to show how it is done.

For any signal plotted in the time domain, where the horizontal axis of the plot is expressed in units of seconds, milliseconds, microseconds, etc., the signal's rate of change at any given point will be the *slope* or *pitch* of the waveform, mathematically defined as its *rise over run*. A great aid to discerning slope at any location on a waveform is to sketch a straight line visually matching the wave's slope at that point, then use locations along that straight line to more easily discern how far it rises (or falls) over some "run" of time. We call this straight line a *tangent line*.

[3] The Scottish physicist James Clerk Maxwell referred to this current as *displacement current*.

[4] For example, a *current probe* converting a sensed current into a voltage the oscilloscope may directly sense, or a *shunt resistor* placed in the circuit developing an oscilloscope-measurable voltage drop for any current passing through.

Here we see an example of a waveform with sloping sections. In the first image we see a specific location on the waveform where we wish to measure voltage rate-of-change ($\frac{dV}{dt}$):

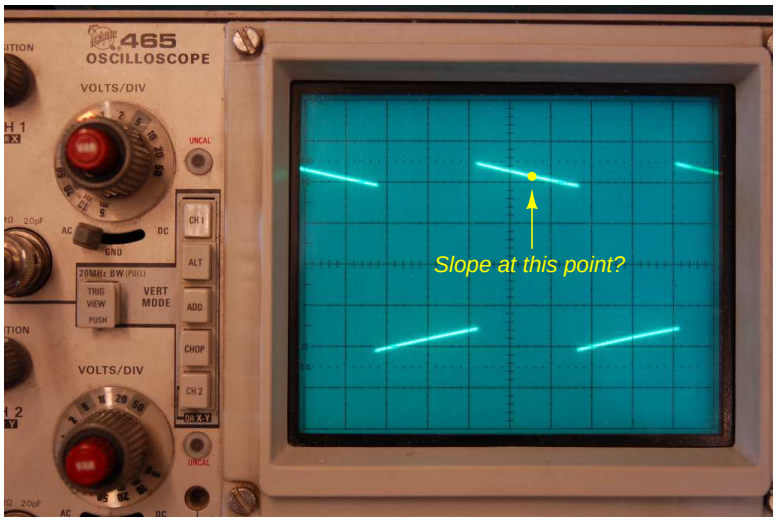

Next we see a tangent line drawn to match the slope of the waveform at the specified location, with convenient points for fall/run measurements taken on that line against the oscilloscope grid's major divisions. In this example, the oscilloscope's vertical sensitivity has been set for 0.5 Volts per division, and the horizontal timebase for 0.2 milliseconds per division:

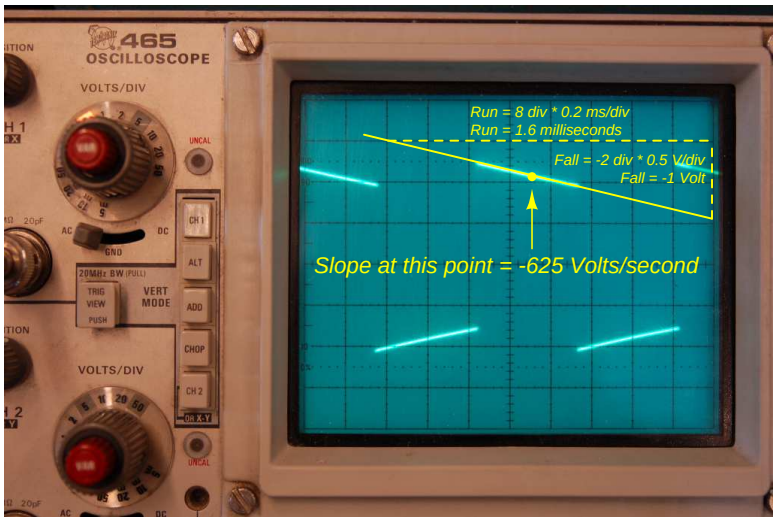

As we can see, the tangent line falls 2 vertical divisions (-1 Volt) over a timespan of 8 horizontal divisions (1.6 milliseconds), yielding a $\frac{dV}{dt}$ quotient of -625 Volts per second, which may also be expressed as -0.625 Volts per millisecond. The negative sign is important, as it distinguishes this particular rate-of-change as *falling* rather than rising over time.

Such rate-of-change measurements are necessarily approximate, as they require us to visually gauge where a tangent line may be overlaid on the waveform's oscillograph, and also to visually assess the slope of that tangent line using the grid provided on the instrument's display screen. However, in most applications extremely precise rate-of-changes are not necessary, and such techniques suffice quite well.

Below is another example of a waveform with sloping sections, the oscilloscope configured for 500 milliVolts per division on the vertical axis and 5 milliseconds per division on the horizontal:

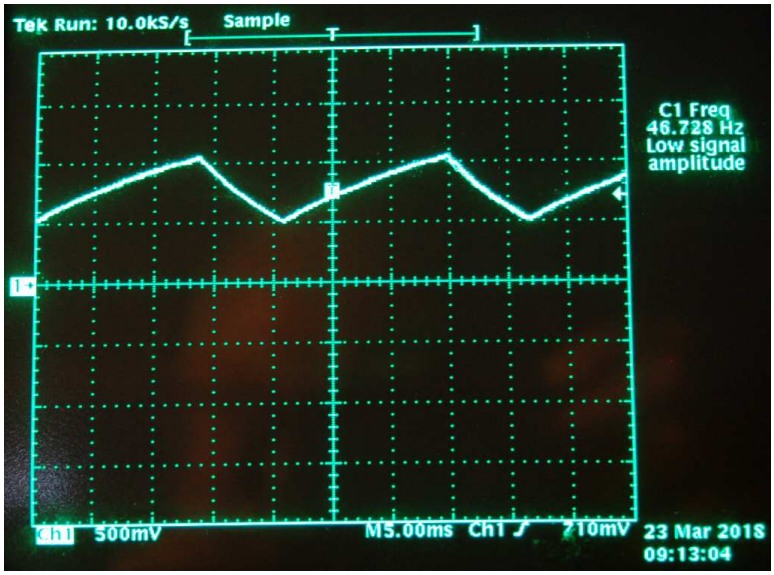

Approximating the slope for each rising section of this wave, we count one division of rise over 2.8 divisions of run, or 500 milliVolts rise over 14 milliseconds of run. This is a rate-of-rise of +35.7 Volts per second.

Approximating the slope for each falling section of this wave, we count one division of fall over 1.4 divisions of run, or −500 milliVolts fall over 7 milliseconds of run. This is a rate-of-fall of −71.4 Volts per second.

If these rates-of-change appear suspiciously large compared to the actual amplitude of the waveform, which barely crests over +1 Volt on the oscillograph, bear in mind that we are calculating *rates of change* for voltage and not absolute values of voltage itself. This is analogous to the distinction between *speed* and *distance*: traveling at a rate of 30 kilometers per hour does not necessarily mean you will travel 30 kilometers, as the actual distance traveled depends on how long that speed is sustained. A voltage rising at a rate of 35.7 Volts per second would indeed rise 35.7 Volts *if given a full second to do so*, but since each rising/falling portion of this waveform is so short in duration the actual amount of rise or fall in each case is only one-half of one Volt. It is therefore perfectly appropriate to consider any $\frac{dV}{dt}$ value as being the *speed* at which a voltage increases or decreases over time, distinct from the actual value of that voltage at any particular moment in time.

Here we see another oscillograph, this one zoomed into the rising edge of a square wave. For this measurement the oscilloscope was configured for 20 milliVolts per division of vertical sensitivity and a timebase of 250 nanoseconds per division on the horizontal:

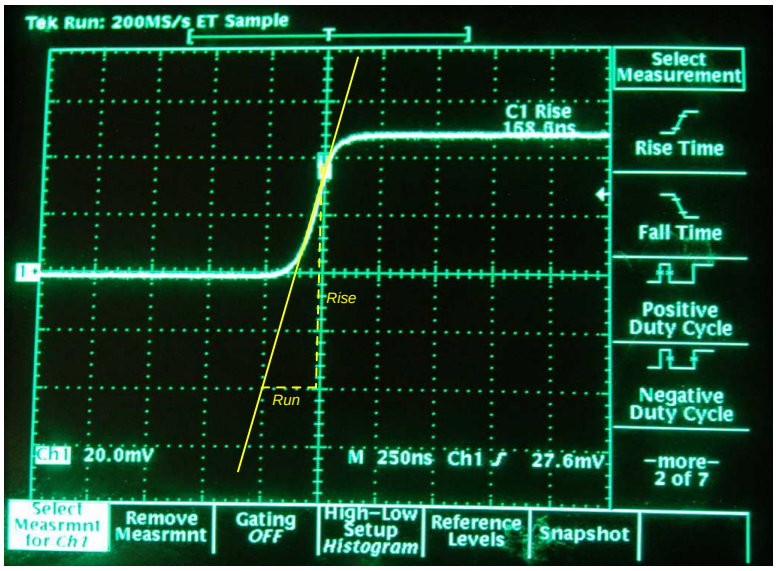

The tangent line overlaid on this screenshot for the purpose of measuring the pulse edge's rate-of-change rises approximately 4 vertical divisions over a run of 1 division, which is 80 milliVolts of rise over 250 nanoseconds of run. The pulse edge's rate-of-change, therefore, is approximately +320,000 Volts per second, or +320 Volts per millisecond, or +0.32 Volts per microsecond (all equivalent expressions of $\frac{dV}{dt}$).

Chapter 3

Tutorial

3.1 Matching source and load voltages

A common problem in electrical and electronic circuit design is properly matching energy sources to loads. Many loads function correctly only when energized within a limited range of voltage or current. Likewise, many electrical sources are limited in the voltage and current levels they are able to output. When the available source is not a good match for the desired load, we must interpose some circuitry between the two to shift the source's output parameters.

In direct current (DC) circuits when source voltage exceeds the desired load voltage, an easy solution is to connect a *dropping resistor* between the source and load, this resistor sized such that it drops the excess source voltage to leave the load with precisely enough:

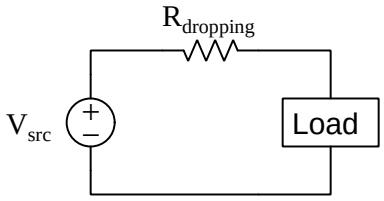

If the load requires a stable supply voltage over a range of load current values and/or a range of source voltage values, we may build an "active" version of the dropping resistor circuit using a *transistor* to drop the excess source voltage, the transistor being driven by a constant-voltage reference (here, a Zener diode and resistor network). This is the basis for most *linear voltage regulator* circuits:

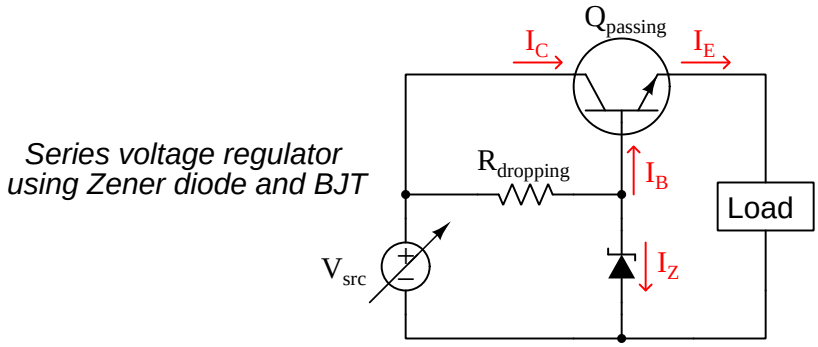

Two major limitations plague these voltage-reducing circuits, though. The obvious limitation is that these networks can only make load voltage something less than the source voltage, and not vice-versa. The second limitation is that they are – by design – inefficient due to energy "wasted" in the form of heat through the use of dissipative voltage-dropping components. Energy efficiency is defined as the ratio of usable energy out versus total energy in ($\frac{E_{out}}{E_{in}}$), and for any system operating in steady-state conditions energy efficiency may be calculated as the ratio of usable *power* out versus total *power* in ($\frac{P_{out}}{P_{in}}$). Since the current into and out of most linear regulators is nearly the same, their energy efficiency is approximately equal to the ratio of load voltage to source voltage ($\approx \frac{V_{load}}{V_{src}}$).

Alternating current (AC) circuits have long benefited from the action of *transformers* to step down or step up source voltage as necessary to match load voltage requirements. Not only is a transformer able to step voltage in either direction (up or down), but it does so with little wasted energy. As voltage steps up, current steps down by approximately the same proportion; as voltage steps down, current steps up. These proportions are a function of the number of "turns" in each of the transformer's windings (coils) because Faraday's Law of Electromagnetic Induction sets voltage proportional to turns for any given change in magnetic flux ($V = N\frac{d\Phi}{dt}$):

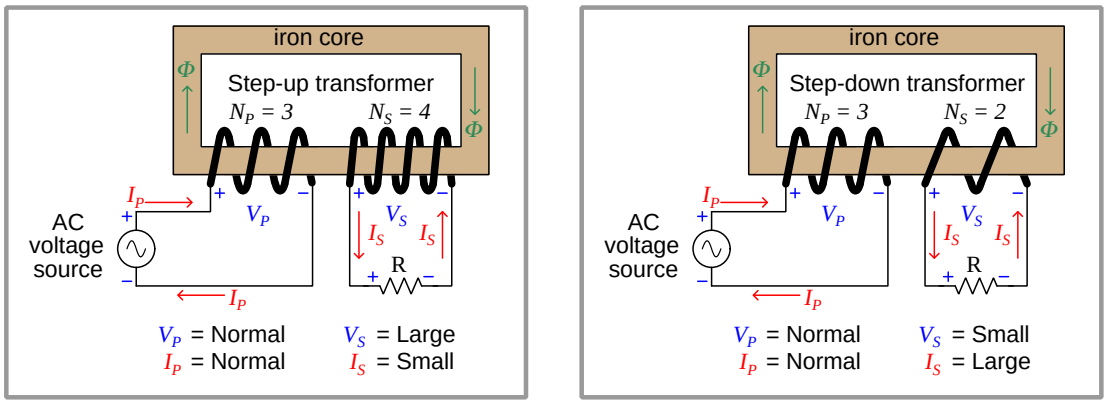

Voltage and current "step" ratios always complement one another due to the *Law of Energy Conservation*. In order for energy to be conserved in a circuit operating in steady-state conditions, $P_{out} = P_{in}$. If we assume negligible dissipative "losses" (i.e. perfect energy efficiency) then P_{out} must equal P_{in}. By Joule's Law we can say that $I_{out}V_{out} = I_{in}V_{in}$, and by algebraic manipulation we may show the complementary ratios of voltage and current: $\frac{I_{out}}{I_{in}} = \frac{V_{in}}{V_{out}}$.

This, more than any other reason, is what positioned AC over DC as the dominant technology for long-distance electric power transmission. Transformers efficiently convert electrical energy into high-voltage/low-current form to avoid I^2R power dissipation over transmission line wires, and then efficiently convert to low-voltage for safe end-use[1].

Nineteenth-century electrical technology offered no competitive DC equivalent to the AC transformer, and so AC reigned supreme in early electrical power system design. However, with the advent of solid-state devices capable of rapidly switching on and off with no mechanical wear, we now have the ability to "transform" DC in a similar manner. Furthermore, these switch-based power conversion circuits offer adjustable voltage/current ratios, unlike transformer ratios which are fixed by the number of turns in the primary and secondary windings. Such circuits are called *DC-DC converters*, *switching converters*, or *switch-mode power supplies* (*SMPS*).

[1] From the Second Law of Thermodynamics we know that every instance of energy conversion results in some of the energy being diverted into less-useful forms (e.g. heat), but some energy-conversion processes are better than others in this regard. Modern AC power transformers and DC-DC converter circuits routinely exhibit energy efficiencies well over 90%. For example, the United States' Department of Energy (DoE) released upgraded transformer efficiency levels in 2016, with the lowest prescribed efficiency being 97.89% for 15 kVA transformers and going upward to 99.28% efficiency for 1000 kVA (three-phase dry-type AC power transformers, all).

A comparison of two DC voltage regulator circuits used to condition the energy from a 15-Volt DC source to a load requiring 3.3 Volts (at 120 milliAmperes) appears below, the upper voltage regulator being linear and the lower voltage regulator being switch-mode in design:

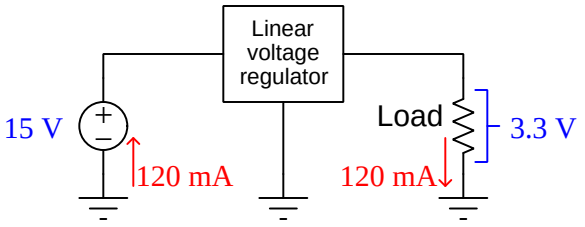

Efficiency = (0.396/1.8) = 22%

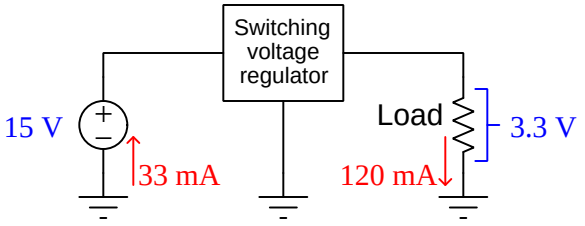

Efficiency = (0.396/0.495) = 80%

Linear voltage regulators essentially behave as series-connected resistances, dropping excess voltage from source to load. This is why their energy efficiency ratings are so low. Switch-mode regulator circuits, by contrast, *transform* DC voltage and current such that a decrease in voltage is accompanied by a boost in current, much like the action of transformers in AC circuits.

Switching power converters operate on the principle of electrical energy storage, typically in capacitors and/or inductors. These energy-storing components are alternately connected to the source and load by semiconductor switches in such a way that they either boost (increase) or buck (decrease) voltage received from the source and deliver that transformed voltage to the load. Transistors and diodes comprise the switching elements, both of which dissipate little energy when operated as on/off (discrete) switches. Capacitors and inductors generally have little parasitic resistance to dissipate energy. Together, they are able to convert between different levels of voltage and current with very high efficiency, which means when used to step voltage in a particular direction they also step current in the other direction by nearly the same proportion, just like an AC transformer, and for the exact same reason (the Conservation of Energy).

3.2 Capacitor and inductor review

Inductance and *capacitance* refer to modes of electrical energy storage: inductance being the storage of energy in a magnetic field associated with current passing through a conductor, and capacitance being the storage of energy in an electric field associated with voltage existing between two distinct conductors. These phenomena take purposeful form in the construction of *inductors* and *capacitors*, respectively. Inductors typically consist of coils of wire wrapped around a magnetic core material, while capacitors consist of metal surfaces[2] separated by an electrical insulator called a *dielectric*.

A small inductor (left) and small capacitor (right) appear in the photographs below:

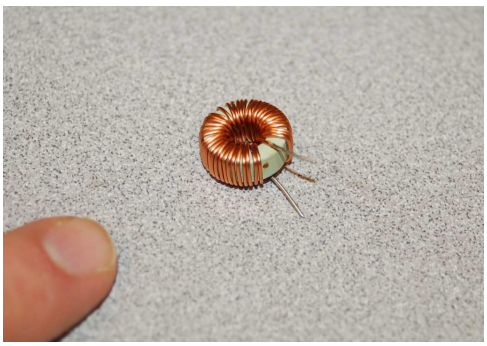

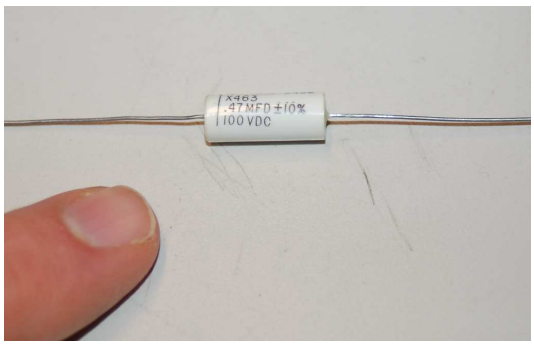

The amount of energy stored in any inductor's magnetic field is a function of current through the inductor's coil, since the strength of that magnetic field varies directly with current. Similarly, the amount of energy stored in any capacitor's electric field is a function of voltage between the capacitor's plates, since the strength of that electric field varies directly with voltage. Energy, as we know, is always conserved, which means it cannot simply appear or vanish without being transferred *from* or *to* somewhere. It follows logically, then, that neither the amount of current through an inductor nor the amount of voltage across a capacitor may vary without energy transfer from or to some other portion of an electric circuit.

[2] Typically the "plates" of a capacitor are sheets of thin metal foil wrapped into cylinders to conserve space.

Energy transfer to or from any electrical component requires that component act as either a *source* or a *load*, and these roles are defined by the relationship between voltage polarity and current direction for that component.

Inductors store energy in magnetic fields, and these fields relate directly to current. Therefore, in order for an inductor to absorb energy from the circuit it must function as a load (i.e. electric charge carriers entering at higher energy and leaving at lower energy, from + to −), and in order to release energy it must function as a source (i.e. charge carriers entering at lower energy and leaving at higher energy, from − to +). Zero voltage across an ideal[3] inductor means a condition of stasis where the magnetic field (and therefore stored energy) neither accumulates nor depletes:

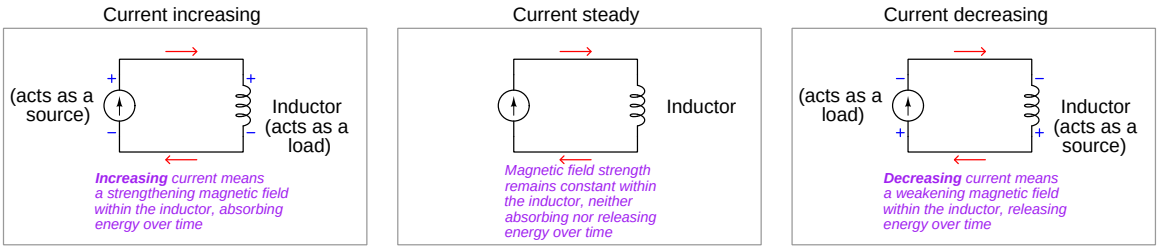

Capacitors store energy in electric fields, and these fields relate directly to voltage. Therefore, in order for a capacitor to absorb energy from the circuit it must function as a load (i.e. electric charge carriers entering at higher energy and leaving at lower energy, from + to −), and in order to release energy it must function as a source (i.e. charge carriers entering at lower energy and leaving at higher energy, from − to +). Zero current through an ideal[4] capacitor means a condition of stasis where the electric field (and therefore stored energy) neither accumulates nor depletes:

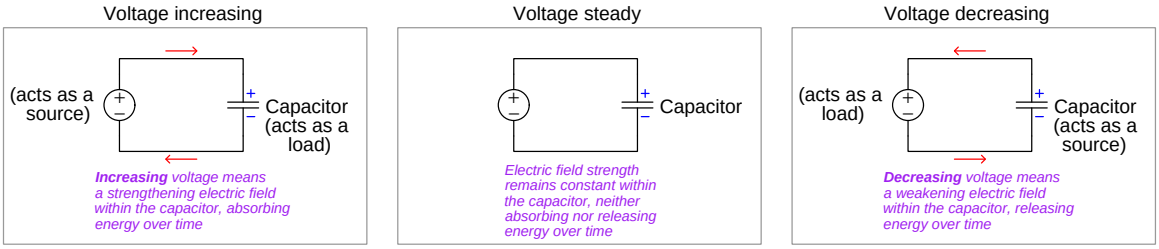

Note which parameters reverse and which maintain the same orientation for each of these energy-storing components. For inductors, current remains flowing in the same direction and voltage polarities reverse as the inductor switches roles as source versus load; for capacitors, voltage remains in the same polarity and current reverses direction as the capacitor switches roles as source versus load. *In each case, the parameter directly related to the amount of energy stored in the component is the one conserved – precisely because energy must be conserved – while the other parameter reverses as needed to facilitate energy transfer into or out of the component.*

[3] An ideal inductor being made of superconducting wire having no resistance to create its own voltage drop.

[4] An ideal capacitor being one with no "leakage" resistance through its dielectric to pass current.

When we compare and contrast the properties of inductance versus capacitance, we see a consistent symmetry between the two, and from this we may see how inductance and capacitance are *complementary properties* of electric circuits:

Capacitance (C)	**Inductance** (L)
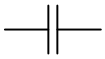 NEMA symbol IEC symbol	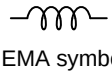 NEMA symbol 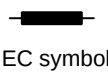 IEC symbol
An *electric field* is created by a *voltage* existing between two conductors	A *magnetic field* is created by a *current* passing through a conductor
Capacitance is enhanced by the area of the conductors, and a narrow spacing between them	Inductance is enhanced by the length of the conductor, and forming a coil shape with multiple turns
An electric field stores energy as a function of voltage: $E_C = \frac{1}{2}CV^2$	A magnetic field stores energy as a function of current: $E_L = \frac{1}{2}LI^2$
Acts as a *load* when *voltage increases* over time ; current flows into + terminal and out of − terminal	Acts as a *load* when *current increases* over time ; voltage drop + where current enters and − where current exits
Maintains a constant voltage when passively storing energy	Maintains a constant current when passively storing energy
Acts as a *source* when *voltage decreases* over time ; current flows out of + terminal and into − terminal	Acts as a *source* when *current decreases* over time ; voltage rise − where current enters and + where current exits
Current is able to *instantly* cease or reverse direction while voltage cannot; current switches direction when capacitor transitions between source and load	Voltage is able to *instantly* cease or reverse polarity while current cannot; voltage switches polarity when inductor transitions between source and load
The amount of current is proportional to the *rate* (speed) at which energy is being absorbed by or released from the electric field	The amount of voltage is proportional to the *rate* (speed) at which energy is being absorbed by or released from the magnetic field
Current is equal to the capacitance times the *rate-of-change* (speed) that voltage varies over time: $I = C\frac{dV}{dt}$	Voltage is equal to the inductance times the *rate-of-change* (speed) that current varies over time: $V = L\frac{dI}{dt}$

3.3 BJT and FET review

DC-to-DC power converter circuits rely on the rapid switching of current through different components, and for this task employ transistors as solid-state switching elements. Both bipolar junction transistors (BJTs) and field-effect transistors (FETs) see application in power conversion circuit designs. This section reviews the basic operational principles of both transistor types.

Transistors are semiconductor devices allowing one electrical signal to control another. Specifically, transistors work by allowing a relatively weak electrical signal to control the electrical conductivity of a solid material, and this modulated conductivity permits or denies passage of a much stronger electrical signal. This is the basic principle of *amplification*: where a smaller signal exerts control over a larger signal. Two major types of transistors dominate modern electronic circuits, *bipolar junction transistors* (BJTs) and *field-effect transistors* (FETs). Each has its own unique principle of operation[5], and although these principles differ markedly, both types of transistors perform similar functions within similar circuits.

Bipolar junction transistors function on the basis of *charge carrier injection*, where a weak electrical signal injects charge carriers into a region of semiconducting material previously lacking plentiful carriers. This injection of foreign charge carriers turns a formerly non-conductive material into a conductive material, allowing the transistor to pass a stronger (i.e. amplified) signal.

Field-effect transistors function on the basis of *channel width modulation*, where the effective width of an electrically conductive channel (i.e. pathway) varies with the application of an electric field from a weak voltage signal. This channel may be "pinched off" or "turned on" by this controlling voltage signal, allowing the transistor to conduct a stronger (i.e. amplified) signal.

[5] Most functional characteristics of transistors may be deduced from a proper understanding of these operating principles. Important behaviors such as the transistor's default state (i.e. whether or not it conducts current in the absence of a "driving" signal), as well as the necessary drive signal types to turn the transistor on (or off) also logically follow from these principles. When first learning about transistors, try to follow the logic connecting each transistor's operating principles with its functional characteristics. Avoid making the common mistake of merely committing facts to memory, devoid of understanding.

Bipolar junction transistors consist of a three-layer "sandwich" of differently-doped semiconductor materials, denoted either NPN or PNP. These three layers form two dissimilar "PN" junctions, hence the name *bipolar*. The middle layer (called the *base*) acts as a controlled-conductivity bridge between the other two layers (called the *emitter* and *collector*) conducting the controlled current. When unstimulated, base-emitter and base-collector depletion regions together prohibit current between collector and emitter. However, if an external source forward-biases the base-emitter junction to collapse that depletion region, charge carriers will be "emitted" by the emitter into the extremely thin base layer where most of them coast through the other side of the base and become "injected" into the depletion region of the base-collector junction, thus making the transistor conductive from end to end. Bipolar junction transistors are *minority-carrier* devices because the majority-type charge carriers emitted by the emitter layer are "minority" type once they enter the base (e.g. holes emitted from a P-type emitter, being injected in the N-type base).

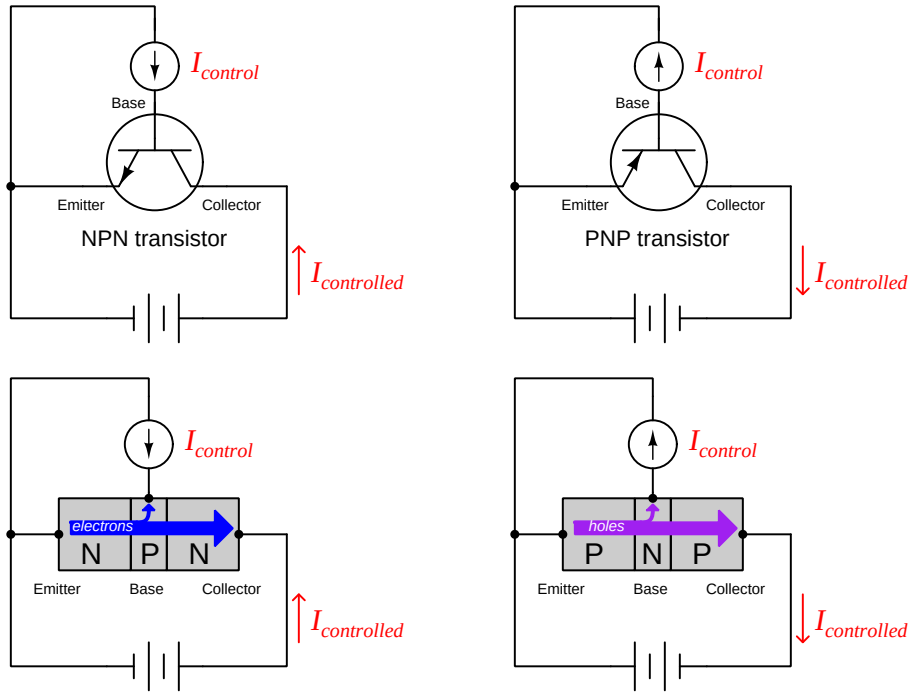

The key to turning on a BJT is to forward-bias the base-emitter junction in order to initiate this charge-injection process. Without this injection of charge carriers into the depleted base-collector region, the transistor defaults to its "off" state because there is no continuous region of charge carriers naturally stretching from emitter to collector.

Field-effect transistors consist of a channel formed of one type of semiconductor material (i.e. either P-type or N-type) that is either flanked or interrupted by a region of the opposite doping type. An electric field applied transverse to this channel affects the quantities and types of charge carriers residing within the channel. Like bipolar junction transistors which exist in NPN and PNP forms, field-effect transistors may be manufactured with either N-type channels or P-type channels. However, unlike BJTs, there are several distinct design variations for FETs: *junction* FETs and two different forms of *insulated-gate* FETs (commonly called *MOSFETs* for their metal-oxide-semiconductor construction). For the sake of simplicity, we will show an *enhancement-mode* MOSFET in the following illustration because it behaves most similarly to a BJT.

Instead of emitter, base, and collector terminals, FETs have *source*, *gate*, and *drain* terminals. The controlled current flows between drain and source, while the gate terminal provides a point to apply the controlling voltage signal. Field-effect transistors are *majority-carrier* devices because the charge carriers leaving the source to pass through the channel (e.g. holes through a P-type channel, electrons through an N-type channel) are native to that channel's doping, never traversing a region of opposite doping:

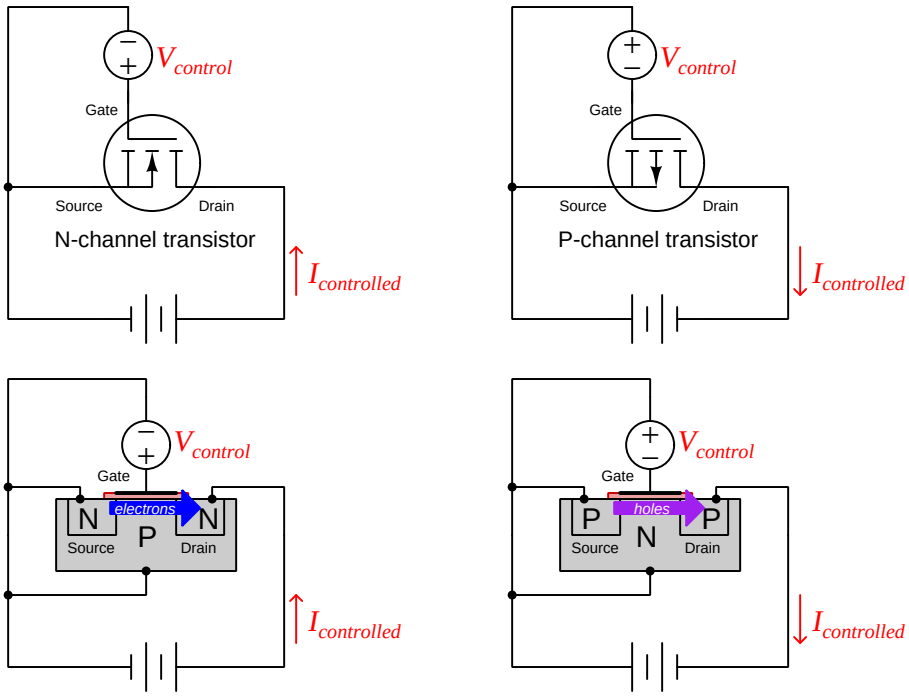

The key to turning on a FET is to electrostatically attract more majority carriers into the channel by giving the gate terminal a polarity opposite the channel type. Reversing this polarity forces the FET to turn off. Different forms of FET (JFET, depletion-mode, enhancement-mode) possess different natural states. JFETs are normally-on devices which may only be turned off (i.e. depleted). D-type MOSFETs are normally-half-on devices which may be turned on more (i.e. enhanced) or turned off more (i.e. depleted). E-type MOSFETs – the type shown above – are normally-off devices which may only be turned on (i.e. enhanced) similar to the behavior of a BJT.

Although the MOSFET is fundamentally a four-terminal device, they are often constructed with the substrate and source terminals connected by wire, which means the controlling voltage must be impressed between the gate and source terminals:

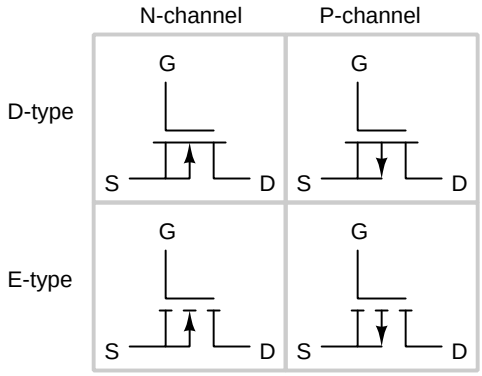

An interesting consequence of this internal connection is that it results in a single PN diode-like parasitic junction connected between the source and drain terminals:

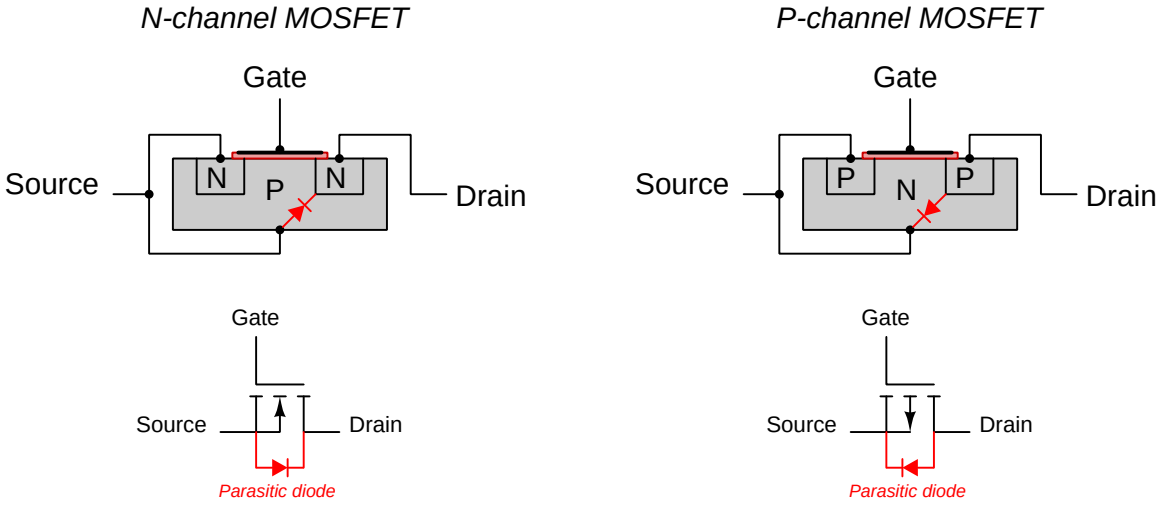

When the source and drain terminals are properly biased[6], this parasitic diode will be reverse-biased and have no effect on the transistor's operation. However, if source-drain polarity ever reverses, this diode will become forward-biased and short past the transistor's channel!

[6]The arrow inside of a MOSFET symbol actually shows you which way this parasitic diode points! For any PN junction the tip of the arrowhead points toward the "N" side and the wide end of the arrowhead is on the "P" side. Note how the wide end of the N-channel MOSFET's arrow connects to the source terminal, as does the anode of the parasitic diode; likewise the tip of the P-channel MOSFET's arrowhead connects to the source terminal, as does the cathode of the parasitic diode. Your task in deciding how to connect any MOSFET into a circuit is to keep this parasitic diode reverse-biased so that the transistor will not be bypassed by it.

Field-effect transistor (FET) and bipolar junction transistor (BJT) behavior differs in several important ways:

Characteristic	**FET**	**BJT**
Charge carrier type	Majority	Minority
Principle of operation	Effective width of a continuous channel modulated by voltage	Charge carriers injected into a depleted region
Drive signal	Voltage only (no DC current)	Voltage and current both necessary
Switching time limit	RC time constant of gate capacitance and drive resistance	Stored charges in base waiting to be swept out
Temperature effect	Greater temperature increases channel resistance	Greater temperature decreases effective resistance
Resting status	Normally-on and normally-off types	Normally-off (only)
Gain	Transconductance $g = \frac{\Delta I_{DS}}{\Delta V_{GS}}$	Beta ratio $\beta = \frac{I_C}{I_B}$
Current regulation	Depletion region "pinch" effect intensifies with increasing V_{DS}	Rate of charge injection into base sets a maximum limit on I_C

Generally speaking, FETs are capable of switching on and off faster[7] than BJTs, and FETs generally offer less resistance[8] to the controlled current when fully turned on than BJTs which means FETs generally have higher current ratings than comparable BJTs. However, comparable BJTs generally have higher voltage ratings than FETs. Perhaps the starkest contrast between the two transistor technologies is their drive requirements: FETs never require a constant current at their gate terminals, whereas BJTs *do* require base current to continuously activate.

[7] The natural resistive-capacitive time delay of an FET's gate may be minimized through external circuit design, but the stored charge within a BJT's base region is largely a function of a transistor's internal construction.

[8] Charge carriers crossing the base-emitter junction of a BJT require an energy investment to do so, and much of this energy later dissipates in the form of heat. No such intrinsic barrier exists within a FET, which means conductance is limited only by charge carrier density and channel area.

3.4 Simple step-down converter

A practical method for controlling average power delivered to an electrical load is to rapidly pulse power on and off to that load. The *duty cycle*[9] of that pulsing determines the percentage of full power received by the load. This technique is called *pulse width modulation* (PWM) and it is very commonly used to control the speed of electric motors as shown in the following diagram:

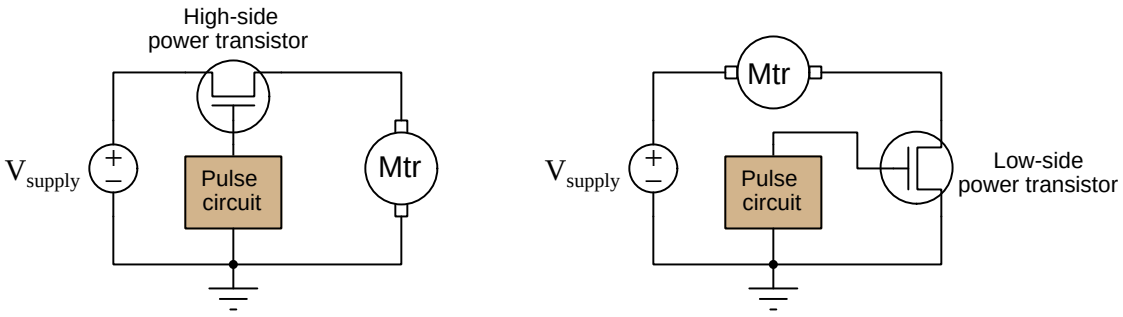

Here, a power transistor driven by a pulse-generating circuit (i.e. square-wave oscillator) functions as a fast on/off switch. When the transistor is fully on it drops negligible voltage and so its power dissipation is nearly zero ($P = IV \approx I0 = 0$); when the transistor turns off it passes no current and so once again its power dissipation is zero ($P = IV = 0V = 0$). Unlike a series dropping resistor or the transistor in a linear voltage regulator, the transistor dissipates negligible power and therefore operates at a cool temperature. Most importantly, we only draw enough energy from the source as is needed by the load, with none "wasted" in the form of heat.

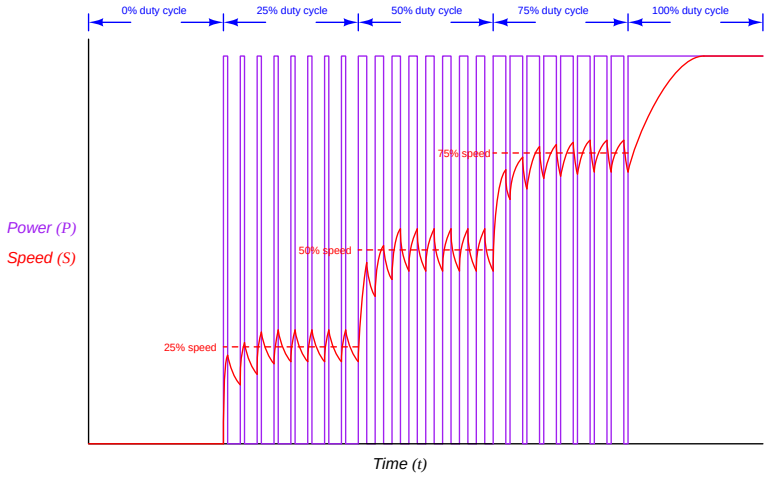

The sawtooth-shaped "rippling" seen in the motor's speed on this graph may be minimized by increasing the pulse frequency. However, it is still the *duty cycle* which determines the average

[9] The duty cycle of a square-edged pulse is the amount of "on" time divided by the total period.

amount of electrical power delivered to the motor and therefore its speed when driving a mechanical load.

Clearly, controlling electric motor power using pulse-width modulation switching is more energy-efficient than doing so using linear voltage-regulation techniques. However, switch-mode power control is not without its disadvantages. Here are some of them:

- Pulsing electrical power on and off to a load results in the load's power fluctuating rather than remaining steady. If the PWM frequency is sufficiently high, this pulsation may be minimized to the point of irrelevance, but for some types of loads this pulsation would be intolerable. For example, in precision measurement applications where we strive for maximum accuracy, the unsteady nature of pulsating voltages and currents complicates their measurement.

- The pulsing voltages and currents, being non-sinusoidal by nature, consist of infinite series of harmonic frequencies as predicted by Fourier's Theorem. These harmonic frequencies tend to "couple" easily to nearby circuitry which makes switch-mode power conversion a source of *electrical noise*. This noise is present not just at the load and its associated wiring, but also at the source feeding energy to the switch-mode power conversion circuit because that source is guaranteed to experience a pulsating load current. This means the noise produced by a switch-mode power conversion circuit will never be strictly confined to it and its load, but in fact will exist within the entire electrical power network sourcing energy to the converter. If that electrical power network sources multiple loads, then the converter's noise will carry to those other loads through the power network.

- Certain power-conversion devices such as transformers, if conveying electrical energy to a switch-mode converter circuit, may suffer excess heating due to the harmonic noise produced by the converter. Certain types of electrical networks, notably three-phase AC power circuits, are adversely affected by specific harmonics[10] of their fundamental frequencies, which means if a switch-mode power converter happens to generate any of those harmonics they may cause unexpected problems elsewhere.

- Some types of loads – especially electric motors – translate these electrical current pulsations into mechanical vibrations. If the PWM frequency is within the audible range, you will be able to actually *hear* the PWM waveform from the audible noise the motor emits!

[10] For example, so-called *triplen harmonics* in a three-phase AC circuit which have frequencies that are integer-multiples of the third harmonic, will manifest as additive neutral current in Wye-connected networks and as circulating current in Delta-connected networks. For a 60 Hz AC power system, triplen harmonics include the 3rd (180 Hz), the sixth (360 Hz), the ninth (540 Hz) of the 60 Hz fundamental.

Suppose we replace the motor in this PWM power control circuit with a resistive load and a capacitor. What we ought to expect with this exchange is a similar response to PWM duty cycle, except it will be load voltage and not motor speed that varies with duty cycle:

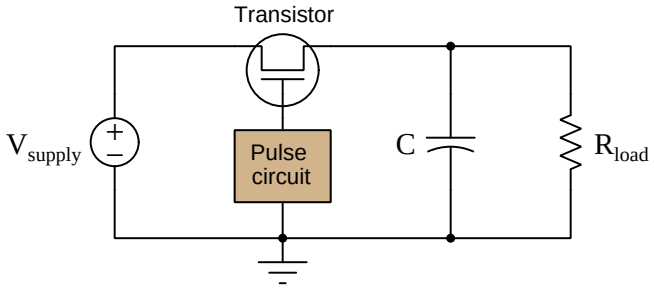

A helpful problem-solving strategy for analyzing switching converter circuits is to annotate all voltage polarities, current directions, and source/load roles within the circuit during the transistor's "on" state versus its "off" state:

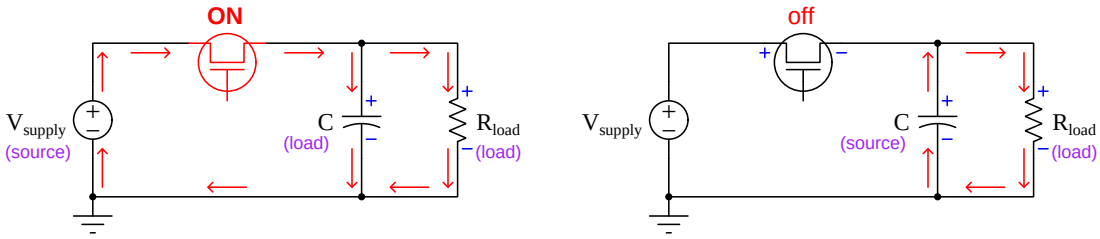

Note how the R_{load} always receives power, and how the capacitor alternately functions as a load and a source. Like the motor's speed, load voltage in this circuit will "ripple" up and down as the transistor switches on and off, but this ripple may be minimized by using a sufficiently high PWM frequency. The important outcome is that we now have a way to reduce the source's voltage down to a lower value at the load without using any energy-dissipating[11] mechanism such as a series resistor.

The Law of Energy Conservation still (and always!) holds true, and average load power must equal average source power for steady-state conditions. Therefore, if the load experiences less voltage than the source while operating at the same power (i.e. $P_{load} = P_{source}$), it must experience proportionately more current than the source. In other words, this circuit functions much like a step-down transformer, except that it can operate with DC rather than requiring AC. Instead of winding turns ratios dictating the voltage step-down ratio as it is in a transformer, the duty cycle of the pulse signal will determine the voltage reduction ratio.

[11] An ideal capacitor dissipates no energy whatsoever, eventually returning 100% of the energy delivered to it. Real capacitors fall short of this perfect mark, but only slightly.

With the step-ratio being adjustable, we may augment this circuit's PWM pulse circuit with negative feedback sensing the load voltage so that it actually functions as a voltage *regulator*. A sample design appears below, using a 555 timer as a sawtooth-wave relaxation oscillator and a comparator to compare this sawtooth voltage signal against the fed-back voltage. The comparator's output provides the PWM signal necessary to drive the power MOSFET:

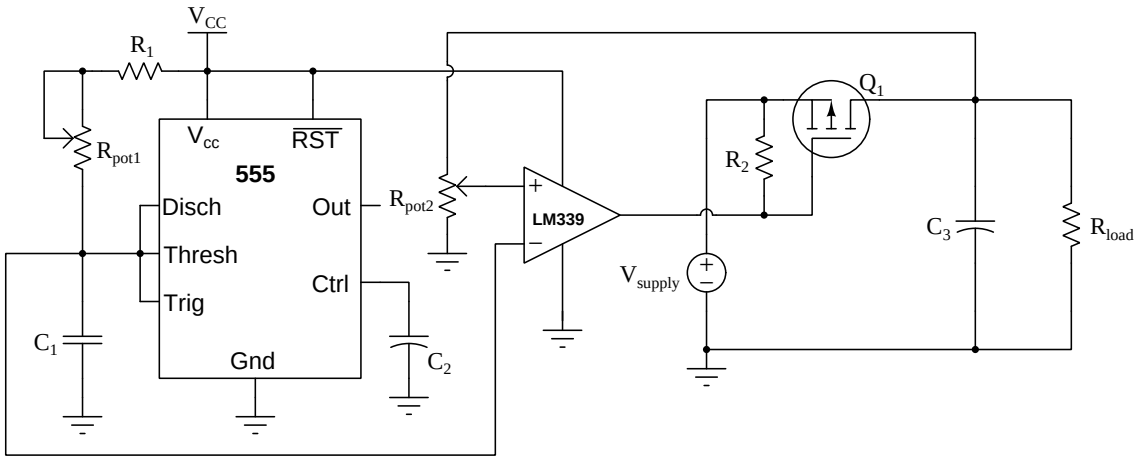

If load voltage decreases for any reason, the comparator will sense less voltage at its noninverting input. This will cause its output to remain "low" for a longer percentage of the period, which in turn causes the P-channel MOSFET to remain on for longer. With the transistor turning on for longer than before, the load voltage will rise to a higher average value and thus compensate for the initial droop. Potentiometer #1 controls the PWM frequency for an acceptable amount of ripple voltage at the load, while potentiometer #2 controls how much of the load's voltage gets fed back to the comparator. Moving the wiper of R_{pot2} in the upward direction causes the comparator to sense a greater load voltage, which in turn causes it to decrease Q_1's duty cycle and thereby decrease load voltage[12].

[12] A helpful technique for analyzing negative feedback loops is to anthropomorphize the comparator's function. If moving R_{pot2}'s wiper upward causes a stronger feedback signal to appear at the comparator's noninverting input, we may say to ourselves that the comparator now "thinks" load voltage has spiked upward, and will "try" to bring that voltage "back down to where it belongs". Of course, a comparator lacks the intelligence of a human being, but anthropomorphic thinking is often helpful to simplify confusing topics such as negative feedback. The essence of the technique is that you imagine yourself having to perform the comparator's function, and ask yourself what you would need to do in response to some stimulus.

3.5 Buck converter

A switch-mode converter circuit designed to output less voltage than it takes in is called a *buck*[13] converter, and does for DC what a step-down transformer does for AC. Below is the schematic diagram for a simple buck converter using both an inductor and a capacitor as energy-storage elements:

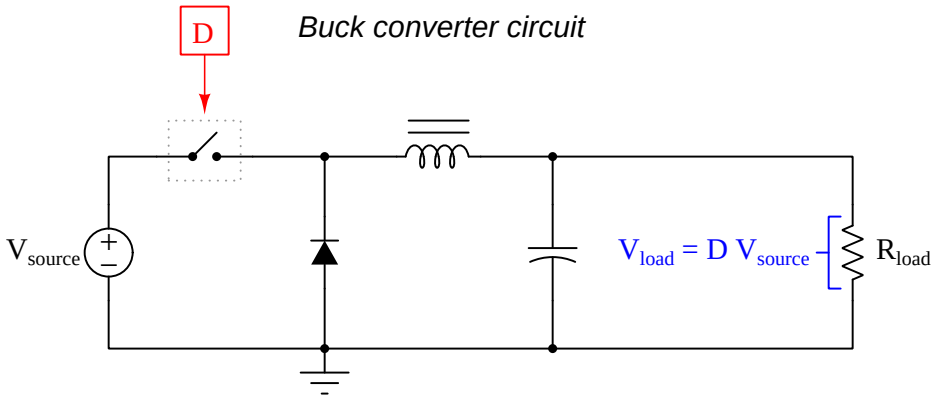

D represents a pulse width modulation (PWM) signal with a duty cycle value driving a switching element such as a transistor. Connected to a very light load (or no load at all), the output voltage for this buck converter is simply source voltage multiplied by *D*. For example, if the DC source was 24 Volts and the duty cycle was 40%, the output voltage should ideally be 9.6 Volts. Of course, with a load connected the output voltage will be less than this – how much less depending on how "heavy" the load is as well as the various inefficiencies of the converter circuit's components.

The transistor provides the connection from the DC source to the rest of the circuit, while the diode provides a complete path for inductor current when the transistor turns off. One way to improve the energy efficiency of this converter circuit is to use a diode with a low forward voltage drop, such as a Schottky diode, or replace the diode entirely with a transistor controlled to always assume the opposite state of the first transistor[14]. Energy efficiency will be maximized to the extent that we may minimize all forms of energy dissipation in the circuit.

A practical application for a buck converter is to combine it with some negative feedback to make it self-regulating, then use it to replace any linear voltage regulator circuit for greater energy efficiency. Recall that linear regulators always output less voltage than they take in, because they function by dissipating energy. Since buck converters similarly output less voltage than they take in, but do so with little energy "loss", they may substitute for many applications ordinarily utilizing linear voltage regulators. The improved conversion efficiency is especially useful when the energy source is a battery, or some inherently limited source such as a small solar panel.

[13] The meaning of the word "buck" in this context is to *work against* something. In this type of DC-DC converter circuit, the output voltage is less than the input voltage, and so we can say the converter "works against" the source voltage to provide less to the load.

[14] Using a second transistor instead of a diode will make this converter *synchronous*, referring to the complementary states of the two transistors and the precision signal timing necessary to make this happen.

Applying the problem-solving strategy of voltage, current, and source/load annotation to this converter circuit in both transistor states, we see how the inductor and capacitor exchange roles as sources and loads:

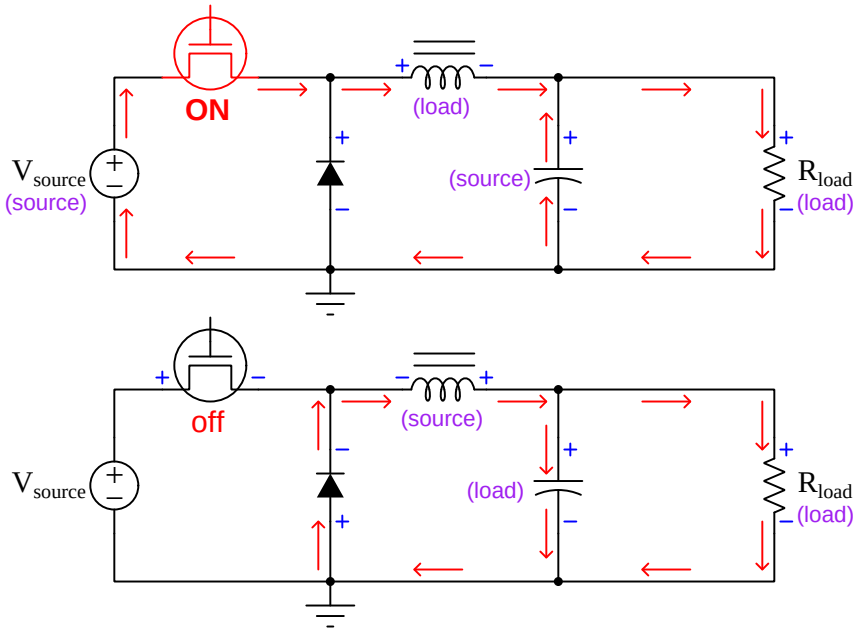

The problem-solving technique of *limiting cases* helps us make sense of the $V_{load} = DV_{source}$ equation: when D is 0% the transistor is off all the time, which means nothing will receive energy from V_{source}. As a consequence V_{load} must be zero. Conversely, when D is 100% the transistor is on all the time which sends full current through the inductor and full voltage across the capacitor such that $V_{load} = V_{supply}$.

Noteworthy limitations of the buck converter circuit include the following:

- The transistor is located on the "high side" of the circuit. This complicates matters for transistor selection and for transistor drive circuitry.
- The source current strongly pulsates between maximum and zero. These fast rise and fall times for the current pulses are equivalent to extremely high-frequency AC waveforms, leading to undesirable "coupling" with neighboring circuits (i.e. interference) if unfiltered by additional components on the front end of the converter circuit.

When using a MOSFET, the positioning of the transistor on the positive rail requires either an N-channel MOSFET with a gate-driving signal pulsing to a higher voltage level than V_{source} or a P-channel MOSFET with higher drain-source "on" resistance than an N-channel MOSFET[15]. Neither transistor option is ideal:

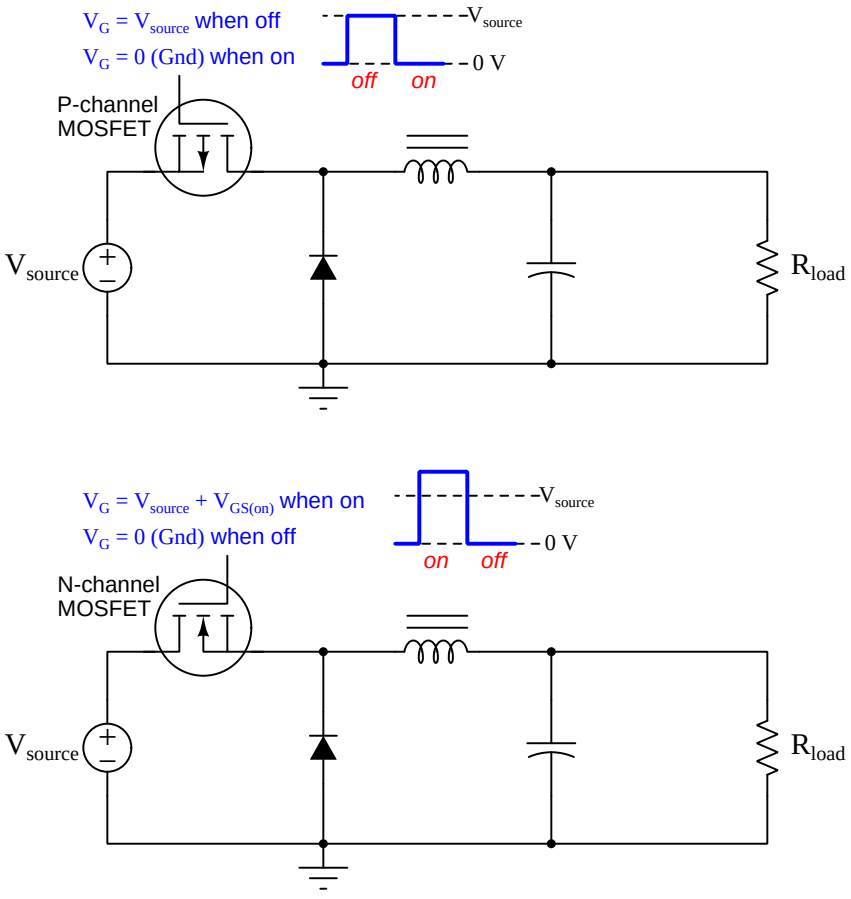

[15] Recall that field-effect transistors are majority-carrier devices, and so the charge carriers in a P-channel FET are holes rather than electrons. Holes have substantially less mobility than electrons, all other factors being equal, and this makes P-channel MOSFETs less efficient (i.e. more resistive, with greater I^2R dissipative losses) as switches than N-channel MOSFETs.

The problem of input current pulsations is common to most DC-to-DC converter designs, and may be minimized by adding a low-pass filter stage to the input of this converter:

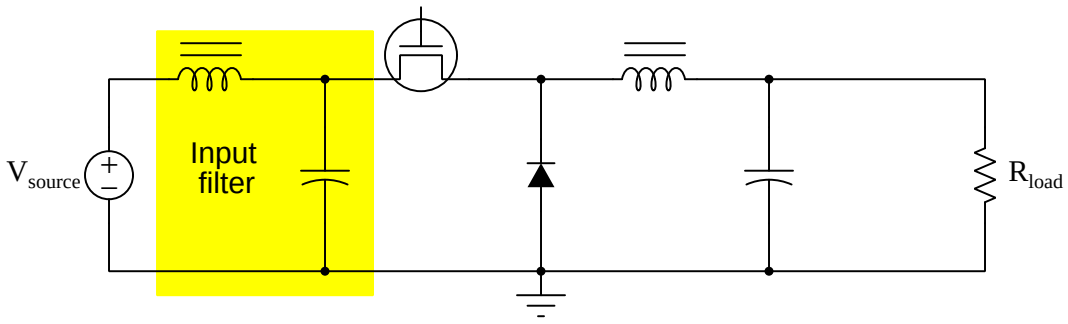

However, this solution adds complexity, bulk, and expense to any converter circuit.

3.6 Boost converter

A switch-mode converter circuit designed to output more voltage than it takes in is called a *boost* converter, and does for DC what a step-up transformer does for AC. Below is the schematic diagram for a simple boost converter using both an inductor and a capacitor as energy-storage elements:

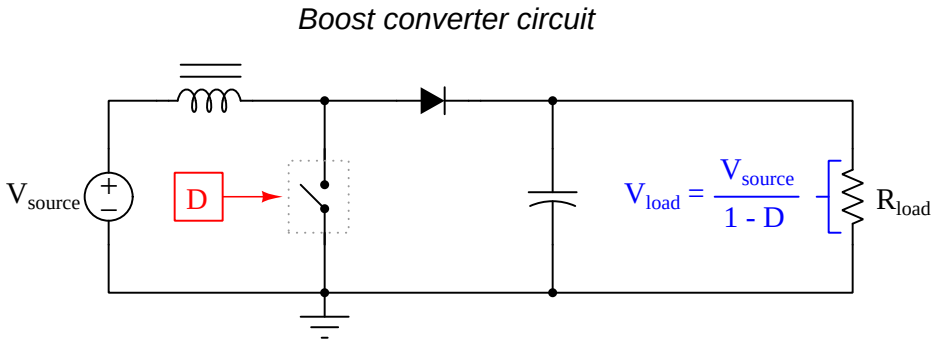

D represents a pulse width modulation (PWM) signal with a duty cycle value driving a switching element such as a transistor. The unloaded output voltage for this boost converter is source voltage divided by $1 - D$. For example, if the DC source was 14 Volts and the duty cycle was 30%, the output voltage should be 20 Volts. Of course, this will only be true if the load is very light or non-existent (i.e. negligible load power). Real loads dissipating energy will cause the output voltage to "sag" below this ideal predicted voltage.

The transistor provides a direct path between the DC source and the inductor for its energization, while the diode prevents the capacitor from discharging through the "on" transistor. As with the buck converter, one way to improve the energy efficiency is to use a diode with a low forward voltage drop, such as a Schottky diode, or replace the diode entirely with a transistor controlled to always assume the opposite state of the first transistor[16]. Energy efficiency will be maximized to the extent that we may minimize all forms of energy dissipation in the circuit.

A practical application for a boost converter is to power electronic circuitry operating from a single-cell battery power source of limited voltage. For example, if the battery is a lithium-ion cell outputting 3.6 Volts and we wish to power some digital logic gates requiring 5.0 Volts, a boost converter may be used to bridge this gap. Moreover, if we equip the converter with some negative feedback in the PWM signal so that it regulates its output voltage, it will maintain a constant 5.0 Volts output even as the battery's voltage gradually declines.

[16] Using a second transistor instead of a diode will make this converter *synchronous*, referring to the complementary states of the two transistors and the precision signal timing necessary to make this happen.

Applying the problem-solving strategy of voltage, current, and source/load annotation to this converter circuit in both transistor states, we see how the inductor and capacitor exchange roles as sources and loads:

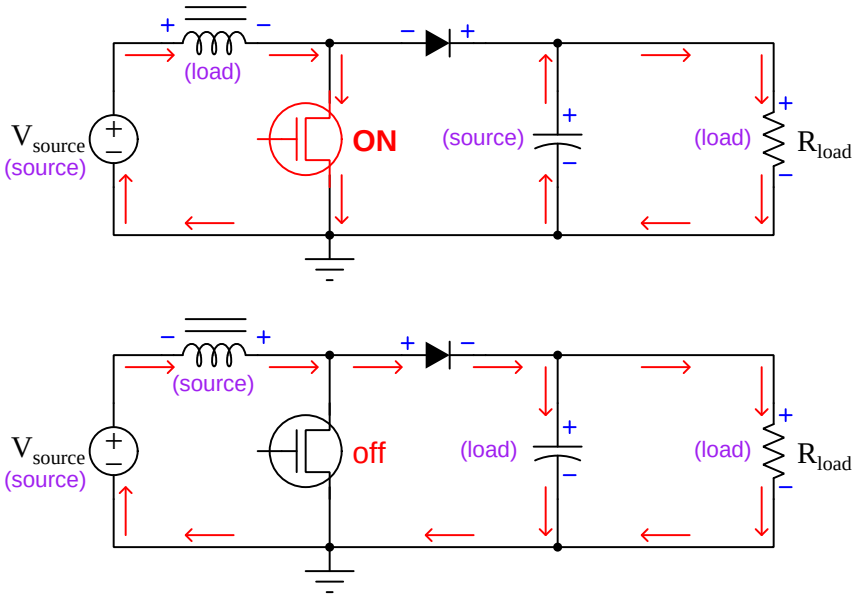

The problem-solving technique of *limiting cases* helps us make sense of the $V_{load} = \frac{V_{source}}{1-D}$ equation: when D is 0% the transistor is off all the time, which means the DC source connects to the load through the inductor and diode. As a consequence $V_{load} = V_{source}$ (assuming negligible diode voltage drop). Conversely, when D is 100% the transistor is on all the time which sends full current through the inductor but prevents the rest of the circuit from receiving any energy, making V_{load} zero. At first this conclusion seems to contradict the equation, but on close inspection we realize the equation gives us an undefined result when $D = 100\%$ (since division by zero is mathematically undefined). A better "limiting case" scenario would be to imagine some duty cycle value very nearly equal to 100% but not fully 100%. This would result in the inductor having the majority of the pulse's period to absorb energy as a load but only a very short time duration to deliver energy as a source. This brief delivery of energy must equate to a fast rate-of-change for inductor current ($\frac{dI}{dt}$), leading to a high voltage output by the inductor ($V = L\frac{dI}{dt}$). This higher voltage gets output to the load as the equation predicts.

It is worth noting that "low-side" the position of the transistor in the boost converter design is much more advantageous than the "high side" transistor location of the buck design. Having one terminal of the transistor grounded permits us to use a (low-resistance) N-channel MOSFET and drive it with a simple pulse signal ranging from 0 Volts to $V_{GS(on)}$.

Also, since the transistor does not completely interrupt the DC source, that source's current will not pulsate fully between zero and full value. Instead, the source's current will merely "ripple" as the inductor absorbs and releases energy. This translates into less potential for interference with neighboring circuits. However, if this ripple is still too great you have the option of adding a low-pass

LC filter stage to the input of the converter just as we saw with the buck converter design.

3.7 Inverting converter

A switch-mode converter circuit able to output either more or less voltage than it takes in is called a *buck-boost* converter. Below is the schematic diagram for an elementary buck-boost converter using both an inductor and a capacitor as energy-storage elements. It also happens to invert the polarity of the source voltage, hence the name *inverting converter*:

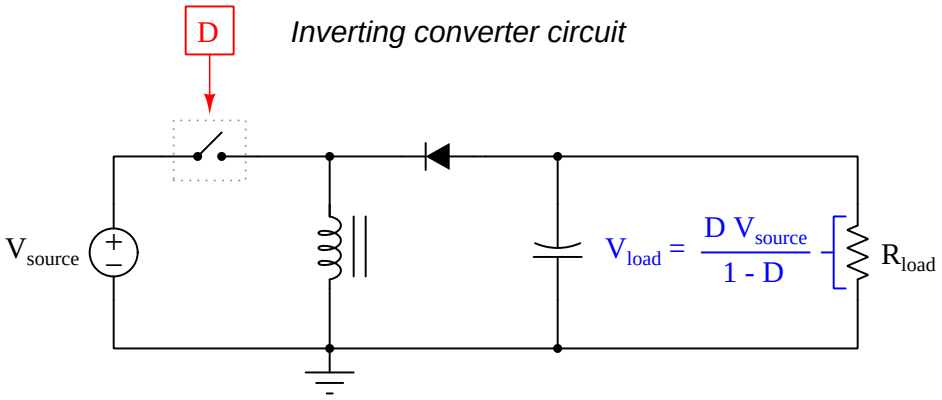

D represents a pulse width modulation (PWM) signal with a duty cycle value driving a switching element such as a transistor. The unloaded output voltage for this buck-boost converter is source voltage multiplied by the duty cycle (D) and divided by the complement of the duty cycle ($1 - D$). For example, if the DC source was 15 Volts and three different duty cycle values of 25%, 50%, and then 75% were applied, the output voltage should be 5 Volts, 15 Volts, and 45 Volts, respectively. The usual caveat of load power applies: this predicted output voltage assumes no power dissipated by the load, which for any realistic application will not be the case. Driving a load, the converter's output voltage will be something less than the voltage predicted by this simple formula.

Just as in the buck converter, the transistor provides the connection from the DC source to the rest of the circuit, while the diode provides a complete path for inductor current when the transistor turns off. One way to improve the energy efficiency of this converter circuit is to use a diode with a low forward voltage drop, such as a Schottky diode, or replace the diode entirely with a transistor controlled to always assume the opposite state of the first transistor[17]. Energy efficiency will be maximized to the extent that we may minimize all forms of energy dissipation in the circuit.

[17] Using a second transistor instead of a diode will make this converter *synchronous*, referring to the complementary states of the two transistors and the precision signal timing necessary to make this happen.

Applying the problem-solving strategy of voltage, current, and source/load annotation to this converter circuit in both transistor states, we see how the inductor and capacitor exchange roles as sources and loads:

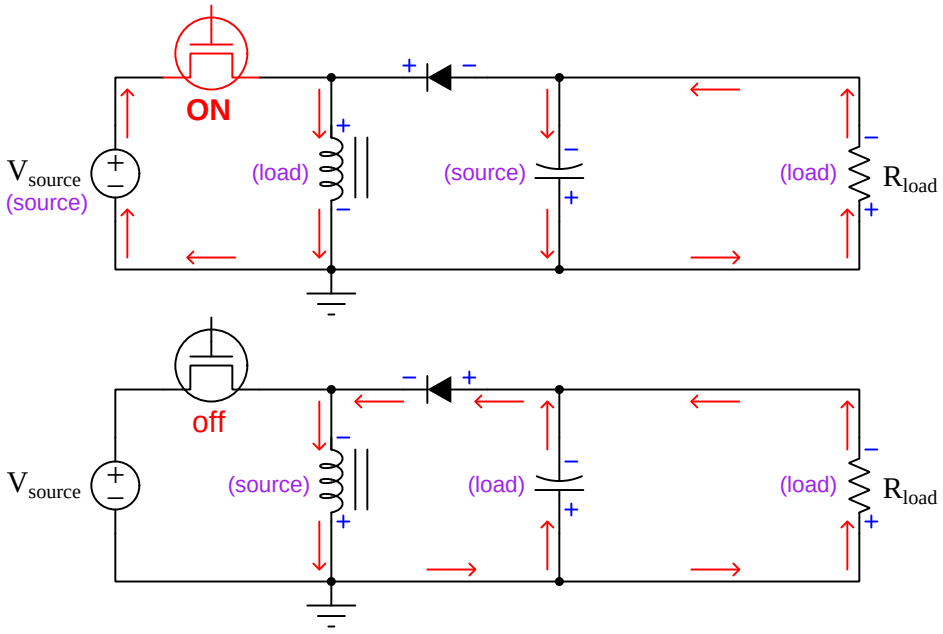

The problem-solving technique of *limiting cases* helps us make sense of the $V_{load} = \frac{DV_{source}}{1-D}$ equation: when D is 0% the DC source is entirely disconnected from the rest of the circuit, and load voltage must go to zero. Conversely, when D is 100% the transistor is on all the time which sends full current through the inductor, but due to the diode's blocking action the rest of the circuit receives no energy, also making V_{load} zero. At first this conclusion seems to contradict the equation, but on close inspection we realize the equation gives us an undefined result when $D = 100\%$ (since division by zero is mathematically undefined). A better "limiting case" scenario would be to imagine some duty cycle value very nearly equal to 100% but not fully 100%. This would result in the inductor having the majority of the pulse's period to absorb energy as a load but only a very short time duration to deliver energy as a source. This brief delivery of energy must equate to a fast rate-of-change for inductor current ($\frac{dI}{dt}$), leading to a high voltage output by the inductor ($V = L\frac{dI}{dt}$). This high voltage, series-aiding with the DC source, results in an even higher voltage at the load as the equation predicts.

Practical limitations of this "inverting" buck-boost converter circuit include the same limitations of the simpler buck converter, all stemming from the location of the transistor on the "high" side of the DC source. This makes transistor operation less than ideal, and it also results in an input current that strongly pulsates, since every "off" phase of the transistor forces input current to zero. As with the buck converter, these problems are not insurmountable, but the solutions add complexity, bulk, and expense to the circuit.

3.8 Ćuk converter

A more recent DC-DC converter design was introduced by Serbian-American power electronics engineer Slobodan Ćuk[18]. It is a buck-boost style of converter, and it also happens to invert the source voltage polarity, but it offers distinct advantages[19] over the traditional buck-boost design:

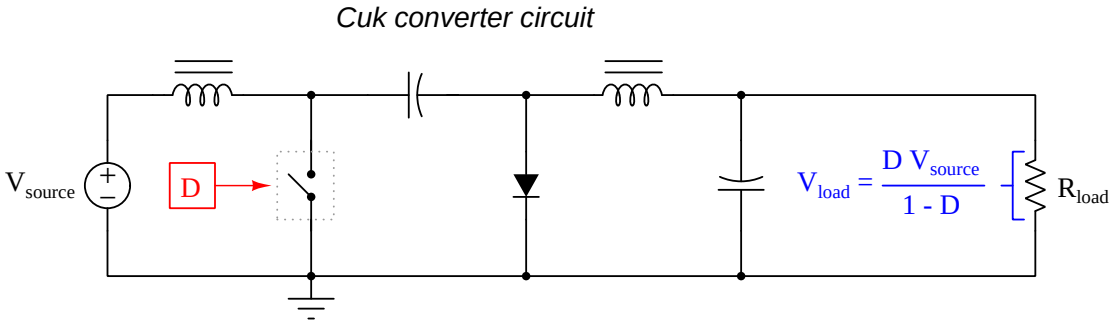

This circuit's characteristic equation is identical to the traditional buck-boost converter, being source voltage multiplied by the duty cycle (D) and divided by the complement of the duty cycle ($1 - D$). For example, if the DC source was 18 Volts and three different duty cycle values of 25%, 50%, and then 75% were applied, the unloaded output voltage should be 6 Volts, 18 Volts, and 54 Volts, respectively. Powering a realistic load will result in an output voltage slightly less than the predicted value(s).

The transistor provides a direct path between the DC source and the inductor for its energization, while the diode prevents the first capacitor from discharging through the "on" transistor. Like all switch-mode power circuits, Schottky diodes or synchronized transistors may be used in place of the simple diode shown in this diagram in order to minimize energy dissipation and improve conversion efficiency.

Note how in the Ćuk converter design both the transistor switch and the diode connect to ground, which is advantageous because it makes these devices "low-side" switching elements. As we saw with the boost converter, having the switching elements connected to ground on one side permits the use of N-channel MOSFETs which are easy to drive and have low channel resistance when on.

[18] Pronounced, "chook"

[19] First and foremost, all the switching elements in a Ćuk converter have grounded terminals (i.e. low-side switches) which means we may use N-channel MOSFETs with modest gate drive voltages rather than high-side P-channel MOSFETs (higher $R_{DS(on)}$) or high-side N-channel MOSFETs (requiring a special higher-voltage PWM driver circuit to provide enough V_{GS} to saturate on). Secondly, Ćuk converters have the ability to minimize ripple by cleverly coupling the two inductors, which will be discussed later in this section.

In Dr. Ćuk's patent for this converter design, he shows the switching elements (transistor and diode) as a single SPDT switch, which is a helpful aid to understanding how this circuit functions:

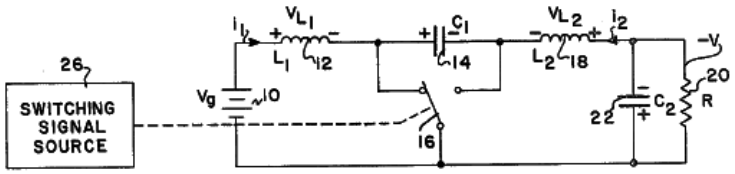

FIG. 5

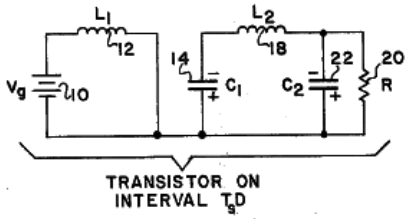

FIG. 5(1)

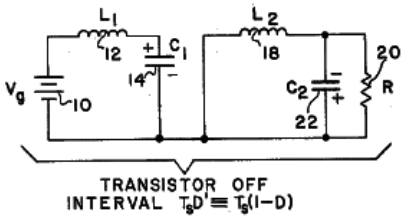

FIG. 5(2)

When the transistor is on (i.e. SPDT switch thrown to the left-hand position), the DC source transfers energy into the first inductor. When this transistor turns off and the diode turns on (i.e. SPDT switch thrown to the right-hand position), the inductor and DC source work together to transfer energy into the first capacitor. This capacitor, during the next transistor "on" phase, transfers its stored energy to the second set of storage elements: the inductor and filter capacitor on the load-side of the converter.

We also see from these diagrams that the first capacitor (C_1) is the main "transfer" device taking energy from the left-hand (input) side of the converter and delivering energy to the right-hand (output) side. In traditional DC-DC converter designs an inductor fulfills this role.

Applying the problem-solving strategy of voltage, current, and source/load annotation to this converter circuit in both transistor states:

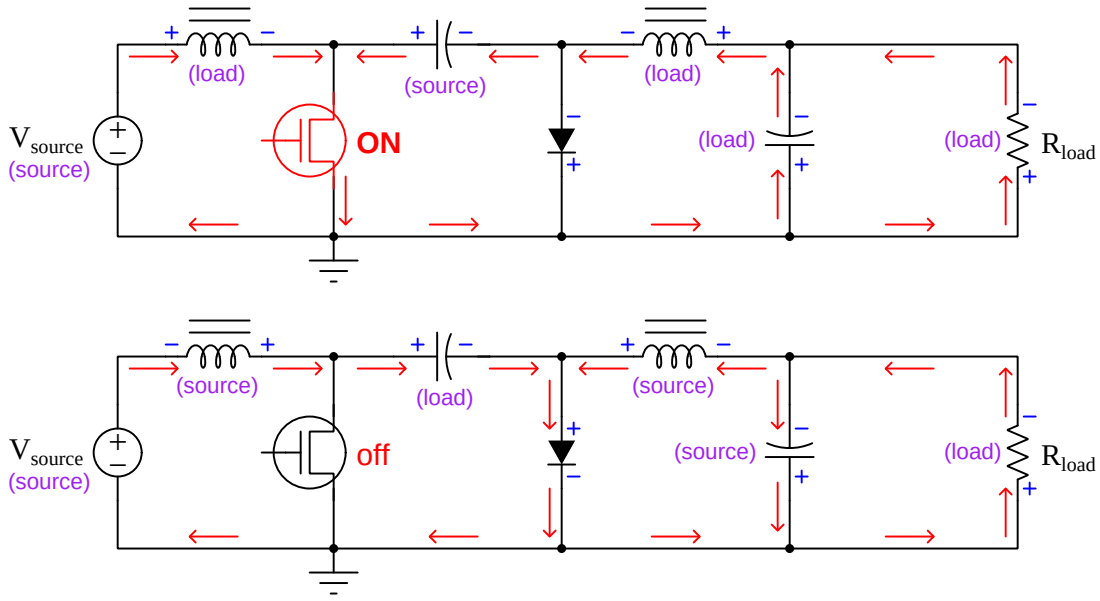

As with all DC-DC converters having the ability to boost voltage as a function of duty cycle, the major factor in voltage boost is the amount of time over which an inductor is forced to release its stored energy: long time durations equate to slow $\frac{dI}{dt}$ rates of change, and thus small amounts of inductor voltage; short time durations equate to fast $\frac{dI}{dt}$ rates and high voltages. It should not surprise us, then, to find the Ćuk converter exhibits the same characteristic equation as the buck-boost converter. When duty cycle is low the front-end inductor has ample time to deliver its energy to the capacitor for transfer, and when duty cycle is high the inductor is hastened to transfer that energy in less time which results in the capacitor being charged up to a higher voltage.

An innovative feature introduced by Dr. Ćuk in his patent for this converter circuit is the ability to "smooth" the ripple current on either the input (source) or output (load) by magnetically coupling the two inductors together:

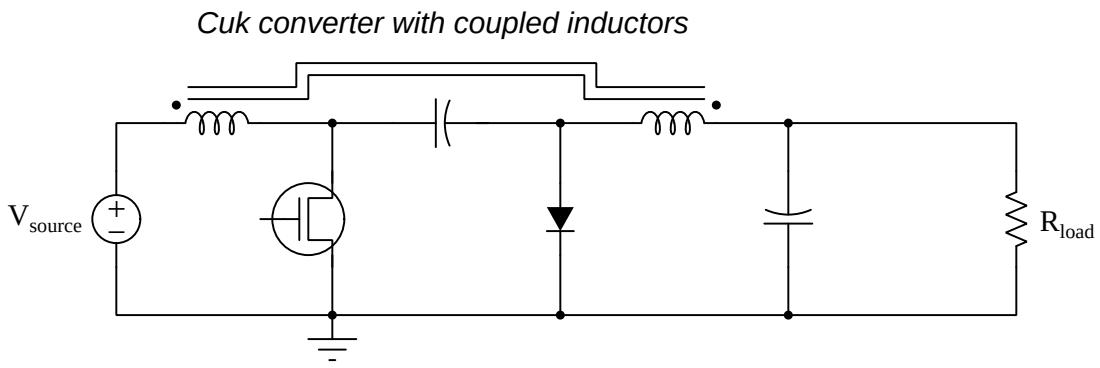

This coupling between the two inductors is intentionally "loose" (i.e. a substantial amount of the magnetic flux for each coil does *not* link with the other coil), transferring some of each inductor's energy to the other, but not all of it. By carefully proportioning this degree of coupling it becomes possible to stabilize current through one of them at the expense of ripple in the other. This is a useful feature when "clean" DC current is more beneficial on one particular side of the circuit than the other, and where additional filtering as a brute-force solution to ripple would be impractical.

3.9 Flyback converter

DC-DC converters typically use self-inductance as a way to store and release energy for the purpose of achieving a ratio of voltage from source to load, but mutual inductance is also useful for this purpose. An example of a DC-DC converter exploiting mutual inductance to transfer energy from source to load is the so-called *flyback*[20] *converter* shown below:

Flyback converter circuit

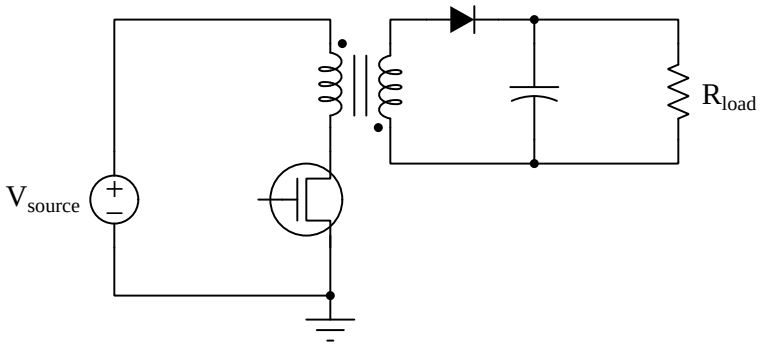

Notice how there is no electrical commonality between the input (source) and output (load) circuits in this converter – the two sides are *electrically isolated* from one another. This electrical isolation is a useful feature for many applications, for example in medical electronic devices where we must ensure there is no possible connection between the patient's body and the electrical power supply for a machine (to avoid shock potential), or to avoid *ground loops* in industrial signal systems. It also makes load voltage polarity arbitrary, unlike the non-isolated converter circuits explored previously in which load voltage polarity was determined by source voltage polarity.

[20] The term "flyback" is an anachronism based on the operation of legacy cathode-ray tube (CRT) television monitors, where the electron beam inside the CRT had to "fly back" across the viewscreen in order to begin retracing its path to update the video frame. A special "flyback transformer" inside the television set developed high voltage for the CRT's operation based on mutual inductance from the deflection current signal's rapid rate-of-change ($\frac{dI}{dt}$) during each fly-back phase.

Applying the problem-solving strategy of voltage, current, and source/load annotation to this converter circuit in both transistor states:

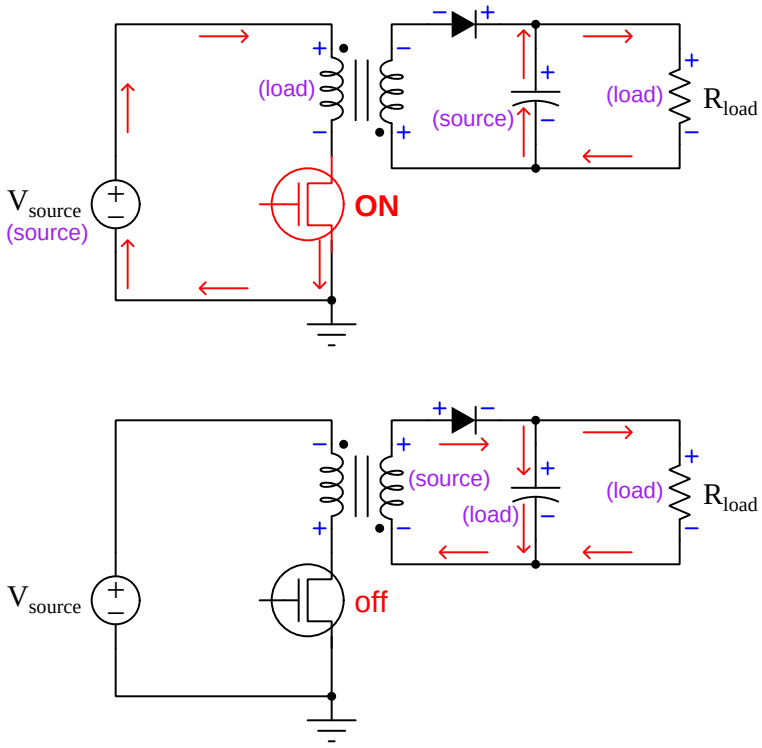

Note how the two inductors *take turns* absorbing and releasing energy. This is decidedly different from how two windings behave in a regular AC transformer: one winding functioning as a load and another as a source *simultaneously*.

3.10 Forward converter

The *forward converter* design resembles the flyback converter, except for the "dot" polarity markings on the coupled inductors:

Forward converter circuit

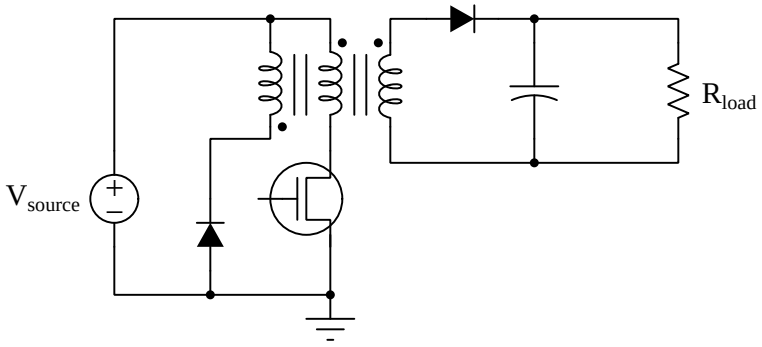

Like the flyback converter, the forward converter enjoys full electrical isolation between the input (source) and output (load) sides. The primary difference between these two converter designs is the timing of the energy transfer. In the flyback converter, energy stored in the left-hand inductor during the transistor's "on" phase would not be released by the right-hand inductor until the transistor turned "off". Here, the polarity is such that energy transfer occurs during the transistor's "on" phase.

A striking difference between the forward converter and the flyback converter is the presence of a third inductor wound around the common core of the other two, connected to the DC source through a diode. It will not become apparent why this third winding or diode exists in the circuit until we analyze its operation in both states of the transistor, as shown on the next page.

Applying the problem-solving strategy of voltage, current, and source/load annotation:

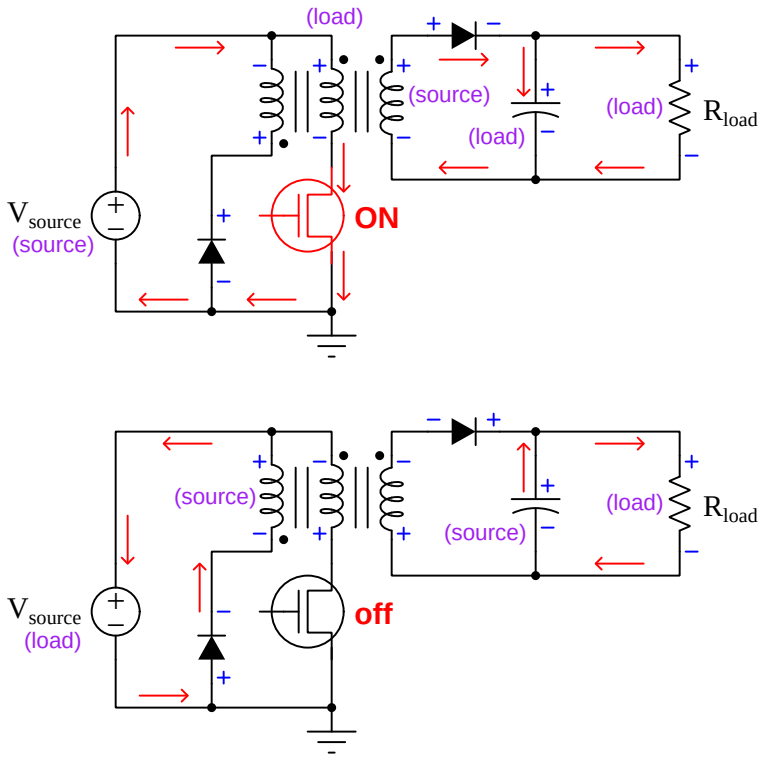

Note how two of these inductors exchange energy at the same time, one functioning as a load (receiving energy from the DC source) and the other as a source (powering R_{load}) simultaneously while the transistor is in its "on" state. This is more in keeping with the function of a transformer, where primary and secondary windings function as source and load (respectively) at the same time. Indeed, it helps explain the name of this converter circuit: "forward" meaning that energy is transferred to the load during the time when the DC source is advancing energy into the magnetic core, as opposed to a "back" converter which would transfer energy to the load when the DC source becomes disconnected and the magnetic field is diminishing (retreating).

The third coil and diode is there to help *reset* the magnetic flux in the core back to zero during the time the transistor is off. If it were not for this third winding and diode, some residual magnetic flux would remain in the core after every cycle of the transistor, and during each additional "on" phase that magnetic flux would accumulate to greater levels (since an *increase* in flux is necessary to induce a voltage in the right-hand winding sufficient to act as a source, by Faraday's Law of Electromagnetic Induction: $V = N\frac{d\Phi}{dt}$), eventually *saturating* the core with magnetic flux. Providing a pathway for current gives the energy stored in that residual magnetic flux a place to go, namely back to the DC source. If the DC source is a chemical battery, this reverse current will re-charge the battery with a small amount of energy every time the transistor switches "off". Other sources (e.g. solar photovoltaic cells) would not benefit from this reset current and would instead dissipate the core's energy as heat.

3.11 Continuous and discontinuous converter operation

All switch-mode conversion circuits exploit the principle of stored and released energy using inductors and/or capacitors, typically some combination of these two component types. Inductors store energy within magnetic fields, their magnetic field strength being a function of current ($E_L = \frac{1}{2}LI^2$). Capacitors store energy within electric fields, their electric field strength being a function of voltage ($E_C = \frac{1}{2}CV^2$). Therefore, the amount of energy stored within an inductor may be inferred by measuring its current at any given point in time, while a capacitor's store of energy is evident by its measured voltage.

If we operate a DC-DC converter circuit such that all of its reactive (energy-storing) components retain a substantial amount of their energy during all portions of the switching cycle, we call that operation *continuous mode*. If, by contrast, we allow one or more of the energy-storing components to become exhausted of energy during any portion of the switching cycle, we call that operation *discontinuous mode*. This is to say, if a converter operates in such a way that every one of its inductor currents never falls to zero, and every one of its capacitor voltages never falls to zero, then we say it is operating in continuous mode. Conversely, if any reactive component within a converter circuit becomes depleted of energy at any point in time during a switching cycle (i.e. any inductor current falls to zero, any capacitor voltage falls to zero), that converter is operating in discontinuous mode.

A useful analogy for the operation of any switch-mode power converter circuit's reactive components (L and/or C) is to envision each of those energy-storing components as being a cup used to transfer water from one reservoir to another. Switching frequency is analogous to the speed at which the cup moves back and forth between the two streams, filling from one and emptying into the other. Duty cycle is analogous to the relative amount of time spent filling from one reservoir versus emptying into the other. In converter circuits containing more than one reactive component (as most do!) the first cup transfers water from a reservoir to the top of an open funnel, where that funnel gradually empties itself into something else. If the circuit in question just contains two reactive components, that "something else" is the load. If the circuit contains more than two reactive components, the "something else" may be another funnel.

Continuous-mode operation is where the switching happens often enough that the cup and funnel(s) never entirely fill nor entirely empty, but always contain some water. Discontinuous mode is where we let the cup fill to overflowing, and/or entirely pour out its contents. A cup allowed to overflow or run empty isn't the most effective transfer vessel. A funnel allowed to run dry means the load will see interruptions in power.

The canonical voltage equations usually shown for various types of DC-DC converters all assume continuous-mode operation. If the converter circuit in question ever falls into discontinuous mode, its characteristic output voltage equation will no longer be accurate. That said, discontinuous mode is not necessarily a poor or dangerous state. In fact, it is possible for virtually any well-designed converter circuit to become discontinuous[21] given the requisite supply and/or load conditions.

In most converter topologies the component experiencing the greatest variations in energy state

[21] Imagine a voltage-regulated switch-mode converter circuit with its load suddenly disconnected. The output capacitor, no longer having a load to deplete its energy store, will simply maintain voltage at a constant level while the regulated switching circuit will stop "filling" the inductor because no more energy needs to be given to the capacitor. It's as though someone plugged the spout of the funnel, resulting in the funnel's water level holding steady.

is the inductor (), and for this reason these modes are commonly referred to as *Continuous Current Mode* (CCM) and *Discontinuous Current Mode* (DCM) in honor of the inductor's current being the key variable.

Empirically determining the operational mode of a converter circuit is best done using an oscilloscope to monitor unfiltered voltages relevant to the energy-storing components. Consider the case of an inverting converter, with an oscilloscope probe connected across the inductor:

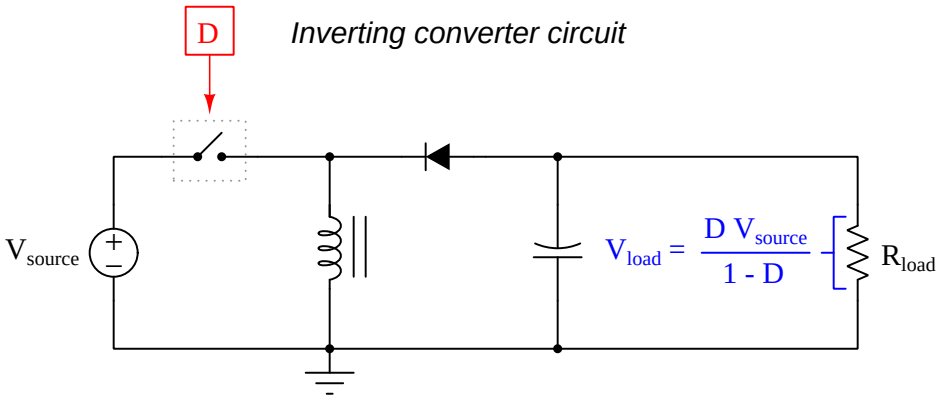

When the transistor turns on, we would expect to see the inductor's voltage immediately jump to and hold at the DC source's voltage level. When the transistor turns off, we expect the inductor's voltage polarity to immediately reverse as it transitions from being a load to being a source, and that voltage will decrease as it delivers energy to the capacitor and to the load:

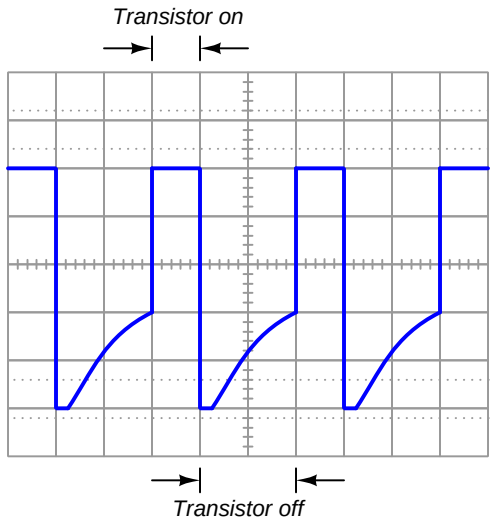

Stored energy within an inductor is a function of its *current* and not its *voltage*, but current probes are rare accessories for oscilloscopes, and for this practical reason we must make do with voltage

measurements. From the oscillograph shown above, though, we have all the necessary information to determine that the inductor never completely exhausts its energy store: note how the curved portion of the trace never completely levels off during the "transistor off" interval. This tells us the inductor is still in the process of discharging (i.e. delivering energy to the capacitor and load) every time the transistor turns back on.

By contrast, consider this next oscillograph for the same inverting converter circuit:

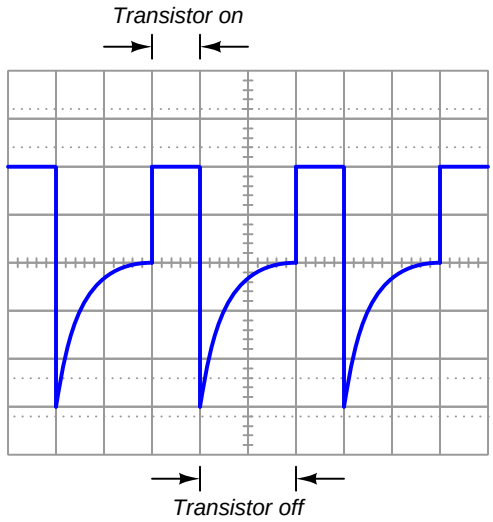

The shape of the curve indicates very clearly that the inductor's store of energy is completely exhausted at the end of each "transistor off" interval.

A good way to conceptualize these different modes of converter operation is to consider the transistor's switching period in terms of *time constant* (τ) for the energy-storing components. If it is a capacitor completely losing voltage during portions of the switching cycle, then we may think of this as the network having an insufficient RC time constant; if it is an inductor completely losing current at various moments in time we may think of this as the network having too short of an $\frac{L}{R}$ time constant. If avoiding discontinuous operation is our goal, then choosing larger C and/or L component values would be an effective solution. We might also consider shortening the transistor switching period (i.e. increase switching frequency) to give these components less time to deplete of energy.

Boost converter circuits are similarly easy to determine the operating mode of by oscillograph. Consider the following boost converter where the oscilloscope probe is connected across the switch (transistor):

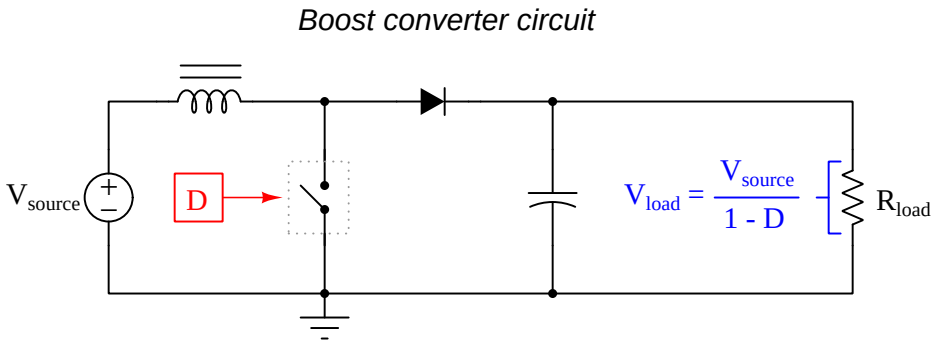

The left-hand oscillograph shows continuous operation, while the right-hand oscillograph shows discontinuous:

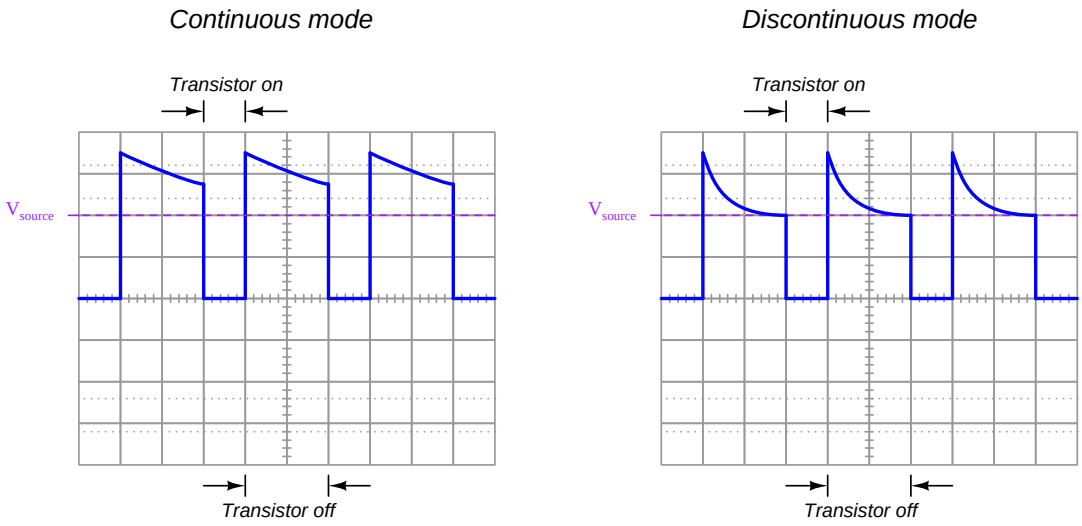

In continuous-mode operation we see the voltage decay over time, but not to the point where it completely levels off. In discontinuous-mode operation we see clear evidence of the inductor completely exhausting its energy store during each "transistor off" interval, where the decay curve achieves a level slope.

Discerning inductor status from voltage (and not current) oscillographs requires we think very carefully about all the components within the converter circuit. Some converter topologies make this easy, but others are less obvious. Consider a buck converter circuit, with an oscilloscope probe connected across the diode:

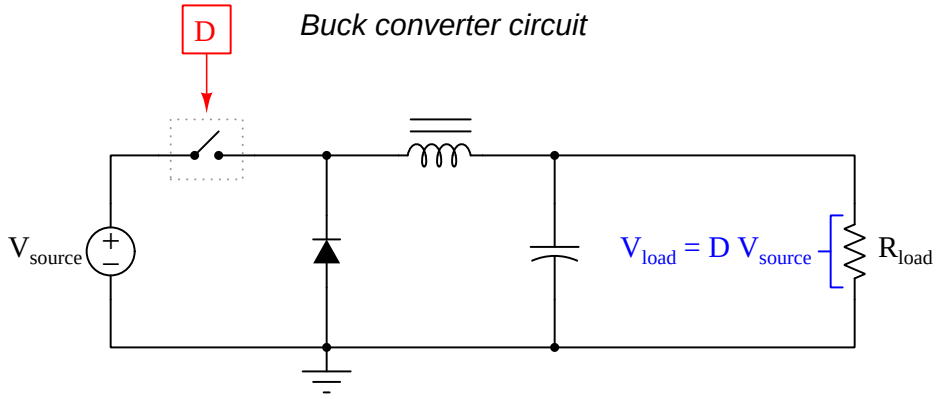

Continuous (left) and discontinuous (right) oscillographs hardly differ in appearance:

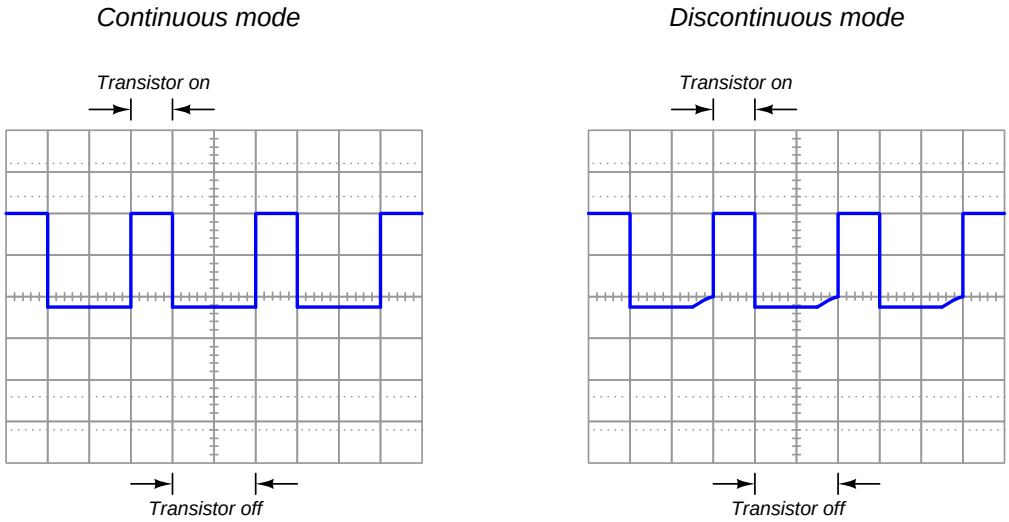

Since the oscilloscope measures diode voltage drop, the only indication of inductor current will be that diode dropping its expected forward voltage during the "transistor off" interval. This voltage will be quite small, approximately 0.7 Volts for a regular silicon rectifying diode and approximately 0.4 Volts for a Schottky diode, and negative in sign due to the fact that the diode's anode terminal connects to ground. Only toward the end of the "transistor off" interval do we see the distinctive difference in the discontinuous oscillograph, where that small voltage decays all the way to zero as the inductor current also decays to zero.

As previously mentioned, discontinuous converter circuit operation is not necessarily problematic. It does generally result in the creation of higher-order harmonic noise owing to the more abrupt changes in circuit voltage and current when the transistor switches state to re-energize the reactive component, but at the same time discontinuous-mode operation may also result in less heat dissipation from the converter circuit because it means periods of time where no current passes through some components.

3.12 Charge pumps

Inductor-based power converter circuits have a few weaknesses that must be acknowledged and understood. First, inductors in general possess more dissipative parasitic effects than capacitors: wire resistance, magnetic core losses, and magnetic coupling being chief among them. Unless made of superconducting wire, all inductors must necessarily possess Ohmic wire resistance, and this of course contributes to inefficiency by dissipating energy in the form of heat. Similarly energy is dissipated in the form of heat when ferromagnetic materials are magnetized and de-magnetized, also contributing to energy inefficiency. Lastly, any lines of magnetic flux not entirely contained within the inductor's core material are liable to "couple" with nearby conductors which not only diverts some of the energy from the converter circuit to adjacent circuits but also creates noise voltages in those circuits. Capacitors, by contrast, do suffer from equivalent series resistance (ESR) but generally not to the degree that inductors do, and they tend to contain their electric fields much more naturally than inductors contain their magnetic fields.

One way to avoid the inherent dissipative mechanisms of inductors in power converter circuitry is to eliminate the inductors entirely and find ways to perform the same conversion using only capacitors. Such circuits are called *charge pumps* because capacitors absorb and release electric charge when properly switched between source and load, much like a bucket may be filled with and emptied of water. Charge-pump circuits typically enjoy very high energy efficiency and broadcast comparatively little noise to adjacent circuits, but are more limited in current-sourcing capability than inductor-based power converters.

At the heart of many charge pump circuits is a component called a *flying capacitor*, which is a capacitor with both terminals connected to the rest of the circuit through switches (usually transistors functioning as on/off switches). A "flying" capacitor has no permanent connection to either the source or the load, but rather "flies" back and forth between both. An interesting application of a flying capacitor unrelated to power conversion is when a voltage signal having a large common-mode potential must be measured by a voltage-sensing device unable to tolerate that high common mode:

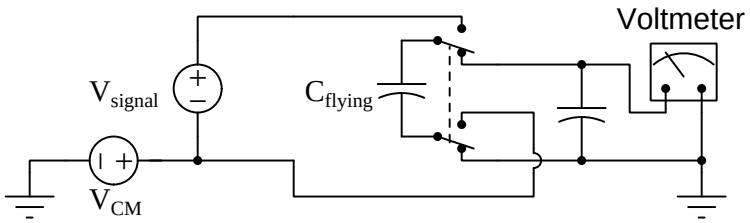

Every time the SPDT switches move to the "up" position the flying capacitor equalizes voltage with V_{signal}; every time the SPDT switches more to the "down position" the flying capacitor equalizes voltage with the voltmeter and its captive capacitor. Cycling the switches back and forth rapidly results in the voltmeter settling at a value equal to V_{signal} with no direct connection to V_{signal} and complete isolation from the common-mode voltage V_{CM}.

In the following circuit we see a flying capacitor used as an energy-transfer device rather than a signal-transfer device, in this case *inverting* the source's polarity to output the same amount of voltage to the load but with an opposite polarity with respect to ground. The four switches shown in this generic diagram represent four transistors driven by a pair of complementary square-wave signals such that only two transistors are ever conducting at any given time:

Inverting charge pump circuit

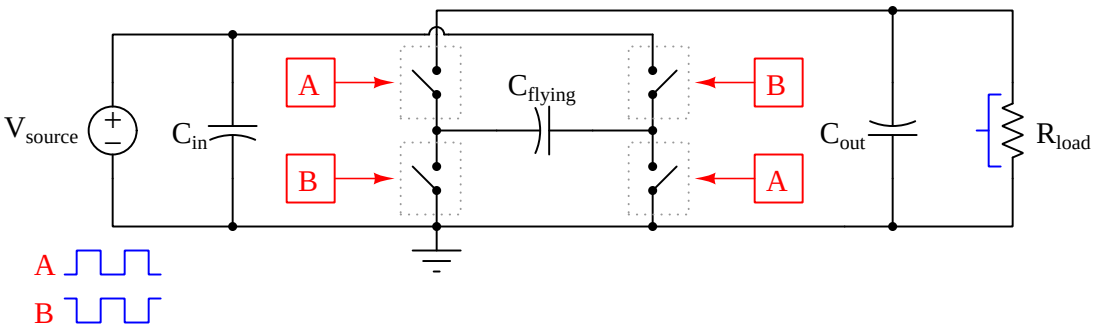

Applying the problem-solving strategy of voltage, current, and source/load annotation to this converter circuit in both transistor states, we see how the "flying capacitor" alternately acts as source and load:

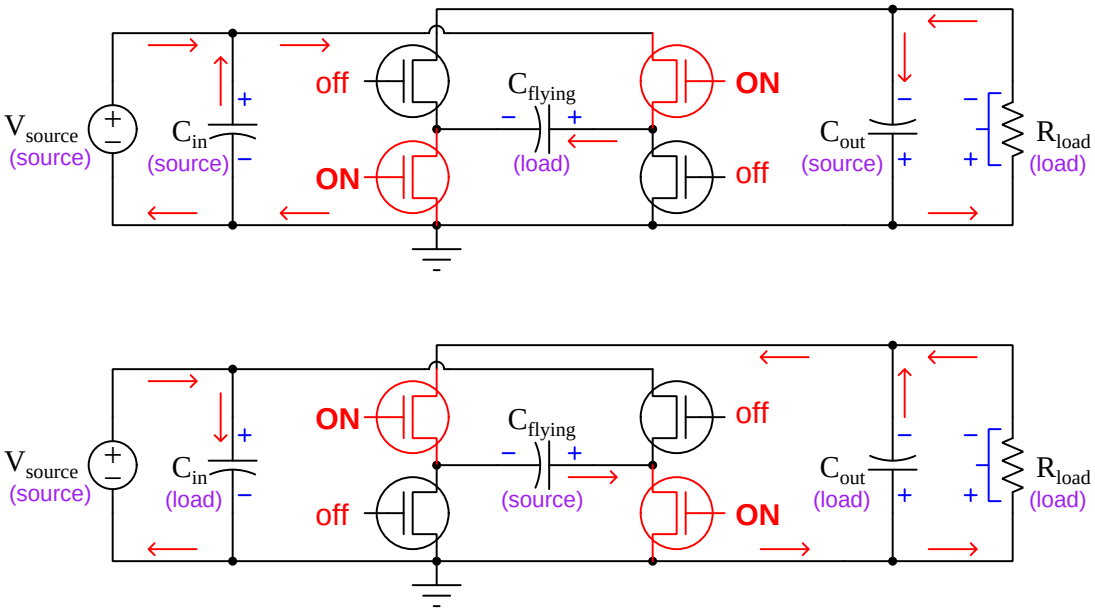

An interesting feature of the inverting charge pump circuit is that it works in reverse as well, so long as the switches (transistors) are capable of conducting bidirectional current. In the previous example we saw a negative-ground source deliver energy to a positive-ground load, but here we see the same circuit allow a positive-ground source deliver energy to a negative-ground load:

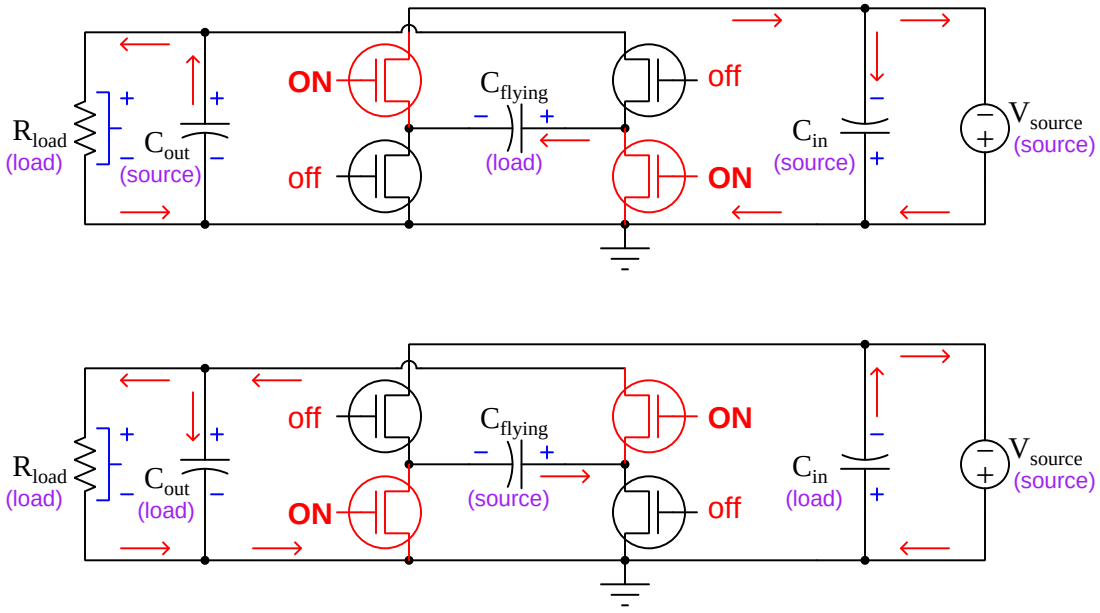

Flying-capacitor circuits are also useful for boosting voltage to a load, as shown in the following schematic diagram:

Boost charge pump circuit

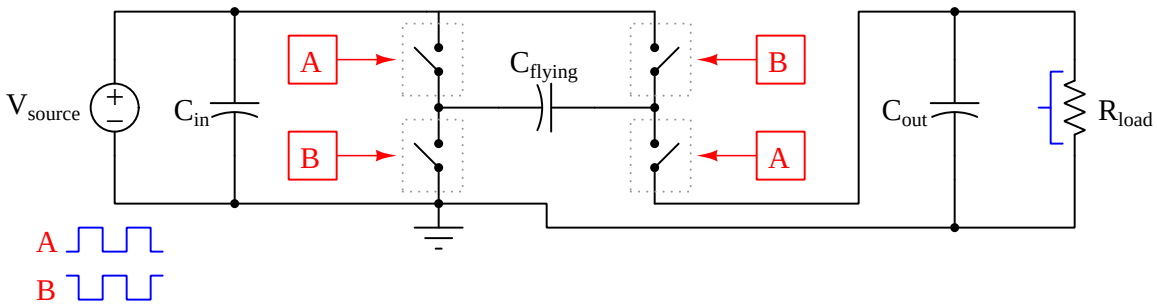

Applying the problem-solving strategy of voltage, current, and source/load annotation to this converter circuit in both transistor states, we see how the "flying capacitor" alternately acts as source and load to boost voltage from source to load:

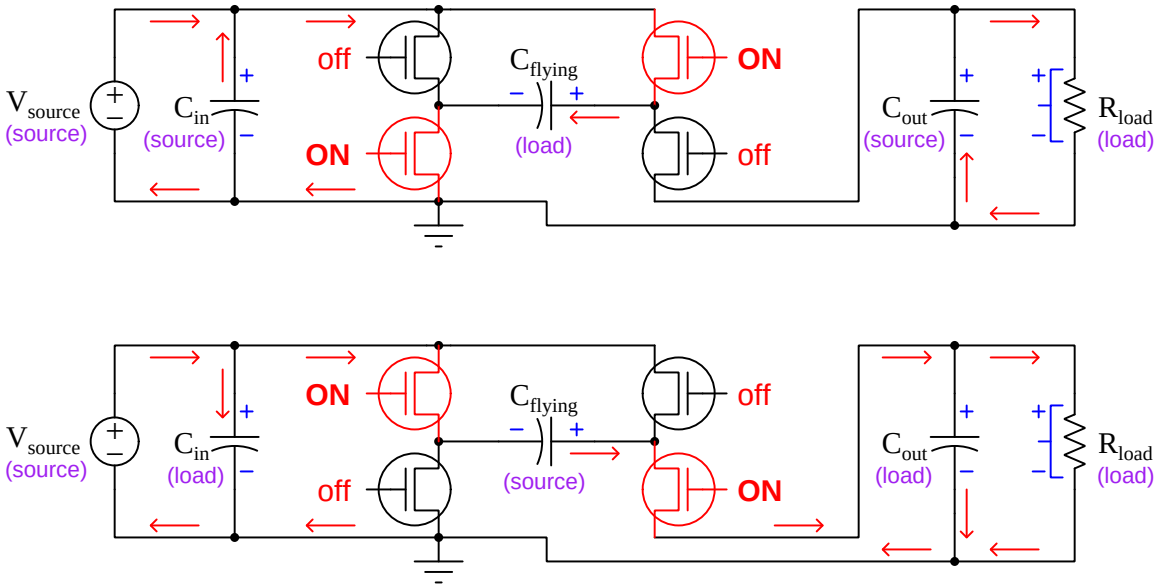

The "B" transistors turn on to connect the flying capacitor in parallel with the source to receive energy from the source, while the "A" transistors turn on to connect the flying capacitor *in series* with the source to present a higher voltage (at the same polarity) to the load. Voltage boost for this particular circuit is a 2:1 ratio of (unloaded) output voltage to source voltage, maximum.

As with the inverting charge pump circuit, we may reverse the locations of source and load so long as the transistors are capable of switching bidirectional current. If we do this, the boost charge pump circuit becomes a *buck* charge pump:

Buck charge pump circuit

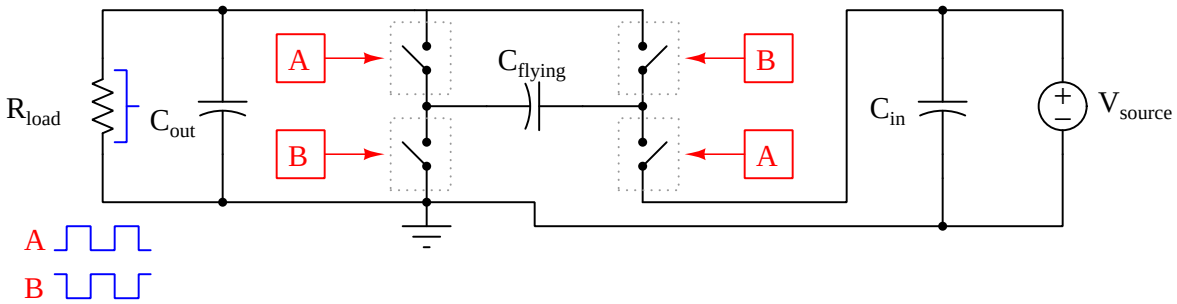

Applying the problem-solving strategy of voltage, current, and source/load annotation to this converter circuit in both transistor states, we see how the "flying capacitor" alternately acts as source and load to diminish voltage from source to load:

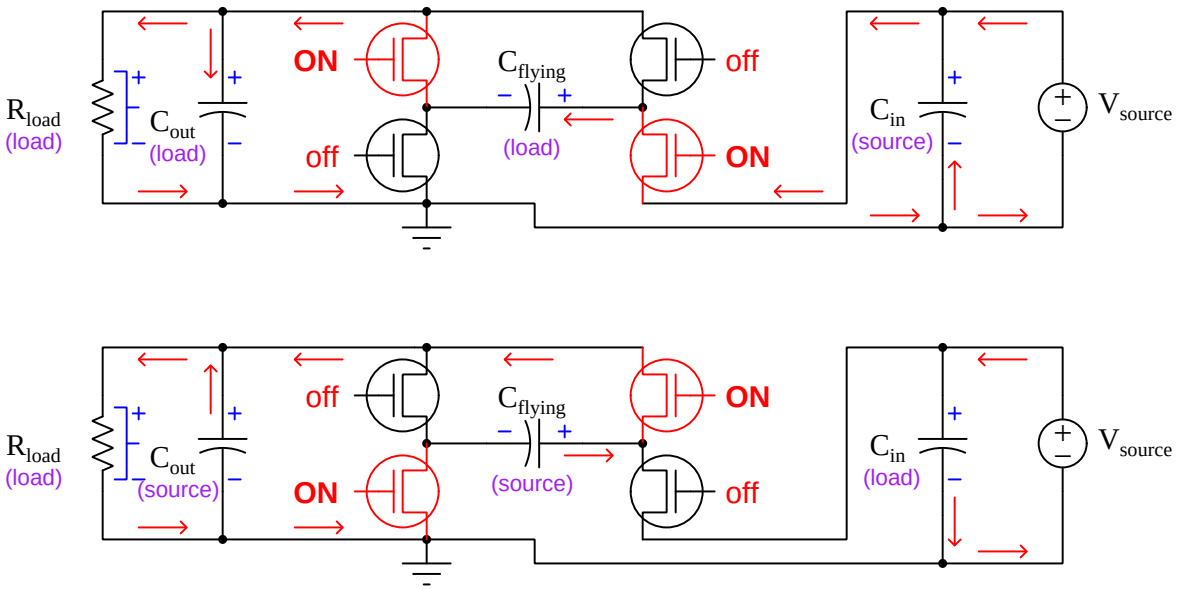

The "A" transistors turn on to connect the flying capacitor in series with $C_{out}||R_{load}$ parallel combination, each capacitor energizing to a fraction of the DC source's voltage. The "B" transistors turn on to connect the flying capacitor in parallel with C_{out} to deliver less voltage but more current to the load than the DC source provides, the ratios of each being 2:1 maximum (i.e. outputs half as much voltage and twice as much current as the source delivers).

3.13 Hybrid switching-linear converters

Switch-mode power converter circuits enjoy the advantages of high energy-conversion efficiency as well as adjustable step ratios, which is why they are so commonplace in modern electronic systems. Their principal disadvantage, however, is the generation of significant *electrical noise*, owing to the on/off pulsing of current as the transistor switching elements oscillate between conduction and non-conduction. These pulsations are non-sinusoidal in nature, which means they are equivalent to an infinite series of harmonic frequencies as predicted by Fourier's Theorem.

Frequencies typically employed for PWM switching of these transistors range in the tens of kiloHertz, and of course the harmonics' frequencies will be integer-multiples of the fundamental switching frequency. For example, if a DC-DC converter circuit uses a PWM frequency of 35 kHz, the expected harmonics will be 70 kHz for the second harmonic, 105 kHz for the third harmonic, 140 kHz for the fourth harmonic, 175 kHz for the fifth harmonic, and so on. The highest-order harmonics will have magnitudes corresponding to how steep the rate-of-rise and rate-of-fall is of the PWM pulse waveform. Since rapid turn-on and turn-off of a converter circuit's transistor(s) is an important factor in maximizing energy efficiency[22], it therefore follows that the most efficient converter circuits tend to create the richest harmonic spectra.

Both the fundamental switching frequency and the associated harmonic frequencies are all considered *noise* in a circuit that is ostensibly a *DC to DC* power converter. Many DC-powered circuits are adversely affected by such noise, and so the use of switch-mode power converters can be problematic. Clearly, a good solution to this problem is *low-pass filtering*, the purpose being to pass DC but block as much AC as possible so that none of this switching noise escapes the converter circuit to cause problems elsewhere. A common solution is to install LC filters before and/or after a switching converter as shown here:

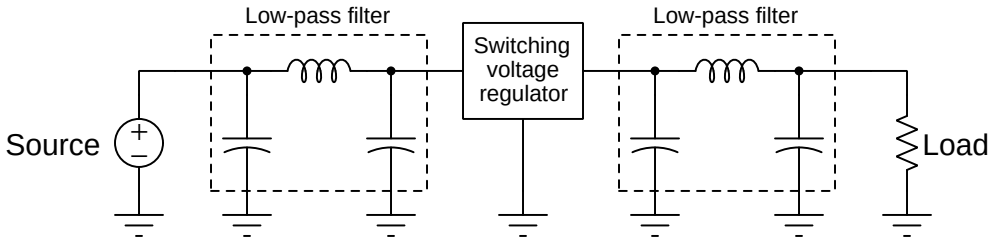

Filter networks containing resistors are not desirable for power applications such as this, because any resistance in a filter would dissipate energy in the form of heat, thereby reducing energy efficiency. The simplest non-resistive low-pass filter is simply a parallel-connected capacitor which relies on intrinsic circuit impedances to drop the undesired (i.e. AC) voltages, but combining series inductors with parallel capacitors makes for more effective filter networks with sharper roll-off rates which in turn more effectively screens out high-frequency AC noise. If superior low-pass filtering is desired, any number of L-C stages may be added to the filter networks until noise is attenuated down to acceptable levels.

[22] The more time that a transistor lingers in its "active" mode between cut-off and full-on saturation, the longer it behaves as a *resistance* dissipating energy in the form of heat. Therefore, in order to minimize this heat dissipation we must force the transistor's state to switch as rapidly as possible between on and off modes.

When striving for the lowest possible AC "noise" at the output of a switching converter, such as precision measurement applications where AC noise would complicate the accurate measurement of DC signals, it may be impractical to sufficiently eliminate that noise by low-pass filtering alone. In principle, a low-pass filter with enough stages of inductance and capacitance should minimize noise to any degree desired, but in practice this may not be an option due to size and/or weight constraints. Inductors tend to be large and heavy components, and the need for capacitors with minimal Equivalent Series Resistance (ESR) may eliminate electrolytic types which relegates the designer to using lower-loss capacitors that are physically larger than electrolytic capacitors of equal capacitance and voltage rating.

AC switching noise appears on the output of a DC-DC converter as high-frequency *ripple* riding on top of the DC voltage as viewed on an oscilloscope:

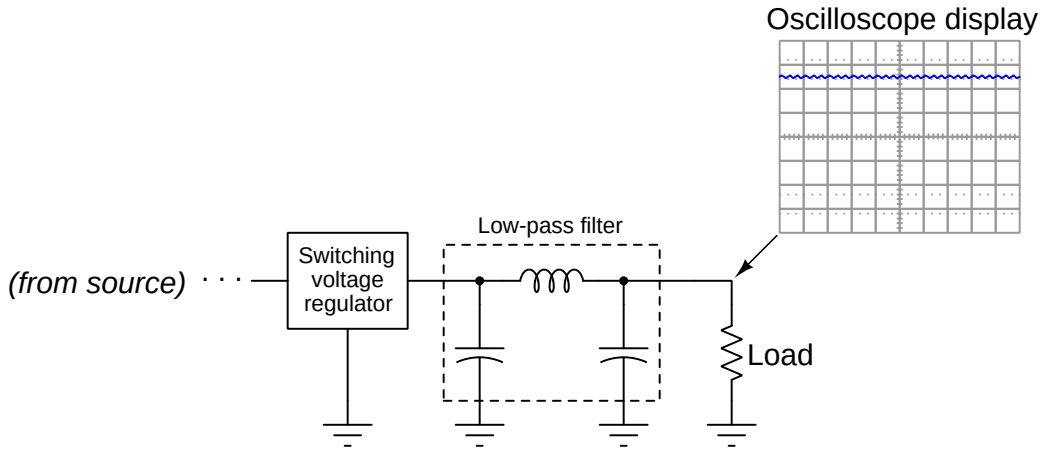

The frequency of this ripple voltage is a function of the PWM switching frequency within the converter, and being non-sinusoidal in shape it is equivalent to the PWM signal's fundamental frequency superimposed with all associated harmonic frequencies. A fortunate feature of this ripple's frequency typically being in the kiloHertz range is that it is much easier to filter out using low-pass filter networks than the ripple created by half-wave or full-wave rectification of AC line voltage (60 Hz and 120 Hz, respectively for 60 Hz AC power systems). Simply put, there is a larger span of frequency separating the desired signal (DC, 0 Hz) from the undesired signal (ripple) in switching converter circuits compared to brute-force power supply circuits. This means filter networks for switching converters may be smaller in size, lighter in weight, and use smaller-valued inductors and capacitors than equivalent filter networks in brute-force power supplies. However, the higher frequency of a switching converter's noise means it couples more readily to neighboring conductors via electric and magnetic fields, which means we generally cannot tolerate as much of it as we might if it were only 60 Hz or 120 Hz.

A compromise between energy efficiency and low load noise may be struck by inserting a linear voltage regulator between the switching circuit's output filter and the load, the voltage regulation point of the linear regulator being just below the DC voltage regulation point of the switch-mode regulator:

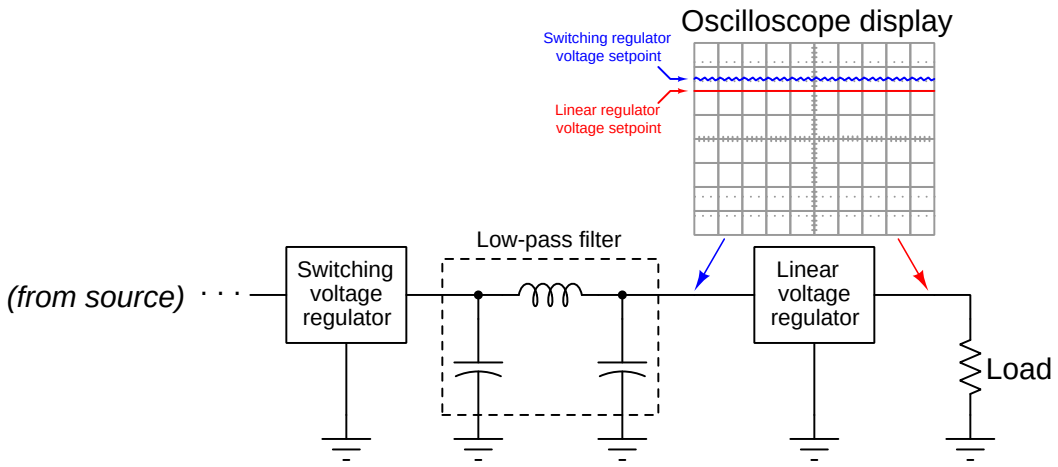

The difference between the switching regulator's DC voltage setpoint and the linear regulator's DC voltage setpoint must be greater than the *dropout voltage* of the linear regulator, to ensure there is enough "headroom" for that linear regulator to function properly, but that difference in voltage should be as small as possible to minimize power dissipated by the linear regulator, that dissipated power being equal to $I_{load}(V_{switching} - V_{linear})$.

What the linear voltage regulator does in this circuit is "chop off" the upper part of the voltage output by the switching regulator, which includes the ripple (noise), leaving nearly-pure DC for the load's use. There will, of course, still be a very small amount of ripple on the linear voltage regulator's output because linear voltage regulators do not have perfect *ripple rejection*, but what ripple will make it through the linear regulator will be quite minimal especially when compared against the ripple reduction offered simply by adding a few more stages of low-pass filtering after the switching regulator.

3.14 Integrated circuit power converters

Modern semiconductor manufacturing technology enables the integration of control and power-switching semiconductor devices on the same wafer or die of silicon, for convenient design of power converter circuits. Rather than piece together one's own design for a boost, buck, inverting, or other configuration of power converter, it is possible to purchase a single IC with all the control and (at least some) of the switching transistors inside, then connect to that IC a handful of external components such as inductors, capacitors, transformers, diodes, etc. Below we may view the "die shot" of one such controller IC, offered graciously by the creators of the website *zeptobars* (`https://zeptobars.com`) under the Creative Commons Attribution 3.0 Unported License.

Power Integrations model LNK364PN switch-mode controller, die size 2455 μm by 1027 μm
Image source: `https://s.zeptobars.com/LNK364PN-HD.jpg`

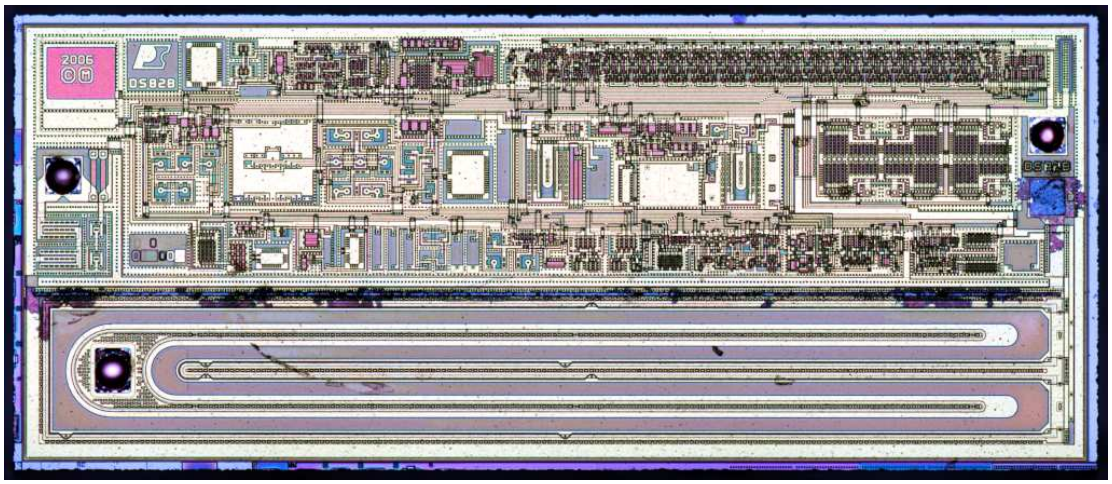

Note that large N-channel MOSFET occupying the entire width of the die at the bottom edge, rated for over 200 mA peak current and upwards of 700 Volts. Not all switch-mode voltage regulator ICs have the switching transistor built-in, electing instead to provide a digital output terminal for the PWM output signal which the user then connects to the gate of an external power MOSFET or base of a power BJT.

Shown here is a schematic diagram for a typical switch-mode "buck" voltage regulator with a built-in power transistor:

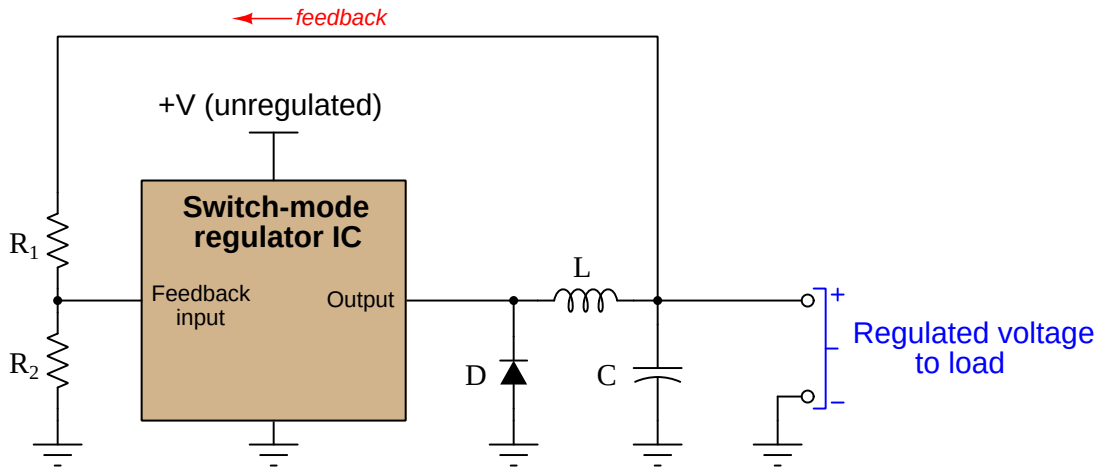

Note the external diode, inductor, and capacitor connected to this IC for the purpose of exchanging energy and achieving the desired power conversion from a higher DC voltage to a lower DC voltage in classic "buck" fashion. These components must be chosen on the basis of the desired load voltage and load current, and typically dwarf the IC itself in terms of physical size.

Note also the existence of a "feedback" conductor and voltage divider network (R_1 and R_2) providing the IC with information on how much voltage is seen by the load. In order for this IC to function as a true voltage *regulator* it must be able to sense load voltage and make any necessary corrections to its PWM duty cycle to overcome any deviations from the desired output voltage. This is a *negative feedback system*, very similar to typical operational amplifier circuits employing the same principle to achieve desired voltage gains. The voltage-division ration provided by these two resistors establishes the voltage regulation setpoint: if the divided voltage is a small fraction of load voltage, then the regulator will act to increase load voltage; if the divided voltage is a large fraction of the load voltage, then the regulator will act to decreases load voltage. If we dispense with the two resistors entirely and simply connect the feedback conductor between the load and the IC's feedback input terminal, it will regulate load voltage at the lowest possible setpoint value for the IC.

For example, consider an IC designed to regulate load voltage such that the potential at its feedback input terminal always rests at 1.23 Volts with reference to ground. If we were to choose equal-valued resistances for R_1 and R_2, the load voltage would end up being regulated at 2.46 Volts in order to maintain the feedback input potential at 1.23 Volts. If we were to choose 10 kΩ for R_1 and 4.7 kΩ for R_2, the load voltage regulation point would end up being 3.847 Volts, again in order to make the feedback input potential 1.23 Volts at the IC.

Chapter 4

Historical References

This chapter is where you will find references to historical texts and technologies related to the module's topic.

Readers may wonder why historical references might be included in any modern lesson on a subject. Why dwell on old ideas and obsolete technologies? One answer to this question is that the initial discoveries and early applications of scientific principles typically present those principles in forms that are unusually easy to grasp. Anyone who first discovers a new principle must necessarily do so from a perspective of ignorance (i.e. if you truly *discover* something yourself, it means you must have come to that discovery with no prior knowledge of it and no hints from others knowledgeable in it), and in so doing the discoverer lacks any hindsight or advantage that might have otherwise come from a more advanced perspective. Thus, discoverers are forced to think and express themselves in less-advanced terms, and this often makes their explanations more readily accessible to others who, like the discoverer, comes to this idea with no prior knowledge. Furthermore, early discoverers often faced the daunting challenge of explaining their new and complex ideas to a naturally skeptical scientific community, and this pressure incentivized clear and compelling communication. As James Clerk Maxwell eloquently stated in the Preface to his book *A Treatise on Electricity and Magnetism* written in 1873,

> It is of great advantage to the student of any subject to read the original memoirs on that subject, for science is always most completely assimilated when it is in its nascent state . . . `[page xi]`

Furthermore, grasping the historical context of technological discoveries is important for understanding how science intersects with culture and civilization, which is ever important because new discoveries and new applications of existing discoveries will always continue to impact our lives. One will often find themselves impressed by the ingenuity of previous generations, and by the high degree of refinement to which now-obsolete technologies were once raised. There is much to learn and much inspiration to be drawn from the technological past, and to the inquisitive mind these historical references are treasures waiting to be (re)-discovered.

4.1 Battery substations

Early developers of electric power systems faced a major technical problem: how to efficiently convey electricity from a generating station to low-voltage loads located many miles away. Operating generators and long-distance transmission lines all at the same low voltage as the loads required massive amounts of current in accordance with Joule's Law ($P = IV$), and this required very heavy-gauge (and very expensive!) metal transmission conductors. High-voltage generation was practical, and transmission of power at high voltages would allow for lower currents and smaller (cheaper) conductors, but high voltages were unsafe for distribution to residential loads. Somehow voltage and current would have to be *transformed* to allow the best of both scenarios: high voltage for long-distance transmission, but low voltage for loads distributed within a city.

Of course the invention of the electromagnetic transformer and AC electricity solved this problem by providing an efficient and reliable means of "stepping" voltage and current by desired proportions. Prior to these inventions, however, engineers had to be more creative. DC-to-DC motor/generator sets were one solution, but expensive and relatively unreliable. An entirely different approach to the problem of DC-DC conversion employed electrochemical batteries, as described by historian Thomas Hughes in his book *Networks of Power – Electrification in Western Society, 1880-1930* on page 85:

> Another ingenious attempt to solve the distribution-transmission problem involved storage batteries. Storage batteries were used to improve the load factor; in addition, batteries were utilized for d.c. transmission. Direct current taken from generators at over a thousand volts was transmitted over substantial distance (at reasonable cost because of the high voltage) and was then fed into batteries connected in series; each battery took an appropriate increment of the high voltage in a manner similar to the connecting of arc lamps in series in a high-voltage circuit. After charging, the batteries were disconnected from the generator and connected in parallel to obtain low voltage appropriate for the distribution network. `[page 85]`

An article in *Western Electrician* magazine from June 1890 describes this particular system – operated by the Chelsea Electricity Supply Company in England – in some detail, beginning on page 342. Note the frequent use of the acronym *e.m.f.* representing *electro-motive force* which was once a popular term for *voltage*:

> The system of distribution adopted by the Chelsea Electricity Supply company of England in which storage batteries are provided of sufficient capacity to supply the whole demand for current, says the London *Electrician*, is worthy of careful study and attention; because. although somewhat complicated at first sight, it is free from the objections which have been so strongly urged against other systems.
>
> The system consists essentially in the establishment of a number of sub-stations, which are distributed over the area of supply, each furnished with a large storage battery, which is divided into two parts, so that one may be charged while the other part is feeding the distributing network. At the time of greatest demand both halves of the battery are coupled up in parallel with each other to the distributing network. When being charged the half battery at each sub-station has its cells coupled in series throughout, and the several sub-stations are also connected in series by the charging cable so that a

high tension current is employed to effect the charging. The discharging halves of the batteries have their cells coupled in sets of 55 for 100-volt circuit, which are then coupled in parallel to the distributing network. Continuous current transformers are to be put down at each sub-station. In some tests which have already been made with this machine the efficiency was 75 to 80 per cent. The transformers absorb the full power of the engines and dynamos, and are intended to be coupled to the distributing feeders, parallel with the batteries, at the time of maximum demand during the winter months, they will be regulated so as to work at full load – that is, their e.m.f. will be slightly higher than that given by the batteries, the object of which is to avoid running the engines with light loads, and to maintain the batteries as fully charged as possible. It was stated at the meeting of the Chelsea company in February last that the coal consumption per board of trade unit delivered to the customers was, on the average, about 11 pounds, which is a much lower figure than that recorded with any other system of supply, and which may be accounted for by the fact that the plant is never running except on full load. One of the features of the system is that the number of charging dynamos in operation always equals the number of half batteries being charged. The dynamos and half batteries are connected all in series, and when a battery becomes fully charged and is cut out of the charging circuit the corresponding dynamo is stopped at the generating station. This will be effected automatically eventually; but at present it is done by the attendant, who sees by his voltmeter and ammeter the moment when a battery is cut out of circuit. The number of charging dynamos already fixed is five, four in operation, that being the number of sub-stations, and one spare. The dynamos are mounted on the same bed plate and coupled direct to Willans compound engines, flexible couplings are employed, and the dynamo spindle is insulated from the engine by the interposition of an ebonite disk and collars. The dynamos are of the Victoria type, giving an e.m.f. at their terminals of about 530 volts and 75 amperes, at a speed of 424 revolutions per minute. The field magnets of the charging dynamos are excited from a separate dynamo driven by a small engine. The exciting dynamo is compound-wound for constant difference of potential, so that no variation is caused by switching off the magnets of one or more of the charging dynamos. The armatures of the charging dynamos are coupled in series as before explained. A variable resistance is provided in connection with the field coils of each, so that the e.m.f. can be adjusted independently on all the – [page 342]

dynamos. When both halves of a battery become fully charged, and are consequently cut out of the circuit, the exciting circuit of a dynamo is broken, and the armature afterwards short-circuited. For a moment, then, the charging current flows through the inactive armature; but it will be seen that this has no objectionable result. The switches which control the exciting and armature circuits of each dynamo are ranged side by side; and a star wheel placed between them prevents the possibility of the armature being short-circuited before the exciting current has been broken, or the exciting current being put on before the armature has been thrown into the charging circuit.

In order to avoid a back current flowing from the batteries through the armatures of the charging machines consequent on any mistake on the part of the attendant, or a mishap, such as the sudden stoppage of the exciting current, a gravity switch. which breaks circuit if the current falls below a minimum value, is inserted in each of the charging

leads; and to avoid disastrous consequences, should these fail to act in event of a sudden reversal of the current, a safety device is introduced, Fig. 1. It consists of a magnetic cutout, set to act at 80 amperes, which is connected in the charging circuit; a fuse wire, which supports the heavy copper switch bar, is connected in parallel circuit with the magnetic cut-out.

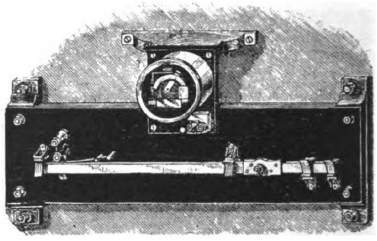

FIG. 1. CHELSEA COMPANY'S SYSTEM.

The four sub stations now in operation are situated at Draycott place, Clabon Mews, Egerton Mews, and Pavilion road. The last named is just being started. Each sub-station contains 454 E.P.S. cells of the size known as 31 L, i.e., 440 for supplying the lamps and 14 back e.m.f. cells. Each cell stands on a separate wood tray, which again rests on oil insulators. A light traveler is fitted up in each battery room, with the help of which any of the cells can be removed for inspection or repair.

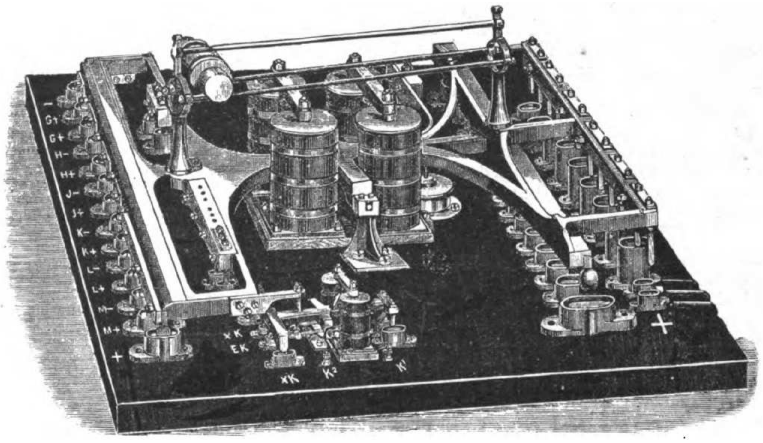

FIG. 2. CHELSEA COMPANY'S SYSTEM—TRANSPOSING SWITCH.

Upon commencing the day's charging the attendant at the generating station turns on a small switch for each substation, after which he switches the charging dynamos into circuit. Each small switch controls a relay at the substation connected with the operating magnet of a large transposing switch, Fig. 2, which is provided for each half battery. The function of the transposing switch is to connect the cells of the half battery either in series and to the charging main, or in parallel and to the distributing feeders, which it does by

means of a number of contacts which are mounted on a rocking framework. When the rocking frame is down on one side. the set of contacts on that side dip into the set of mercury cups and connect the battery in series for charging; and when it is down on the other side the other set of contacts dip into the other set of mercury cups and connect the battery in parallel for discharging. Suppose, then, that the attendant at the generating station has turned on the switch which sets the automatic apparatus in operation, the transposing switch of the first half battery is thereby thrown over from discharging to charging side, and the charging proceeds. In order to automatically cut the battery out of the charging circuit as soon as it is fully charged, a single negative plate in the half battery is selected, and a cap is arranged to catch the gas given off by the plate. The gas is led through a bubbling bottle into a small gasometer, which is adjusted so as to make a contact when the cells are properly charged, which operates a magnet on the transposing switch, causing it to rock over, and connect the cells in parallel onto the distributing feeders. The movement of this transposing switch makes a contact which causes the transposing switch of the second half battery to rock over from the discharging to the charging position, and the charging goes on until the gasometer, which in the meantime has been emptied by an electrically operated valve, has again risen high enough to make contact, when the second half battery is disconnected from the charging mains and thrown onto the distributing feeders parallel with the first half. To prevent the charging circuit being broken a carbon resistance is automatically connected across the terminals of the charging mains before the charged half of the battery is disconnected, and remains connected until the second half of the battery is connected in for charging.

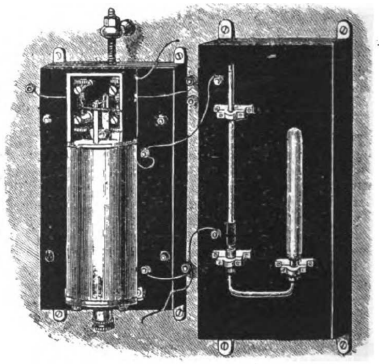

FIG. 3. CHELSEA COMPANY'S SYSTEM.

The regulation of the e.m.f. delivered to the distributing feeders is performed as follows by means of counter e.m.f. cells. By this arrangement the whole of the cells are uniformly discharged, a condition essential to the maintenance of the battery in good order. These regulating cells are automatically switched in and out of circuit. For this purpose a pair of pilot wires are brought from the point on the network where it is desired to maintain a constant e.m.f. These pilot wires are connected at the sub-station to the terminals of a governor, Fig. 3, which consists of a high resistance solenoid – ingeniously compensated for temperature variations – with a balanced core, to the top of which is attached a cross arm, carrying four contacts. When the e.m.f. is normal the cross arm remains in a central

position between four other contacts. If the cross arm is pushed upwards consequent on an increase in the e.m.f., the top contacts on the cross arm close a relay circuit, which operates a step-by-step device, Fig. 4, which at each step causes a back e.m.f. cell to be introduced into the circuit, through the intervention of the mercury switch, Fig. 6.

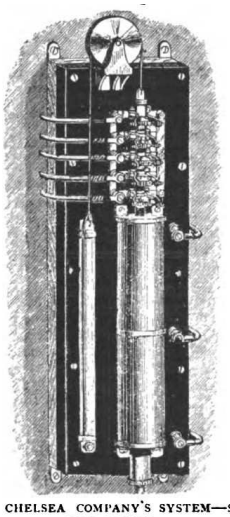

FIG. 4. CHELSEA COMPANY'S SYSTEM—STEP BY STEP DEVICE.

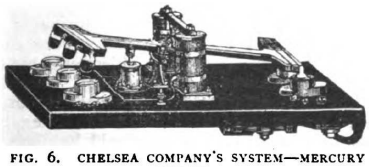

FIG. 6. CHELSEA COMPANY'S SYSTEM—MERCURY SWITCH.

When a sufficient number of cells has been introduced the core moves down breaking the contact, and the action ceases. Similarly, if the e.m.f. on the network falls, the core of the governor descends, the cross arm makes another pair of contacts, and the reverse action cutting out back e.m.f. cells takes place. The solenoid of the governor, through which the current from the pilot wire passes, is compensated against the effect of the surrounding temperature, which would otherwise cause considerable variation in the voltage maintained on the network. Instead of both ends of the operating solenoid being connected directly to the pilot wires, one end is connected through a carbon rod about 1-32 inch diameter, which introduces resistance. This carbon rod passes down the center of a small glass tube which is partly filled with mercury. The mercury is in connection with a closed chamber filled with alcohol, the expansion and contraction of which when the temperature of the room varies causes the mercury to rise or fall in the tube and to short-circuit more or less of the carbon rod. This device is adjusted to keep the total resistance always the same, and thus compensate for variations of temperature.

If the difference between the currents delivered by the two halves of the battery exceed 10 per cent. on the whole, back e.m.f cells are introduced into one half, or cut out of the other through the intervention of the current regulator, Fig. 5, which consists of two solenoids formed by cutting a spiral out of a piece of copper tube. The spirals are placed side by side, and the current from each half of the battery passes through them. The solenoids pull down either one side or the other of a balanced contact rim, which causes, through the medium of the relay, Fig. 5, and the step-by-step device, Fig. 4, a back e.m.f. cell to be switched out of the half-battery which is not supplying its share of current.

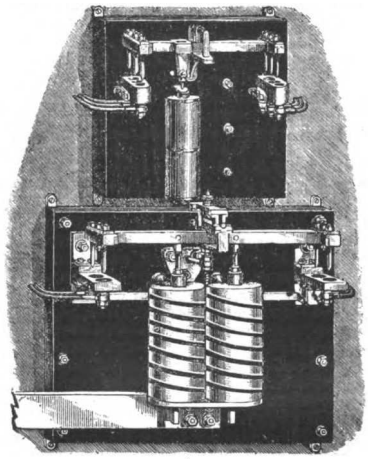

FIG. 5. CHELSEA COMPANY'S SYSTEM—CURRENT REGULATOR AND RELAY.

A large main switch and ammeter for the discharge circuit and an ammeter for the charging circuit is provided, and also a smaller ammeter having flexible leads and a wedge contact, which when forced between the proper contact plates diverts the current from any group of 55 cells through the ammeter.

The present capacity of the station and sub-stations is about 10,000; but the ultimate capacity, reckoning the continuous current transformers, will reach 30,000 8 to l0 candle-power lamps. The company find that 8 to 10 candle power lamps give very satisfactory lighting, and that if 16 candle power lamps are put in, the customer complains of the bill at the end of the quarter. The charge is made by meter – the Aron being generally employed.

The mains are laid underground on the Callender-Webber system. Their experience shows that when a separate channel is provided for each cable, as is the case in this system, long lengths of heavy cable can easily be drawn in or out. A cable, 2 in. in external diameter, 400 yards long, was recently drawn into a channel $2\frac{3}{8}$ in. diameter without difficulty, the cable being drawn 200 yards each way from a central opening. Thirteen men were employed in pulling the cable in. If a separate iron pipe were employed for each cable the cost would be too great; and experience has shown that even if a large

pipe is used the cables drag very heavily one on another if several are drawn into the same pipe. The charging mains are about two miles in length, and the charging e.m.f. reaches 2,200 volts. The boilers are the Babcock-Wilcox type of 100 horse-power nominal. The feed-water is supplied by a Worthington pump, and, in addition, injectors are fitted to each boiler in case of the pump failing from any cause.

In this connection are given some supply curves taken from two of the sub-stations, a glance at which will show that the current consumed during the daytime – though it would result in a loss if it necessitated constant running of plant – yet bears no mean proportion to the whole supply, and might make all the difference between a good dividend and none at all. [page 343]

4.2 Ćuk converter patent

Slobodan Ćuk is widely respected as a power electronics engineer, whose name is immortalized in a specific design of DC-DC converter circuit. The following excerpts come from his US Patent (number 4,184,197 filed 28 September 1977 and granted 15 January 1980):

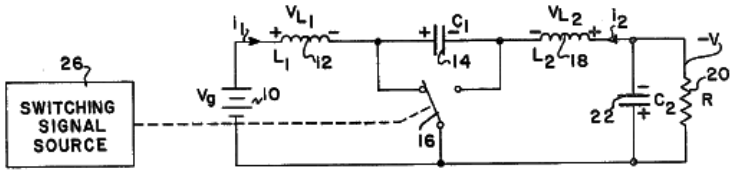

FIG. 5

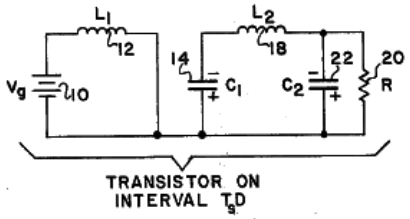

FIG. 5(1)

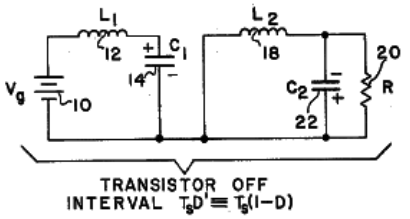

FIG. 5(2)

ORIGIN OF THE INVENTION

The invention described herein was made in the performance of work under a NASA contract and is subject to the provisions of Section 305 of the National Aeronautics and Space Act of 1958, Public Law 85-568 (72 Stat. 435; 42 USC 2457).

BACKGROUND OF THE INVENTION

This invention relates to dc-to-dc converters and more particularly to a switching converter having nonpulsating input and output current.

Several of the prior-art dc-to-dc converters are described here and some of their deficiencies are discussed. Later, after the present invention disclosed herein has been described, comparison of the invention with the prior-art converters will reveal the superior characteristics of the invention.

All of the three configurations discussed here (conventional buck-boost, cascade connected boost and buck, and double-ended, push-pull converters) have the general

DC conversion property, that is increase or decrease of both DC voltage and current. The same property is present in this invention as well, but with nonpulsating input and output current, and if desired isolation between the input and output can be obtained. The input or output current ripple may be reduced to zero.

OBJECTS AND SUMMARY OF THE INVENTION

An object of this invention is to provide a dc-to-dc converter having a maximum performance (high efficiency and low electromagnetic interference) for the minimum number and size of parts.

A further object is to provide a dc-to-dc converter having nonpulsating input and output current.

Another object is to provide a dc-to-dc converter with input or output current ripple reduced to zero.

Still another object is to provide a dc-to-dc converter which satisfies any of the foregoing objects with isolation between the current source and the load.

Yet another object is to provide a dc-to-dc converter which satisfies any of the foregoing objects with or without isolation and with an extended range of step-up or step-down of the input DC voltage.

These and other objects of the invention are achieved in a dc-to-dc converter with two inductances, a first one in series with the input source and a second one in series with the load, and energy transferring means having storage capabilities connected between the two inductances. The energy transferring means is comprised of storage capacitance in series with the two inductances and switching means for alternately connecting the junction between the first inductance and the storage capacitance to ground for the source, and connecting the junction between the storage capacitance and the second inductance to ground for the load. In some applications, both the load and the source may be referenced to the same ground. In those applications where the load and the source may not be referenced to a common ground, isolation between the source and the load is provided by dividing the storage capacitance between two capacitors with transformer coupling between the capacitors. In either case, the two inductances may be windings of a transformer designed for no ripple in either the input current or the output current. For zero output current ripple, the transformer is designed for the condition $n = k$, and for zero input current ripple, the transformer is designed for the condition $h = \frac{1}{k}$, where n is equal to the square root of the ratio of the self-inductances L_{11} and L_{22} of the input and output windings, respectively, and k is the coupling coefficient between the two windings. Alternatively, either of the two inductances may be a tapped inductor to form an autotransformer. In the input, the autotransformer provides an extended range of step-up of the input DC voltage, and in the output it provides an extended range of step-down of the input DC voltage.

The novel features that are considered characteristic of this invention are set forth with particularity in the appended claims. The invention will best be understood from the following description when read in connection with the accompanying drawings.

As typical with patents, the authors begin by describing “prior art” (i.e. the state-of-the-art

prior to the advent of their new invention). Here we read a description of the canonical buck-boost converter topology and its limitations, complete with illustrations:

> DESCRIPTION OF PRIOR ART CONVERTERS
>
> A conventional buck-boost converter is shown in FIG. a together with the corresponding pulsating input and output current waveforms, i_{in} and i_{out}. As seen in FIG. 1, the current drawn by this DC converter is not continuous and "clean" DC; rather it is pulsating, where the pulse amplitude in a higher current (and power) application may be very high. This invariably requires that an input filter F (usually a single section low-pass filter comprised of an inductance L_1 and capacitance C_1) be put in front to smooth out the substantial current pulse component at the switching frequency drawn from the line supply. In that way, electromagnetic interference (EMI) problems generated by the abrupt variation in energy flow (pulsating current) are reduced, and contamination of the environment by the undesired electromagnetic disturbances is alleviated. In operation, a transistor Q_1 is alternately turned on and off. When it is conducting, an inductance L_2 stores energy, and while it is off, the energy stored flows through a diode D_1 to a load R. A capacitance C_2 across the load smooths the output current to the load.
>
> While the input filter F reduces EMI problems at the input, the output current i_{out}, as seen in FIG. 1, is still pulsating, thus requiring a large size of output capacitance C_2 to smooth out that pulsating current ripple. In addition, relatively complicated "floating" switch drive circuitry is required for the transistor in this converter since its emitter is not referenced to ground.
>
> In summary, the important deficiencies of this prior art include the fact that both input and output currents are pulsating (with corresponding EMI problems), that larger size and weight are required, that complicated drive circuitry is required and, as demonstrated later, that significantly lower efficiency results.

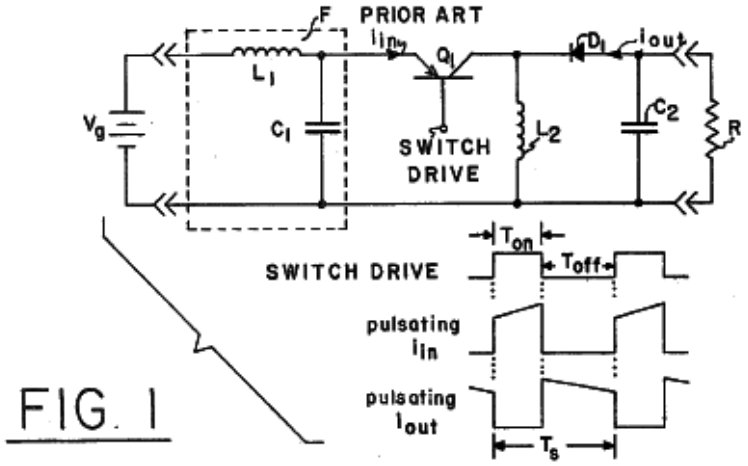

FIG. 1

Chapter 5

Derivations and Technical References

This chapter is where you will find mathematical derivations too detailed to include in the tutorial, and/or tables and other technical reference material.

5.1 Capacitor parameters and dielectric types

Capacitors come in many different types, the primary distinction being the type of substance used for the dielectric (i.e. the insulating layer between the capacitor plates). Rather than reproduce tables of dielectric material types and various performance ratings for each, here we will list several important parameters and for each of these cite the best and worst dielectric materials.

- **Rated voltage**: *the maximum continuous DC voltage across the capacitor based on the dielectric's breakdown limit*
- **Capacitance density**: *the amount of capacitance per unit volume, generally a desirable feature* – "supercapacitors" exploiting such effects as double-layer and pseudo-capacitance are the undisputed champions for this parameter. For conventional dielectrics, aluminum and tantalum electrolytic are best. Worst is vacuum or air.
- **Insulation resistance**: *the desirable ability of the dielectric to oppose "leakage current" through the capacitor* – polystyrene film is best; tantalum electrolytic is worst.
- **ESR and dissipation factor**: *mechanisms of energy dissipation within the capacitor, any of which are undesirable* – polystyrene, mica, and NP0 ceramic is best; tantalum electrolytic is worst.
- **Dielectric absorption**: *the undesirable tendency of a dielectric to maintain a DC bias after rapid discharging* – polypropylene and polystyrene are best; tantalum electrolytic is worst.
- **Stability**: *the degree to which a capacitor's capacitance value undesirably drifts over long periods of time* – NP0 ceramic is best; several dielectric materials tie for worst (High-K ceramic, reduced titanate, multi-layer glass with "V" characteristic).
- **Temperature coefficient**: *the degree to which a capacitor's capacitance value undesirably changes with temperature* – NP0 ceramic is best; High-K ceramic is worst.
- **Voltage coefficient**: *the degree to which a capacitor's capacitance value undesirably changes with applied voltage* – electrolytics are best; Class 2 ceramic and plastic-film are worst. Any dielectric material with *paraelectric* or *ferroelectric* properties (e.g. Class 2 ceramics such as barium titanate) exhibit this trait because their permittivity values vary significantly as the electric field increases and decreases in strength.
- **Polarization**: *whether or not the capacitor requires unidirectional voltage polarity* – all electrolytics are polarized; others are not.

The principal mechanism of energy dissipation within a capacitor is *dielectric loss*, where molecules within the dielectric material shift with a varying electric field. This voltage-induced molecular movement results in the dielectric becoming warm, thus converting electrical energy into thermal energy where it does not return to the circuit. Interestingly, this phenomenon is actually exploited in certain industrial processes such as the heating of plastics via high-voltage and high-frequency AC applied across the insulating material. In a capacitor, though, this effect is universally undesirable. Not only does dielectric loss constitute energy diverted from the circuit, but the heating will cause the capacitor's own capacitance to drift and may even destroy the device if severe enough.

The following table attempts to rank ten different dielectric types according to five important performance parameters. Please note that these rankings are based on some *worst-case* data I found, and that the performance metrics of specific capacitor models may likely vary from the values and orders shown below. As always, manufacture literature and/or empirical measurement are your best sources of information:

Rank	**Temp coeff.**	**Stability**	**Absorption**	**ESR**	**Insul. Resis.**
Best	NP0 = 0.3%	NP0 = 0.1%	PolyP = 0.05%	NP0 = 0.1%	PolyS = 10^6 MΩ
	PolyS = 1%	Mica = 0.1%	PolyS = 0.05%	Mica = 0.1%	PolyP = 10^5 MΩ
	Mica = 1.8%	PolyS = 2%	PolyC = 0.35%	PolyS = 0.1%	NP0 = 10^5 MΩ
	PolyC = 2%	PolyP = 3%	PolyE = 0.5%	PolyP = 0.35%	PolyC = 10^4 MΩ
	PolyP = 2.5%	PolyC = 5%	NP0 = 0.6%	PolyC = 1%	HiK = 10^3 MΩ
	Tant = 8%	PolyE = 10%	Mica = 0.7%	PolyE = 2%	PolyE = 10^3 MΩ
	Alum = 10%	Tant = 10%		HiK = 3%	RedT = 10^2 MΩ
	PolyE = 12%	Alum = 10%		Alum = 8%	Mica = 10^2 MΩ
	RedT = 30%	RedT = 20%	Alum	RedT = 10%	Alum = 10^2 MΩ
Worst	HiK = 80%	HiK = 20%	Tant	Tant = 24%	Tant = 10^1 MΩ

Abbreviations for dielectric types:

- NP0 = Multi-layer NP0 “Class 1” type ceramic
- HiK = Multi-layer High-K “Class 2” type ceramic
- RedT = Reduced titanate
- PolyE = Polyester film (“mylar”)
- PolyC = Polycarbonate film
- PolyP = Polypropylene film
- PolyS = Polystyrene film
- Tant = Solid tantalum electrolytic
- Alum = Aluminum electrolytic

Also note that conservative design principles may help minimize undesirable capacitor traits. For example, if the capacitor is known to have a large temperature coefficient, placing it in an area of the circuit board where you would expect temperature to be most stable will mitigate this shortcoming. Similarly, if the capacitor has a large voltage coefficient a good mitigation strategy is to over-rate the capacitor in terms of voltage (i.e. select a capacitor with a much higher voltage rating than the expected voltage for your application).

Multi-layer ceramic capacitors are extremely popular given their non-polarized performance making them suitable to reversing DC polarities as well as AC, and their ready availability in surface-mount device (SMD) packages. However, the two major classes (NP0 as "Class 1" and High-K as "Class 2") differ widely in performances. Class 1 ceramic capacitors are extremely stable in terms of their capacitance varying with age, with applied voltage, and with temperature. Class 2 ceramic capacitors sacrifice these desirable traits in favor of greater capacitance per physical volume. We typically find Class 1 ceramic capacitors used within signal filter and timing networks where their greater stability is worth their increased size and weight. Popular applications for Class 2 ceramic capacitors include DC power decoupling and radio-frequency interference (RFI) reduction where the precise amount of capacitance is less important than a certain minimum value of capacitance. Interestingly, Class 2 ceramic dielectrics also exhibit *piezoelectric* properties, meaning they produce voltage when physically strained and will physically strain when subjected to voltage. This means vibrations (including sound) will introduce AC noise voltage into the circuit at those same frequencies, and applied AC will correspondingly produce mechanical vibrations within the capacitor!

An important factor in ceramic dielectric types is temperature stability, i.e. how much capacitance is expected to vary as the device's temperature changes. Alphanumeric codes typically not appearing on the devices themselves specify the device's temperature dependence, the table below showing codes for Class 1 (NP0) ceramic dielectric materials:

Base temp. coefficient	**Temp. coefficient multiplier**	**Temp. coefficient tolerance**
C = 0.0 ppm/deg C	0 = −1	G = ± 30 ppm
B = 0.3 ppm/deg C	1 = −10	H = ± 60 ppm
L = 0.8 ppm/deg C	2 = −100	H = ± 120 ppm
A = 0.9 ppm/deg C	3 = −1000	K = ± 250 ppm
M = 1.0 ppm/deg C	4 = +1	L = ± 500 ppm
P = 1.5 ppm/deg C	6 = +10	M = ± 1000 ppm
R = 2.2 ppm/deg C	7 = +100	N = ± 2500 ppm
S = 3.3 ppm/deg C	8 = +1000	
T = 4.7 ppm/deg C		
V = 5.6 ppm/deg C		
U = 7.5 ppm/deg C		

C0G is the most common Class 1 dielectric type, with the best temperature stability of them all. Type C0G dielectric is so popular, in fact, that the designation is often used interchangeably with NP0 as a shorthand for "Class 1" dielectric.

Class 2 ceramic dielectric materials have a completely different alphanumeric coding scheme for representing temperature dependence, as shown in the table below. Rather than specify a temperature coefficient in parts per million, Class 2 dielectrics specify lower and upper bounds of their operating temperature range plus the percentage of rated capacitance value we may expect the capacitance to vary over that range:

Low temperature	**High temperature**	**Capacitance change over range**
X = −55 deg C	4 = +65 deg C	P = ±10%
Y = −30 deg C	5 = +85 deg C	R = ±15%
Z = +10 deg C	6 = +105 deg C	S = ±22%
	7 = +125 deg C	T = +22/−33%
	8 = +150 deg C	U = +22/−56%
	9 = +200 deg C	V = +22/−82%

Note how all three parameters together define temperature coefficient for a given ceramic capacitor. For example, a capacitor with X5R dielectric will drift more per degree of temperature variation than a capacitor with X7R dielectric despite the fact these two dielectric types share the same "R" letter code representing ±15%, because the X5R has a significantly narrower temperature range than the X7R for that same drift percentage. If a type-X7R capacitor drifts 15% in its capacitance as temperature rises from −55 degrees to +125 degrees, but a type-X5R capacitor drifts the same amount with less temperature rise (−55 degrees to just +85 degrees), it means the type-X5R capacitor is more sensitive to changes in temperature and therefore X7R would be a better choice than X5R for applications requiring a more stable capacitance value.

Common Class 2 dielectrics include X5R, X7R, and Y5V.

5.2 Inductor parameters and core types

Inductors come in many different types, the primary distinctions being the type of substance used for the magnetic core and the physical shape of that core. Several important performance parameters exist for inductors, listed below:

- **Insulation resistance**: *the desirable ability of the wire insulation to oppose "breakdown" resulting from high voltage applied across that insulation, either from the wire to the core material, from one wire to another (e.g. primary-to-secondary in a transformer), or from one turn of wire to an adjacent turn of wire*

- **DC resistance**: *Ohmic resistance of the wire used to make the winding(s) of the inductor, generally undesirable because this constitutes a mechanism of energy dissipation within the inductor*

- **Rated current**: *the maximum continuous DC current through the inductor based on wire heating from I^2R dissipation*

- **Saturation current**: *the amount of current at which point the core begins to magnetically saturate*[1]*, and its inductance correspondingly decreases, usually based on a specified percentage of inductance decrease compared to at zero current*

- **Distributed (parasitic) capacitance**: *wire insulation separating winding turns from each other and the winding from the core material (if electrically conductive) constitute parasitic capacitance in the inductor*

- **Self-resonance**: *the frequency at which an inductor will self-resonate based on its inductance reacting with its own parasitic capacitance*

- **Leakage inductance**: *the amount of inductance inherent to magnetic flux escaping the core (i.e. "leaking" flux), which may couple with adjacent conductors*

- **Stability**: *the degree to which an inductor's value undesirably drifts over long periods of time*

- **Temperature coefficient**: *the degree to which an inductor's value undesirably changes with temperature*

- **Rated temperature**: *the maximum operating temperature at which an inductor will still perform within specification, limited by insulation type, thermal expansion coefficients of the wire and core materials, and the Curie temperature*[2] *of the core material*

[1] Ferromagnetic materials assist in the development of magnetic fields as their internal magnetic "domains" align one-by-one with any externally-applied field such as that from an energized coil of wire. If we keep increasing the external magnetic field, we will reach a point where all the ferromagnetic core material's domains are fully aligned with that field, at which point the core can offer no further assistance to increasing field strength. Here the core is said to be *magnetically saturated*, like a sponge fully saturated with water. Any further increases in magnetic field strength (e.g. increasing coil current beyond the saturation point) will proceed unaided by the core.

[2] The Curie temperature is that temperature beyond which a ferromagnetic material begins to lose its magnetic properties.

Multiple energy-dissipation mechanisms exist within inductors, and may be broadly divided into two categories: "copper" (wire) losses and "iron" (core) losses. All wires made of non-superconducting materials exhibit some electrical resistance, and this directly converts electrical energy into thermal energy ($P = I^2R$) where it will never return to the circuit. At very high frequencies where the *skin effect* causes AC current to avoid the center of the conductor and only travel near the surface (skin), the wire's effective cross-sectional area decreases and this leads to greater I^2R losses. Ferromagnetic core materials experience molecular-level motion[3] when the magnetic field varies in intensity, and this represents another way for electrical energy to become converted into heat and never return to the circuit. Magnetic materials dissipate energy in this way proportional to the amount of hysteresis they exhibit: magnetically "soft" materials having narrow B-H hysteresis curves are more energy-efficient as inductor and transformer cores than magnetically "hard" materials having wide B-H hysteresis curves. Furthermore, if the core material happens to be electrically conductive there will be some amount of electric current induced within the core itself as inductor current varies, as a result of electromagnetic induction. These *eddy currents* encounter electrical resistance in the core material, also resulting in I^2R losses.

A variety of core shapes exist for different sizes and purposes of inductor:

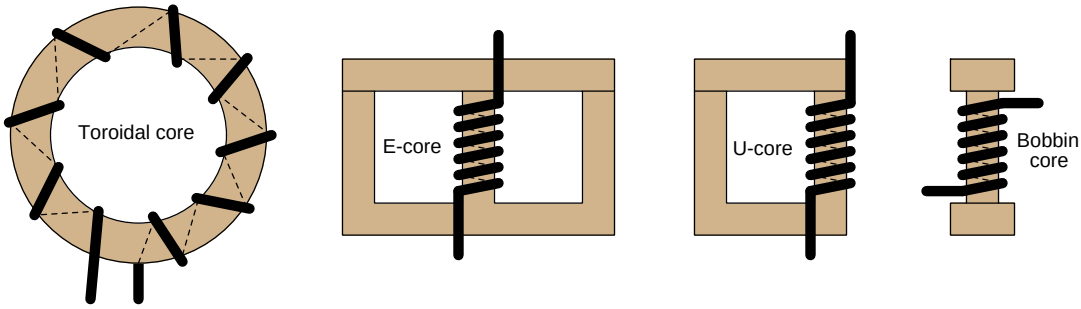

Toroidal cores are exceptionally good at confining the magnetic field to the inductor (i.e. very low "leakage" flux) but are more expensive to manufacture.

Inductor core materials exist to enhance the amount of magnetic flux resulting from electric current passing through the winding(s), as well as to "contain" the magnetic flux so that it remains inside the inductor case and does not link with any conductors outside of the inductor. Some inductors have no solid core material, and are called *air-core* inductors. These have very low inductance values and suffer from unwanted "coupling" with adjacent conductors but enjoy immunity from magnetic saturation. Most inductors use some form of *ferromagnetic* material such as iron or iron alloy.

[3] This effect when viewed on a macroscopic scale is called *magnetostriction*, and accounts for why ferromagnetic inductor and transformer cores may "hum" when their respective windings are energized by AC in the audio-frequency range.

Most inductor designs attempt to minimize energy losses in the magnetic core, one strategy for this being to minimize magnetic hysteresis losses. Hysteresis losses occur when the magnetic domains inside the material have a high degree of *remanence*, necessitating work to be done when magnetizing, unmagnetizing, and re-magnetizing the core material as inductor current varies and/or reverses direction. The amount of energy dissipated in this manner by a magnetic core is proportional to the area enclosed by the B-H curve:

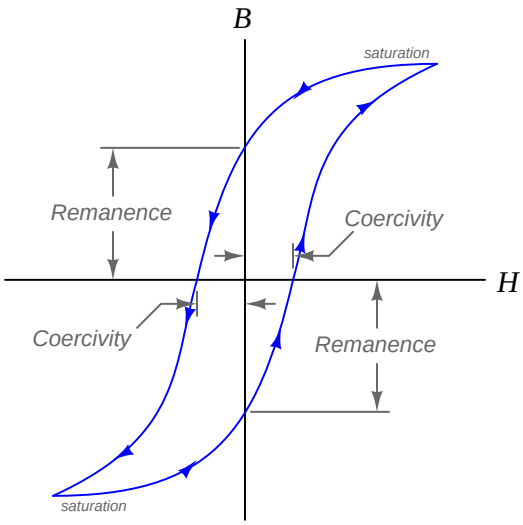

Choosing a magnetic material with a naturally narrow B-H curve helps minimize hysteresis losses. So does "over-building" the core to have greater cross-sectional area than is absolutely necessary, in an effort to spread out the magnetic flux over more material and therefore decrease flux density (B) so the core remains as far from saturation as possible.

Another strategy for minimizing energy losses in the core focuses on the mechanism of *eddy currents* which are electric currents induced in the core material by way of electromagnetic induction. This is a problem when the core material happens to be electrically conductive as well as ferromagnetic (e.g. iron is a good example of this), as these circulating electric currents do no useful work yet still encounter electrical resistance in the core material causing heat dissipation which is energy extracted from the circuit never to return to it. An effective mitigation is to build the core in such a way that it electrically interrupts these currents by not providing a continuous pathway for them to circulate. If the core material is iron, constructing the core as a stack of thin iron sheets coated in (insulating) enamel accomplishes this goal, and we call this a *laminated iron core*. Another strategy is to use powdered iron rather than solid iron. Yet another is to use a material that is ferromagnetic but electrically insulating, for example *ferrite* which is a chemical oxide of iron rather than being metallic iron.

For the end-user rather than the designer of manufactured inductors, an effective strategy for minimizing energy losses is to simply over-rate the inductor in terms of current (i.e. select an inductor with a much higher current rating than the expected current for your application). Using an inductor with a much higher current rating than your expected design current means that inductor will have larger-gauge wire windings, as well as an over-sized magnetic core capable of handling much stronger magnetic fields than you will ever impress upon it. Over-sized windings minimize resistive

energy losses, while the over-sized core effectively shortens and narrows the B-H curve which in turn minimizes hysteresis losses.

5.3 Capacitor parasitics

5.3.1 Model of a real capacitor

An ideal capacitor exhibits only capacitance, with no inductance, resistance, or other characteristics to interfere. Real capacitors exhibit all these phenomena to varying degrees, and we collectively refer to these undesirable traits as *parasitic effects*. The following diagram models some of the parasitic effects observed in real capacitors:

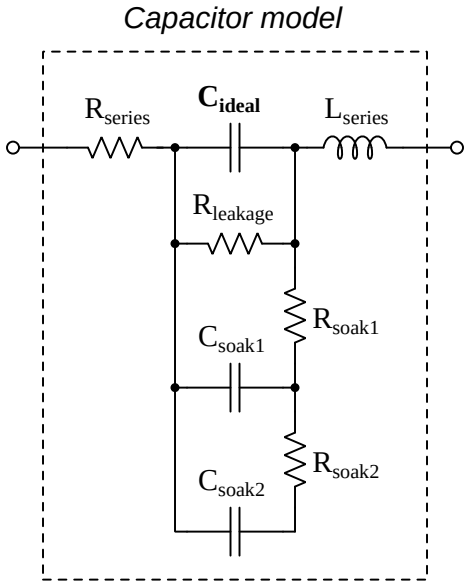

In addition to the capacitance the capacitor is supposed to exhibit (C_{ideal}), the capacitor also has parasitic resistance (R_{series}, also known as *Equivalent Series Resistance*, or *ESR*), parasitic inductance (L_{series}), and additional energy storage in the form of *soakage* (also known as *dielectric absorption*) whereby the dielectric substance itself absorbs and releases energy after relatively long periods of time compared to the main (ideal) capacitance.

Some of these parasitic effects – such as leakage resistance and soakage – affect the capacitor's performance in DC applications. Most of the other parasitic effects cause problems in AC and pulsed applications. For example, the effective series capacitor-inductor combination formed by C_{ideal} and L_{series} will cause *resonance* to occur at a particular AC frequency, resulting in much less reactance at that frequency than what would be predicted by the capacitive reactance formula $X_C = \frac{1}{2\pi fC}$.

Next we will explore common mechanisms for each of these effects.

5.3.2 Parasitic resistance in capacitors

Wire resistance, of course, plays a part in this parasitic effect, but this is generally very small due to the short conductor lengths we typically see with capacitors. More significant is *dielectric losses* – energy dissipation caused by the stressing and relaxation of dipoles within the dielectric material – which act like resistance because energy ends up leaving the component (in the form of heat) and not returning to the circuit. Electrolytic capacitors have an additional source of parasitic resistance, in the form of the electrolytic gel substance used to make electrical contact from the metal-foil "plate" to the surface of the dielectric layer.

Another form of parasitic resistance within a capacitor behaves like a resistor connected in parallel with the ideal capacitance ($R_{leakage}$), resulting from the dielectric not being a perfectly insulating medium. This parasitic characteristic results in a small current passing through the capacitor even when the voltage across the capacitor is steady (i.e. $\frac{dV}{dt}$ is zero).

5.3.3 Parasitic inductance in capacitors

Any time a magnetic field forms around a current-carrying conductor, energy is stored in that magnetic field. We call this magnetic-based energy-storing capability *inductance*, and of course all capacitors must have some inductance due to the wire leads serving as connection points to the capacitor's metal plates. Much of a capacitor's parasitic inductance may be minimized by maintaining short lead lengths as it attaches to a printed-circuit board.

Parasitic inductance is a problem for capacitors in AC applications because inductive reactance (X_L) tends to cancel out capacitive reactance (X_C). If we plot the impedance of a capacitor as a function of frequency, we would expect an ideal capacitor to manifest a straight-line descent on a logarithmic plot. However, what we see is that at a certain frequency the series parasitic inductance resonates with the capacitance leaving only parasitic resistance (ESR), and then past that frequency the inductive effects overshadow the capacitance:

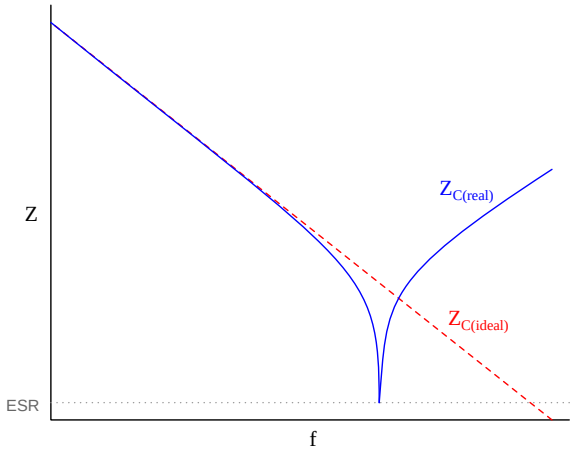

5.3.4 Other parasitic effects in capacitors

Soakage is an interesting effect resulting from dipole relaxation within the dielectric material itself, and may be modeled (as shown) by a series of resistor-capacitor subnetworks. This effect is especially prominent in aluminum electrolytic capacitors, and may be easily demonstrated by discharging a capacitor (by briefly connecting a shorting wire across the capacitor's terminals) and then monitoring the capacitor's DC voltage slowly "recover" with no connection to an external source.

5.4 Inductor parasitics

5.4.1 Model of a real inductor

An ideal inductor exhibits only inductance, with no capacitance, resistance, or other characteristics to interfere. Real inductors exhibit all these phenomena to varying degrees, and we collectively refer to these undesirable traits as *parasitic effects*. The following diagram models some of the parasitic effects observed in real inductors:

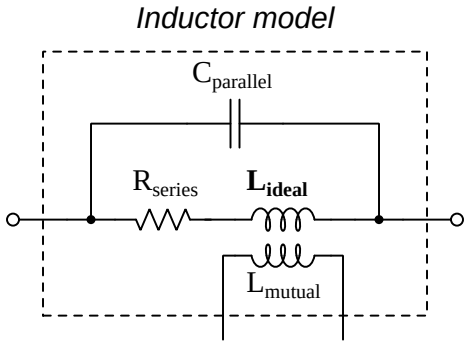

In addition to the inductance the inductor is supposed to exhibit (L_{ideal}), the inductor also has parasitic resistance (R_{series}, also known as *Equivalent Series Resistance*, or *ESR*), parasitic capacitance ($C_{parallel}$), and mutual inductance (L_{mutual}) with nearby wires and components.

Some of these parasitic effects – such as equivalent series resistance – affect the inductor's performance in DC applications. Most of the other parasitic effects cause problems in AC and pulsed applications. For example, the effective inductor-capacitor "tank circuit" formed by L_{ideal} and $C_{parallel}$ will cause *resonance* to occur at a particular AC frequency, resulting in much more reactance at that frequency than what would be predicted by the inductive reactance formula $X_L = 2\pi fL$.

Next we will explore common mechanisms for each of these effects.

5.4.2 Parasitic resistance in inductors

Wire resistance plays a dominant role in this parasitic effect due to the typically long lengths of wire necessary to wind the coil that forms most inductors. Wire resistance is not the only dissipative mechanism at work, though. Other losses include magnetic hysteresis of the iron core material as well as *eddy currents* induced in the iron core. An "eddy" current is a circulating electric current induced within the iron core of an inductor, made possible by the fact that iron is an electrically-conductive material as well as being ferromagnetic. These circulating currents do no useful work, and dissipate energy in the form of heating the iron. They may be minimized by forming the iron core from pieces of iron that are electrically insulated from one another, e.g. forming the iron core from *laminated sheets* or *powdered particles* of iron where each sheet or particle is electrically insulated from the next by a layer of non-conductive material on its outer surface.

The series resistance of an inductor is always frequency-dependent. In DC conditions (i.e. frequency of zero Hertz) there will be the basic wire resistance of the coil at play. As frequency increases from zero, however, both the magnetic core losses from hysteresis and eddy currents also increase which add to the DC resistance to form a larger ESR. At extremely high frequencies the *skin effect*[4] further adds to the inductor's ESR.

[4] At high frequencies, electric current travels more toward the outer surface of a conductor rather than through the conductor's entire cross-section, effectively decreasing the conductor's cross-sectional area (gauge) as frequency rises.

5.4.3 Parasitic capacitance in inductors

Any time an electric field forms between two conductors, energy is stored in that electric field. We call this electric-based energy-storing capability *capacitance*, and of course all inductors must have some capacitance due to the insulating media between wire leads as well as between adjacent turns of wire within the coil (and between the wire turns and the iron core).

Parasitic capacitance is a problem for inductors in AC applications because capacitive reactance (X_C) tends to cancel out inductive reactance (X_L). If we plot the impedance of an inductor as a function of frequency, we would expect an ideal inductor to manifest a straight-line ascent on a logarithmic plot. However, what we see is that at a certain frequency the parallel parasitic capacitance resonates with the inductance to create a nearly-infinite impedance, and then past that frequency the capacitive effects overshadow the inductance:

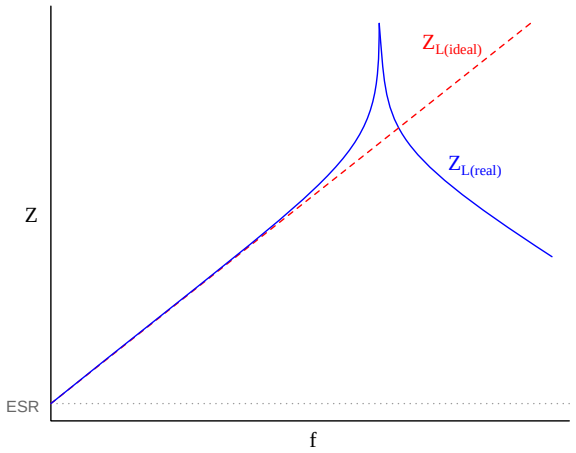

Precious little may be done to eliminate parasitic capacitance within any inductor, whereas parasitic inductance is fairly easy to minimize within a capacitor. This explains why when faced with an equivalent choice between a circuit design using capacitors and a circuit design using inductors, capacitors nearly always win. Simply put, it is easier to make a nearly-ideal capacitor than it is to make a nearly-ideal inductor.

This also explains why the self-resonant frequency of most inductors is much lower than the self-resonant frequency of most capacitors: all other factors being equal, an inductor will have more parasitic capacitance in it than an equivalent capacitor will have parasitic inductance within it, making the LC product greater for the inductor than for the capacitor.

5.4.4 Other parasitic effects in inductors

Mutual inductance occurs whenever adjacent conductors' magnetic fields link with one another, which is difficult to avoid especially in physically dense circuit layouts. This parasitic effect may be minimized by proper placement of inductive components (e.g. keeping them spaced as far apart from each other as possible, orienting their axes perpendicular to each other rather than parallel) as well as by core designs with strong magnetic field containment (e.g. toroidal cores contain their magnetic fields better than rectangular cores).

5.4.5 Inductor testing

The following list shows typical tests for inductors at the time of manufacture:

- **DC winding resistance** – the simplest of these to measure, using a multimeter
- **Winding inductance/impedance** – performed by exciting the winding with an AC source and comparing voltage/current, typically using either an LCR meter or an impedance analyzer
- **Equivalent winding resistance** – based on the ratio of AC voltage to AC current (only the portions of each in-phase with each other) representing winding resistance as well as energy losses in the magnetic core
- **Core saturation** – performed by exciting the winding with a sinusoidal AC current and measuring self-induced winding voltage until a non-sinusoidal waveshape is detected, or more crudely by determining when the self-induced voltage no longer increases proportionately to increased excitation current
- **Insulation resistance** – performed with a "Hi-Pot" (high-potential) insulation tester, generally connecting one terminal to either or both ends of the winding and the other tester terminal to the conductive core material, this infers insulation resistance by applying (typically) several thousand Volts DC while measuring the small current passing through the stressed insulation

A more thorough battery of tests used during transformer design development might include the following:

- **BH loop** – performed by exciting the winding with an AC voltage and measuring the self-induced voltage across it, the phasor portion of that induced voltage shifted by 90 degrees from the current waveform representing the rate-of-change of magnetic flux which may be mathematically integrated to infer the amount of magnetic flux, those two variables being plotted using an oscilloscope's X-Y mode much like "curve tracing" for a semiconductor component ; often performed at different temperatures to fully characterize the magnetic properties of the core material
- **High-voltage impulse** – similar to an insulation resistance test, this uses a high-voltage DC pulse applied across the winding insulation to stress it
- **Volt-second** – performed by applying a DC voltage pulse to the winding while measuring current through it, the ideal response being a linear-ramping current whose magnitude is proportional to the Volt-second product of the applied pulse
- **Overload** – subjecting a winding to dramatic overcurrent to thermally and mechanically stress the inductor as it might be in an overloaded power system scenario
- **External magnetic field rejection** – subjecting the entire inductor to an external (alternating) magnetic field and measuring any voltage induced across any winding, usually tested with the magnetic field applied across different physical axes
- **Vibration** – not an electrical test per se, but a useful stressor when followed by repeated electrical tests to ensure electrical/magnetic properties have not substantially changed following the vibration

Chapter 6

Animations

Some concepts are much easier to grasp when seen in *action*. A simple yet effective form of animation suitable to an electronic document such as this is a "flip-book" animation where a set of pages in the document show successive frames of a simple animation. Such "flip-book" animations are designed to be viewed by paging forward (and/or back) with the document-reading software application, watching it frame-by-frame. Unlike video which may be difficult to pause at certain moments, "flip-book" animations lend themselves very well to individual frame viewing.

6.1 VFD transistor switching sequence

The following animation shows how the "rotating" magnetic field of a three-phase AC induction motor is produced by the interaction of three stator winding sets energized with different phases (A, B, and C) of a three-phase AC power source. A red arrow shows the direction of the *resultant* magnetic field created by the interaction of the three winding sets.

A VFD's power circuitry consists of a *rectifier* section to convert AC into DC, a *filter* section to "smooth" the DC so that it is more steady than it would be otherwise emerging from the rectifier, and finally an *inverter* section where the filtered DC is converted into three-phase AC at the desired voltage and frequency. In this animation we show the inverter stage's six transistors turning on and off as necessary to create the rotating magnetic field vector in the motor, but what we do not show is the high-speed pulsing of those transistors within each of their "on" periods necessary to synthesize a *sinusoidal* output current.

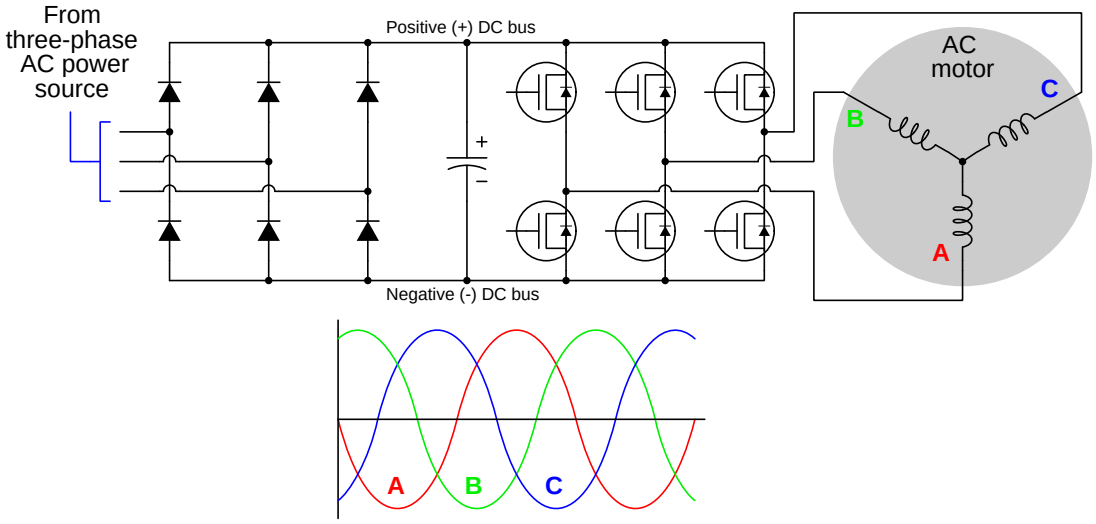
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
+
–
AC
motor
A
B
C
A
B
C

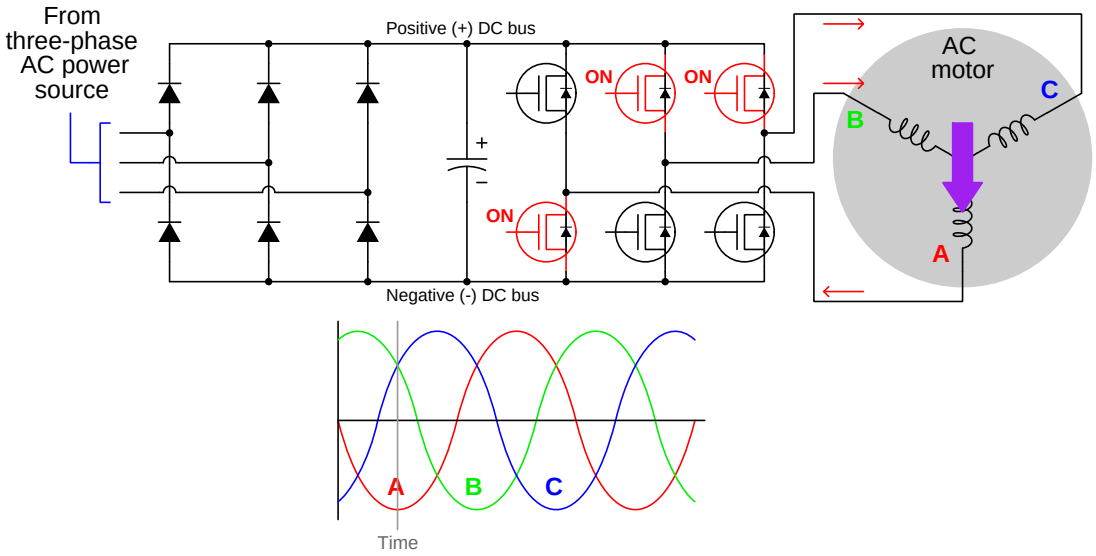
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
AC
motor
A
B
C
A
B
C
Time

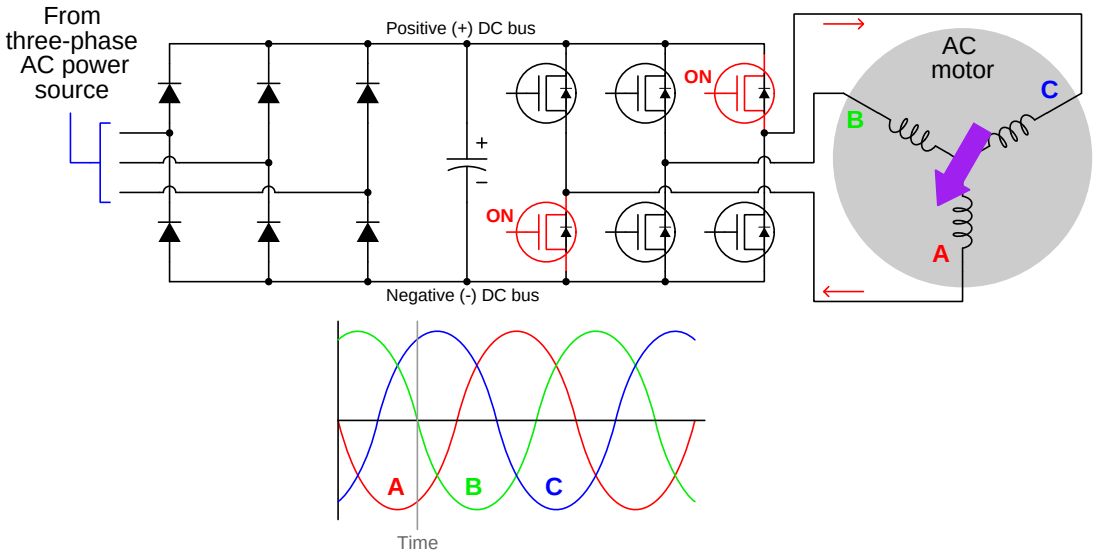
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
+
−
AC
motor
A
B
C
A
B
C
Time

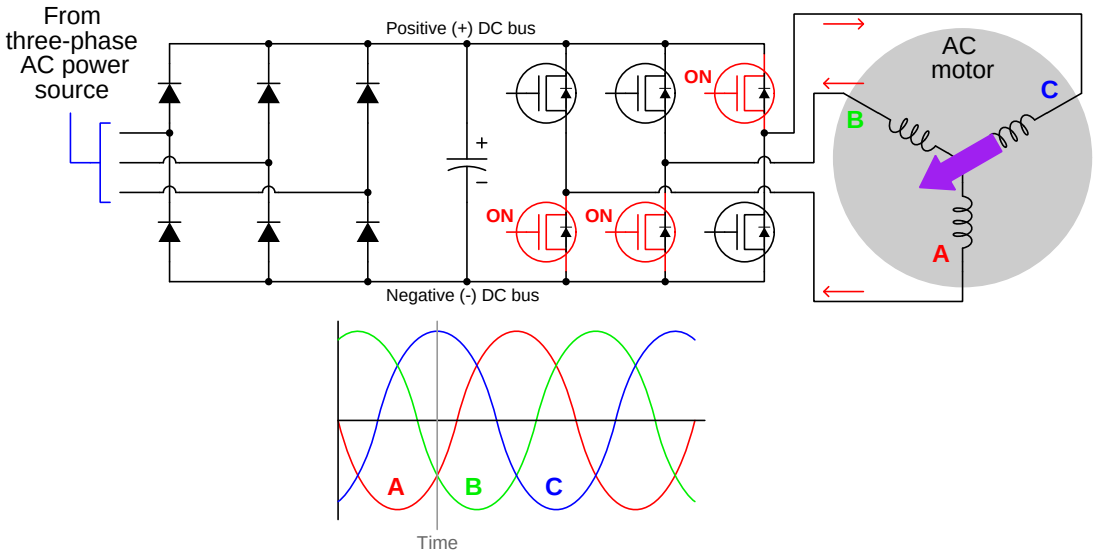
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
AC
motor
A
B
C
A
B
C
Time

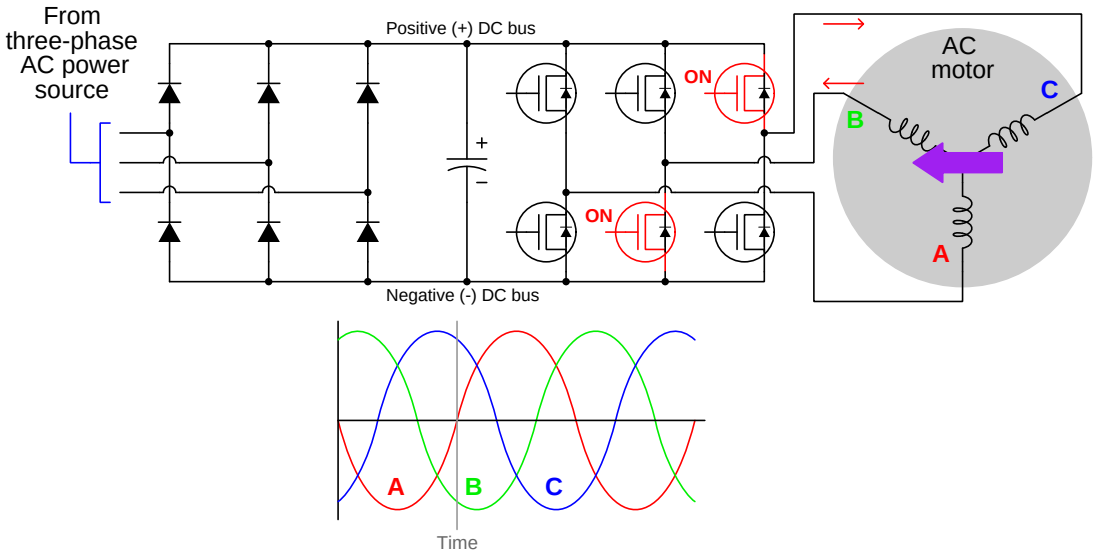
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
AC
motor
A
B
C
A
B
C
Time

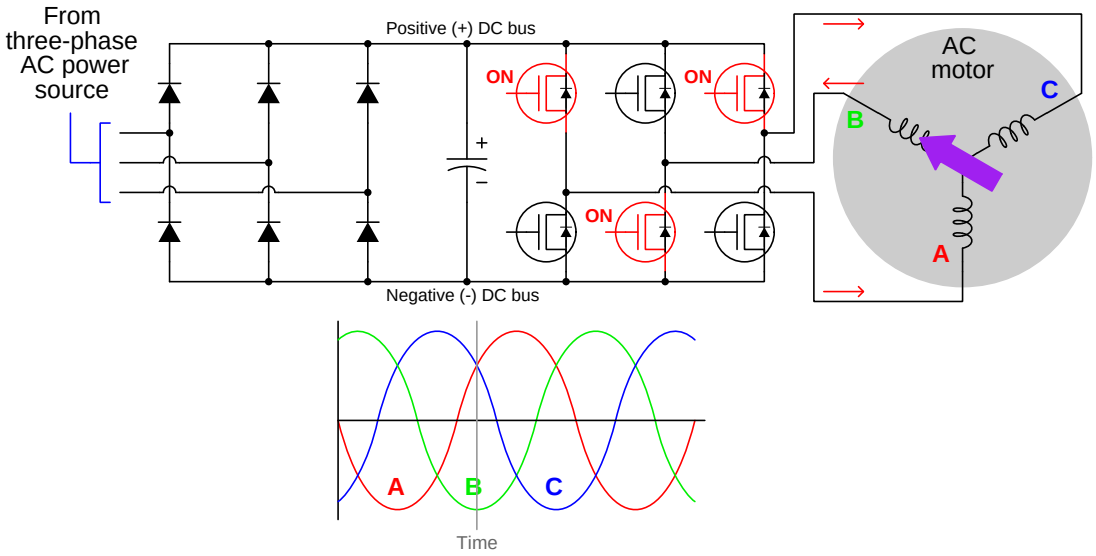
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
+
–
AC
motor
A
B
C
A
B
C
Time

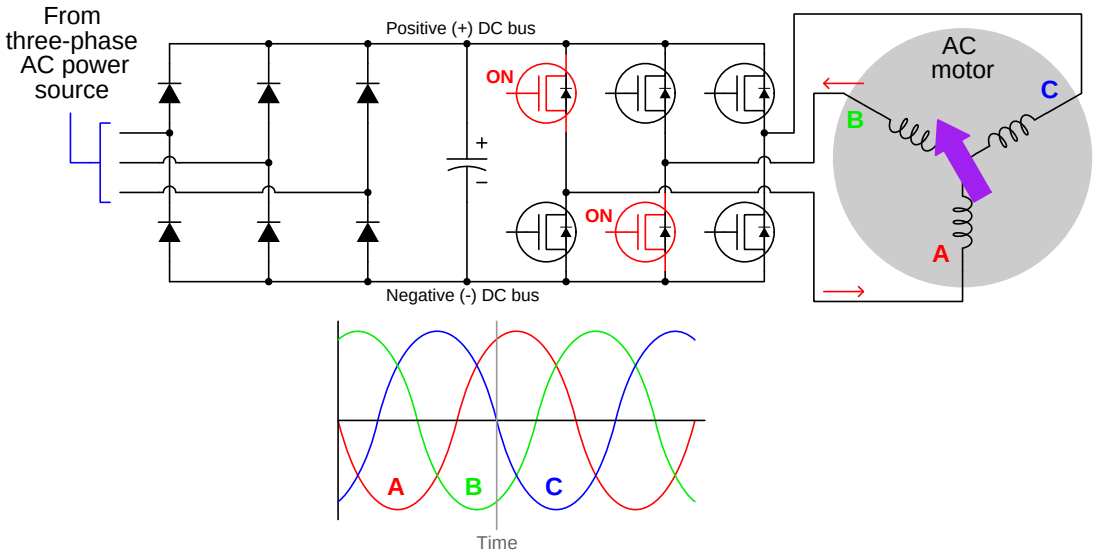
From
three-phase
AC power
source
Positive (+) DC bus
ON
ON
+
–
Negative (-) DC bus
AC
motor
C
B
A
A
B
C
Time

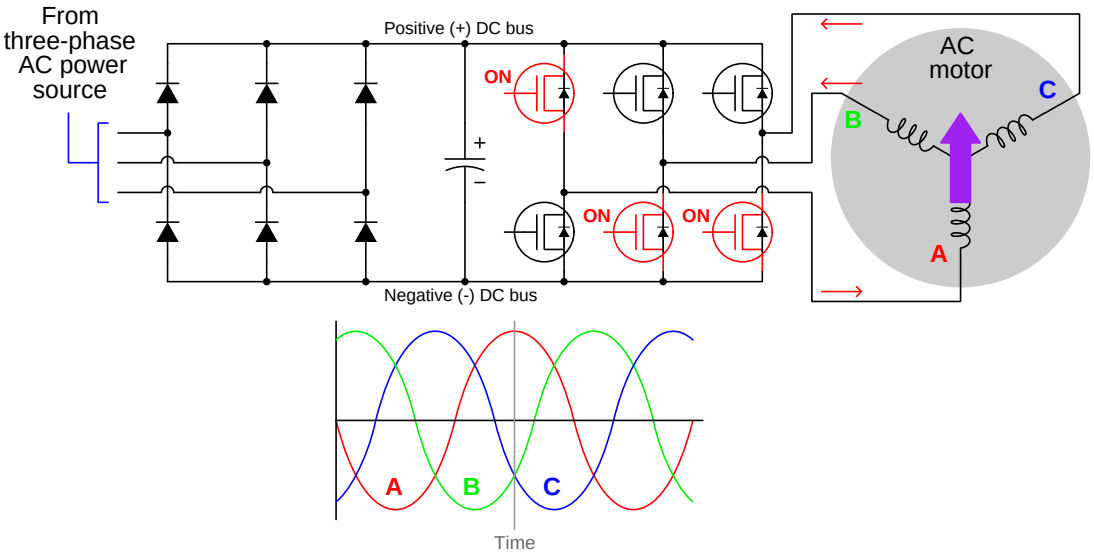
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
+
−
AC
motor
A
B
C
A
B
C
Time

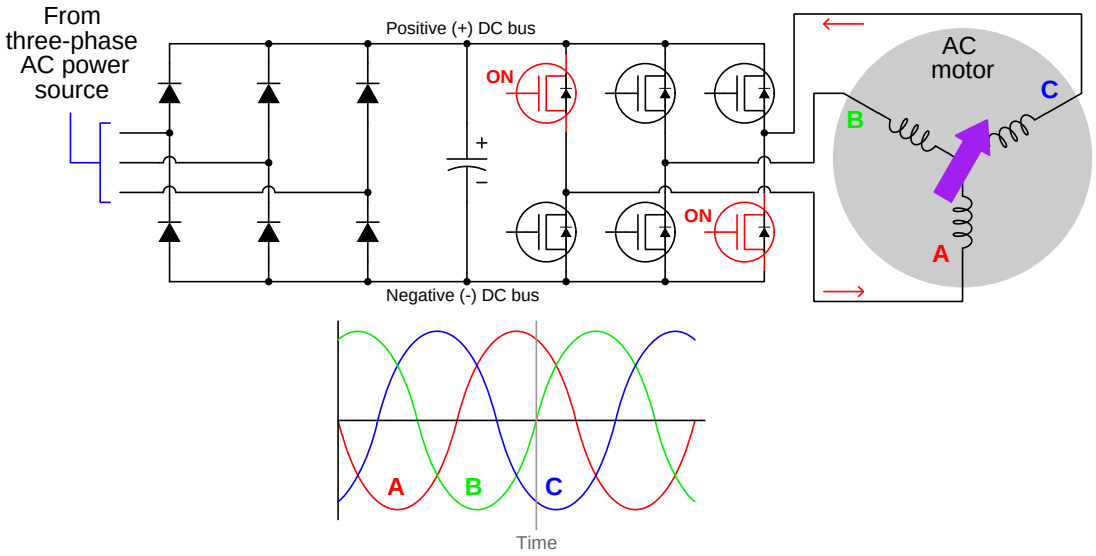
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
+
−
AC
motor
A
B
C
A
B
C
Time

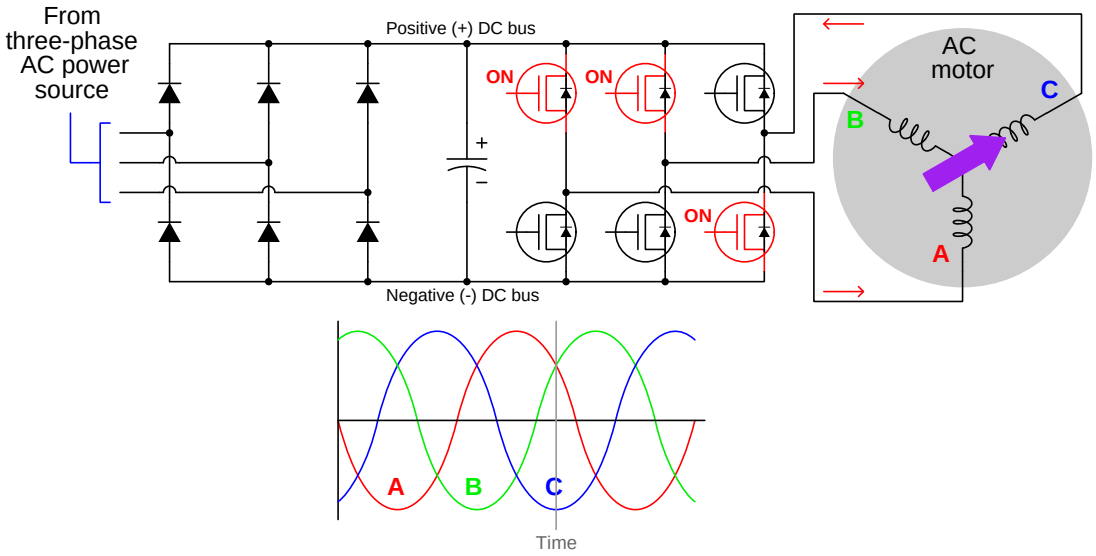
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
+
–
AC
motor
A
B
C
A
B
C
Time

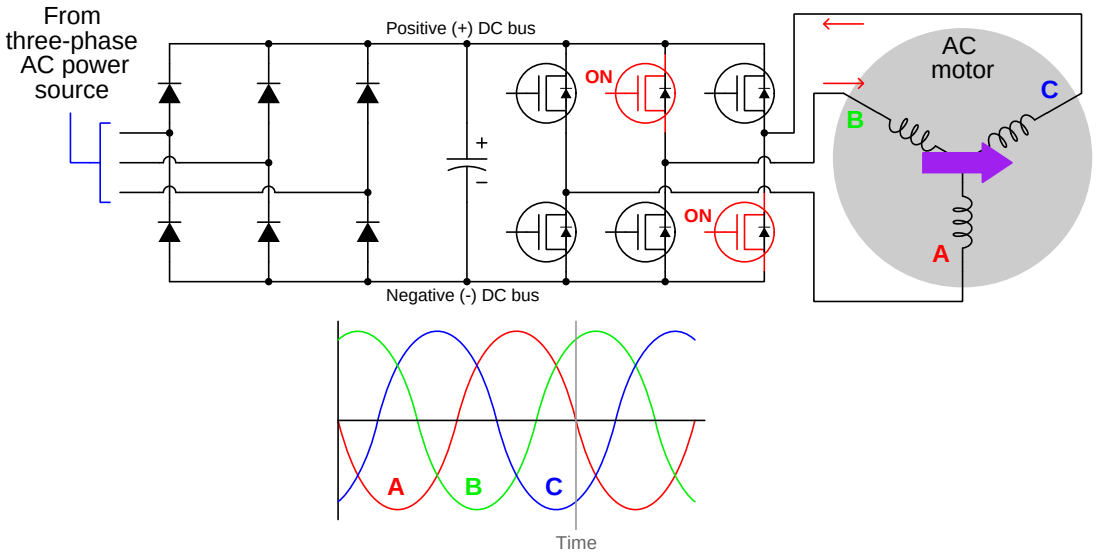
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
AC
motor
A
B
C
A
B
C
Time

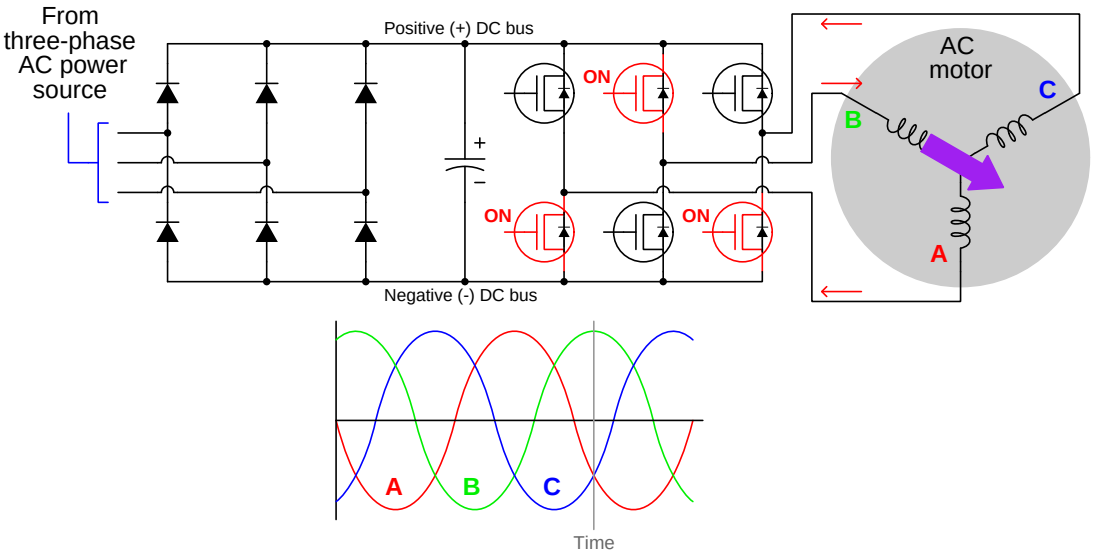
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
AC
motor
A
B
C
A
B
C
Time

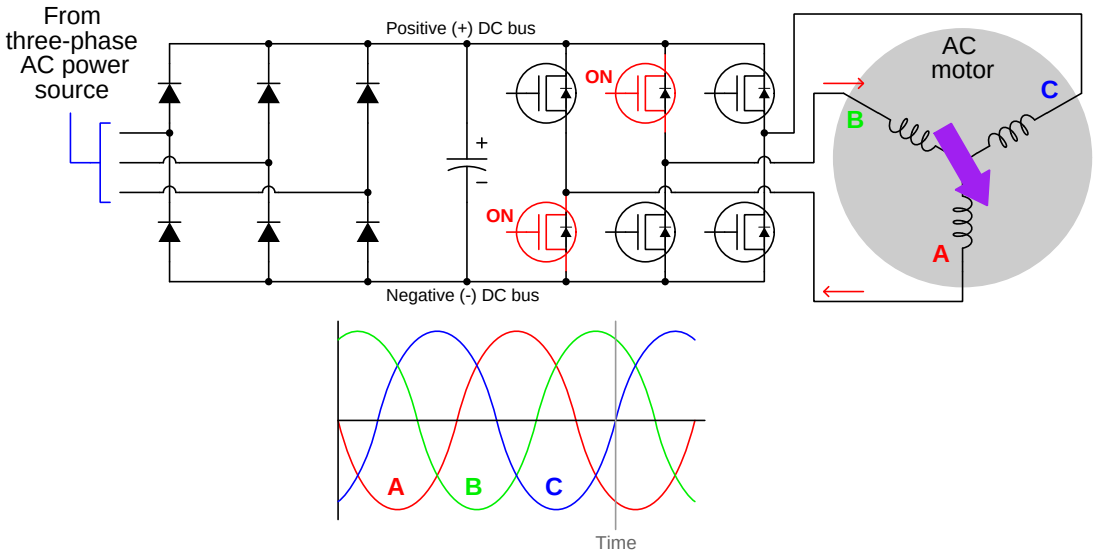
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
+
–
AC
motor
A
B
C
A
B
C
Time

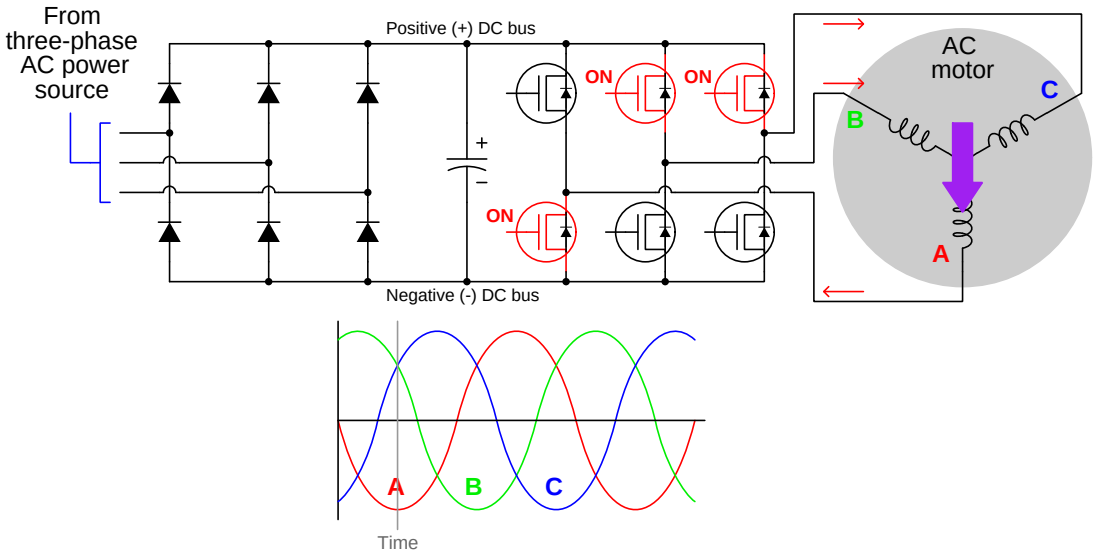
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
AC
motor
A
B
C
A
B
C
Time

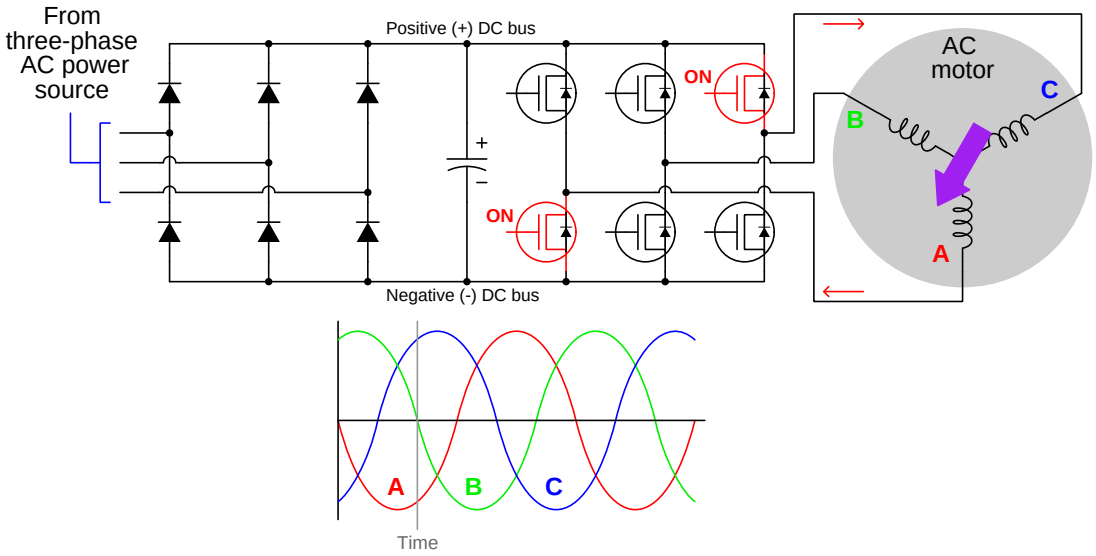
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
+
−
AC
motor
A
B
C
A
B
C
Time

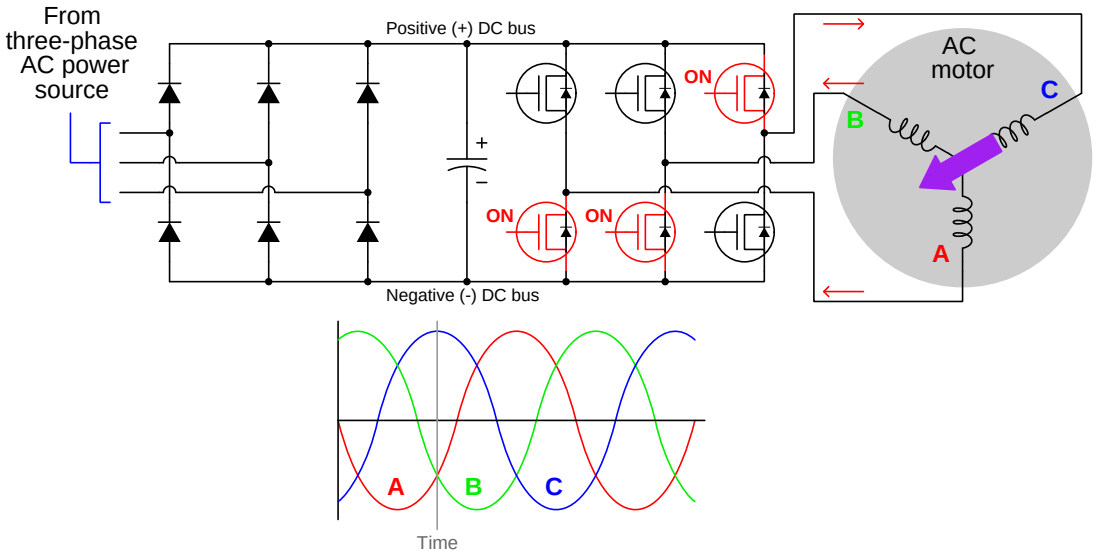
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
AC
motor
A
B
C
A
B
C
Time

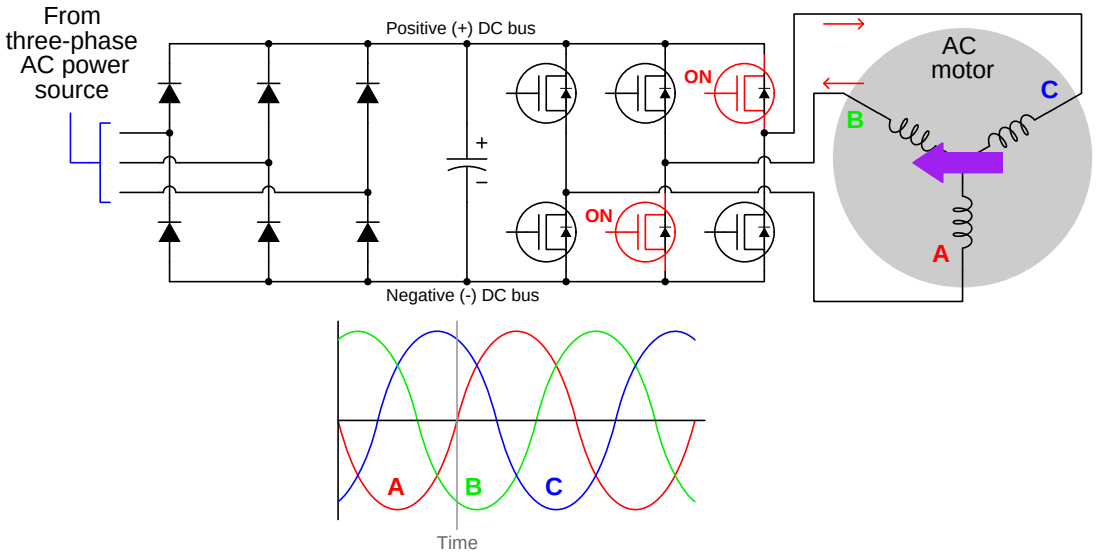
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
AC
motor
A
B
C
A
B
C
Time

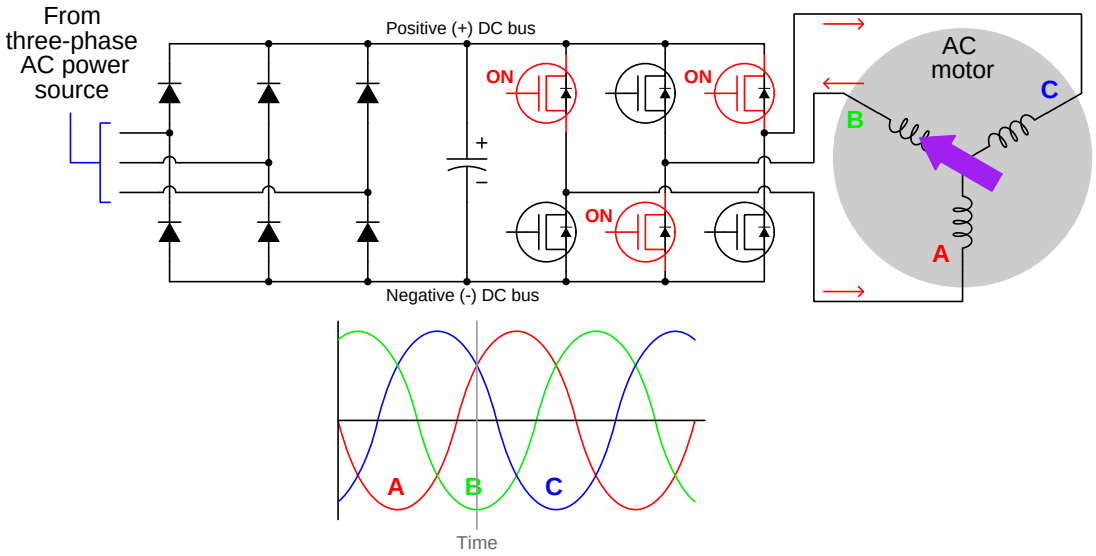
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
AC
motor
A
B
C
A
B
C
Time

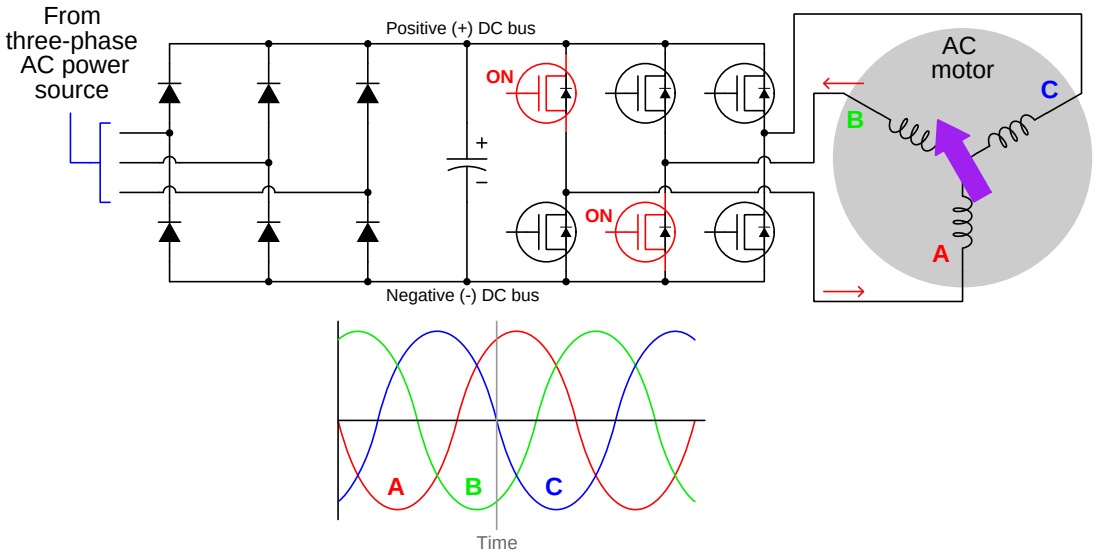
From
three-phase
AC power
source
Positive (+) DC bus
ON
ON
+
−
Negative (-) DC bus
AC
motor
C
B
A
A
B
C
Time

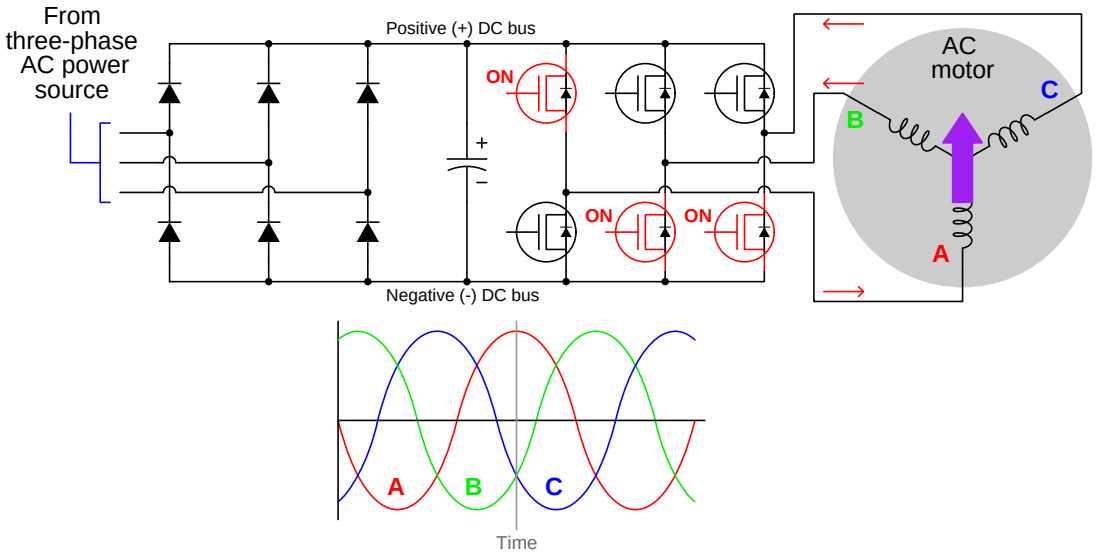
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
AC
motor
A
B
C
A
B
C
Time

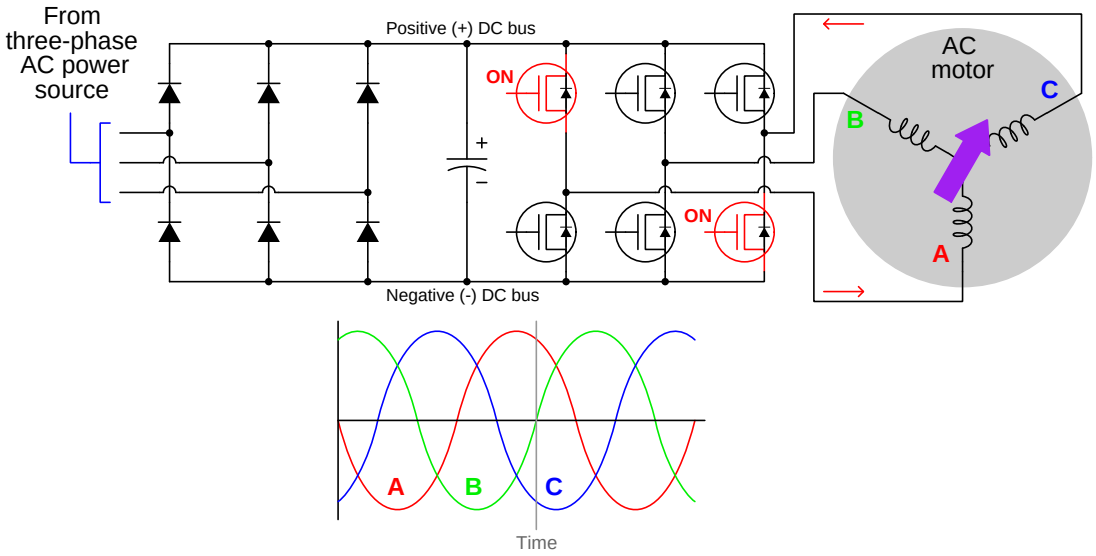
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
+
−
AC
motor
A
B
C
A
B
C
Time

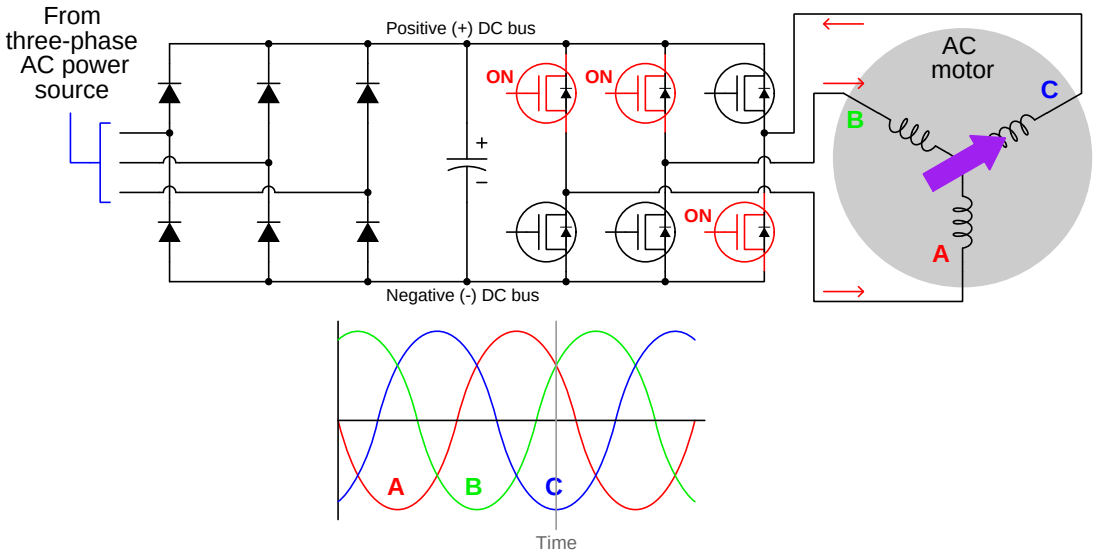
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
+
−
AC
motor
A
B
C
A
B
C
Time

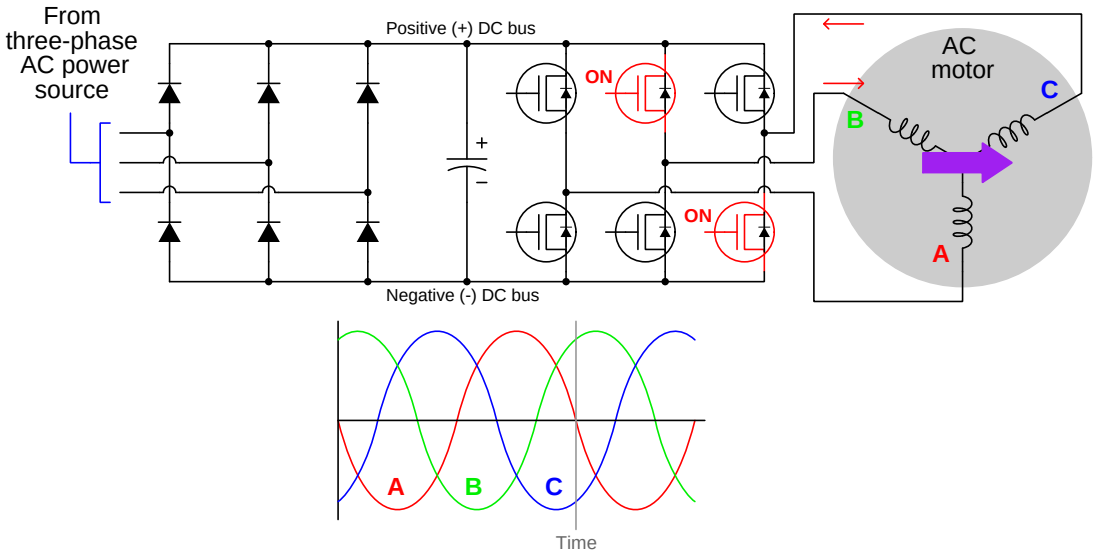
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
AC
motor
A
B
C
A
B
C
Time

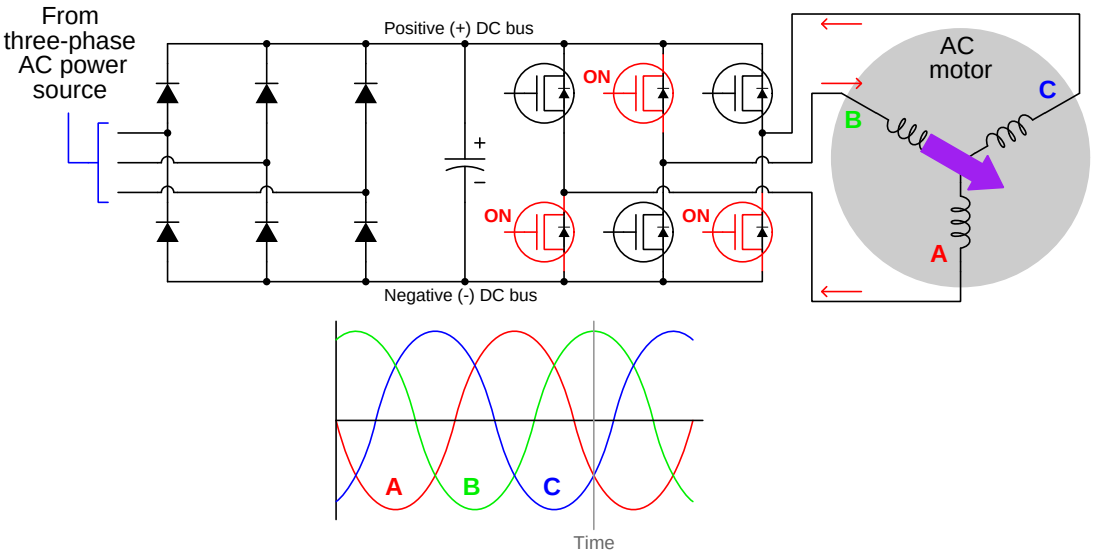
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
AC
motor
A
B
C
A
B
C
Time

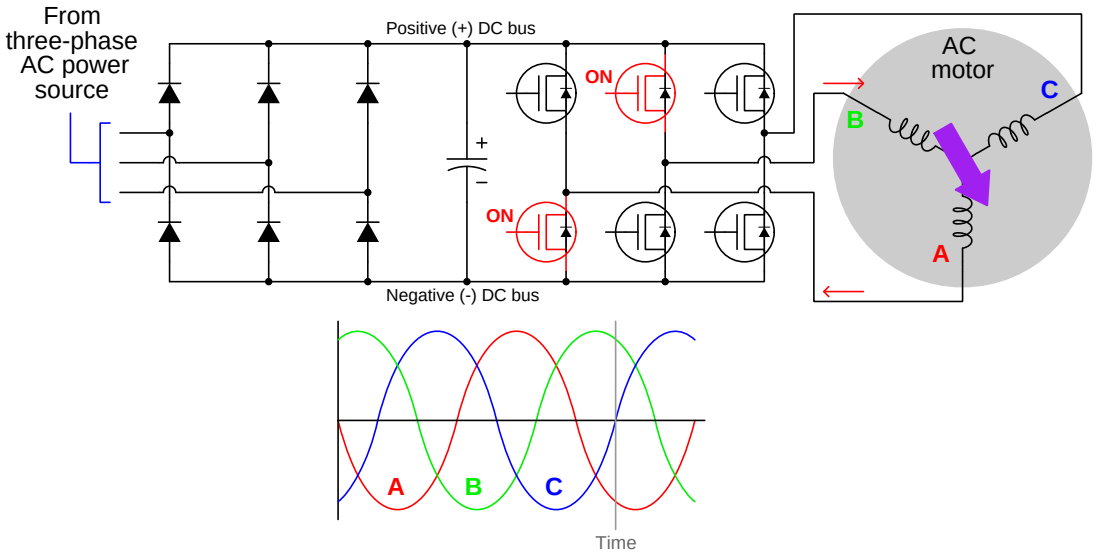
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
+
–
AC
motor
A
B
C
A
B
C
Time

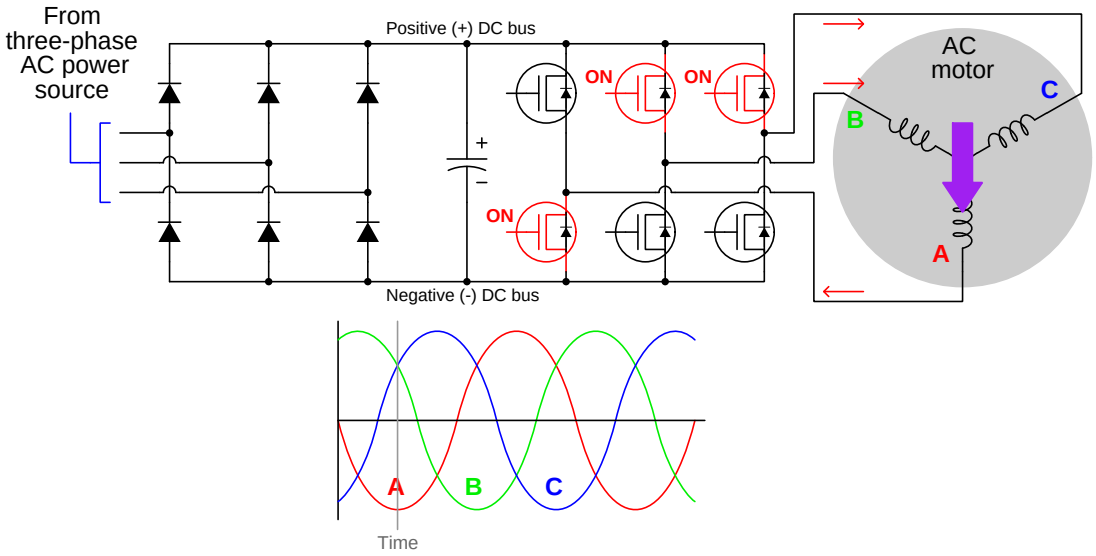
From
three-phase
AC power
source
Positive (+) DC bus
Negative (-) DC bus
ON
ON
ON
AC
motor
A
B
C
A
B
C
Time

6.2 Animation of an NPN BJT turning on

Bipolar junction transistors are naturally "off" devices, having no continuous channel of conductive semiconductor material joining emitter and collector terminals. Instead the NPN or PNP sandwich interrupts that path with two series-opposed PN junctions. However, when the base-emitter PN junction is forward-biased by an external source, majority charge carriers from the emitter region are drawn into the thin and lightly-doped base region where some of them travel out the base but most carry past into the depletion region of the base-collector PN junction to be swept into the collector region. For an NPN transistor this means electrons from the N-type emitter region will be injected into the P-type base region (as "minority" type carriers there) before most of them get swept into the N-type collector region.

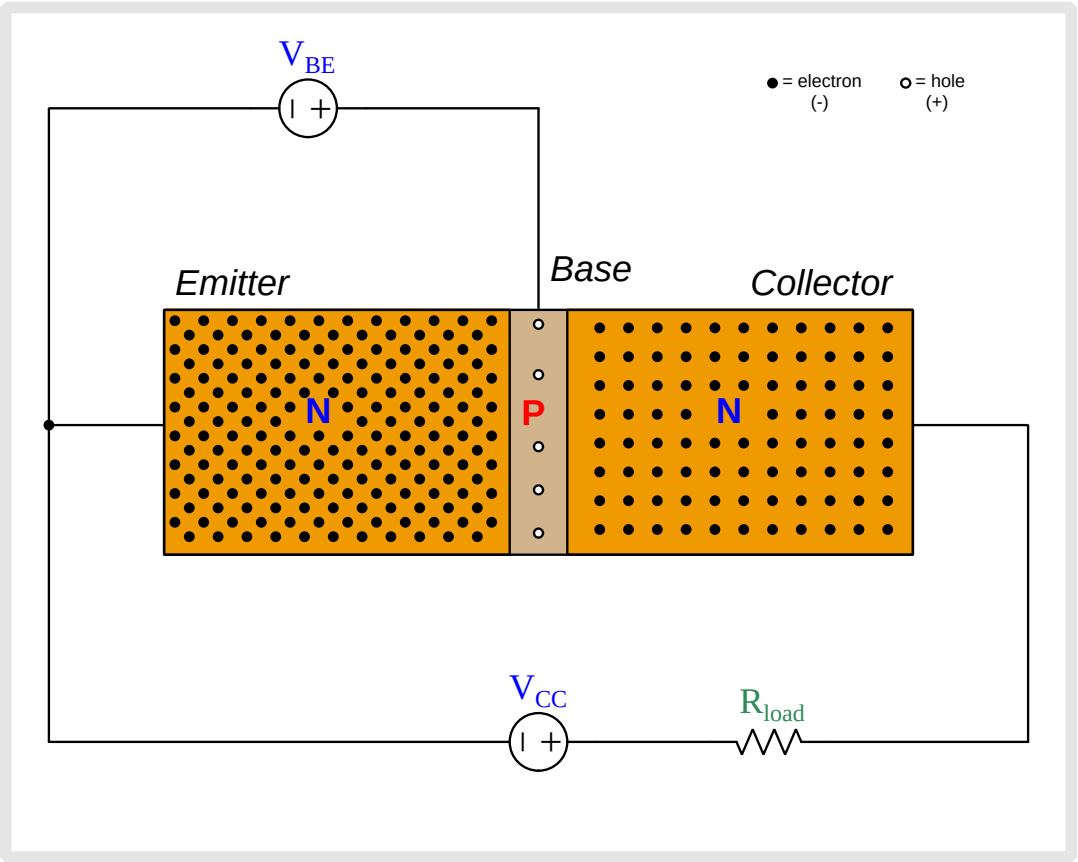

Above we see an NPN bipolar junction transistor connected to two sources. V_{BE} applies a controlling voltage between the base and emitter. V_{CC} is the voltage source which will drive current through the load resistor once the transistor begins to conduct.

The N-doped emitter region is heavily doped and therefore contains many free electrons (negatively-charged) as mobile charge carriers. The collector region is also doped as an N-type semiconductor but the dopant concentration is less than that of the emitter. The P-doped base region is the least-doped of them all with just a few holes (positively-charged) as mobile charge carriers. Note that when load current begins to flow, its direction will be shown in *conventional flow notation* which is the same direction in which holes move but is opposite to the direction free electrons move.

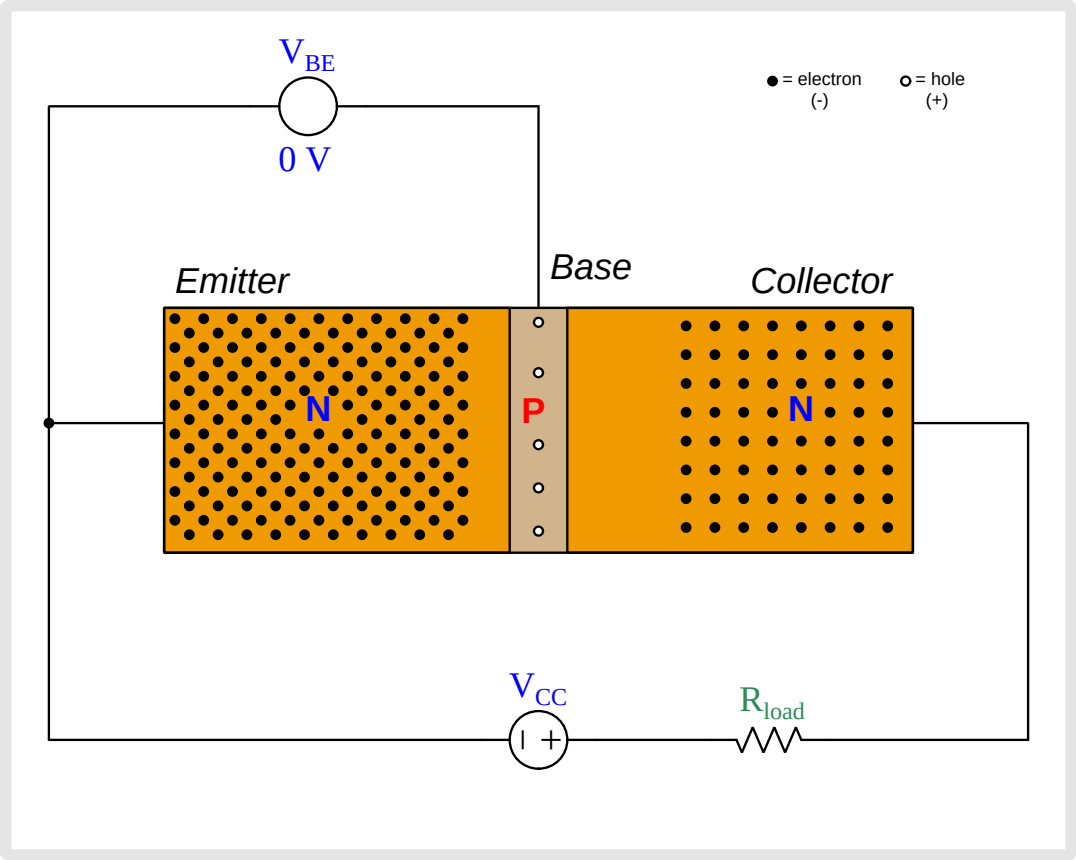

Note the depletion regions shown as areas with no free charge carriers, the base-collector depletion region being wider than the base-emitter depletion region due to the applied reverse-bias potential of V_{CC}.

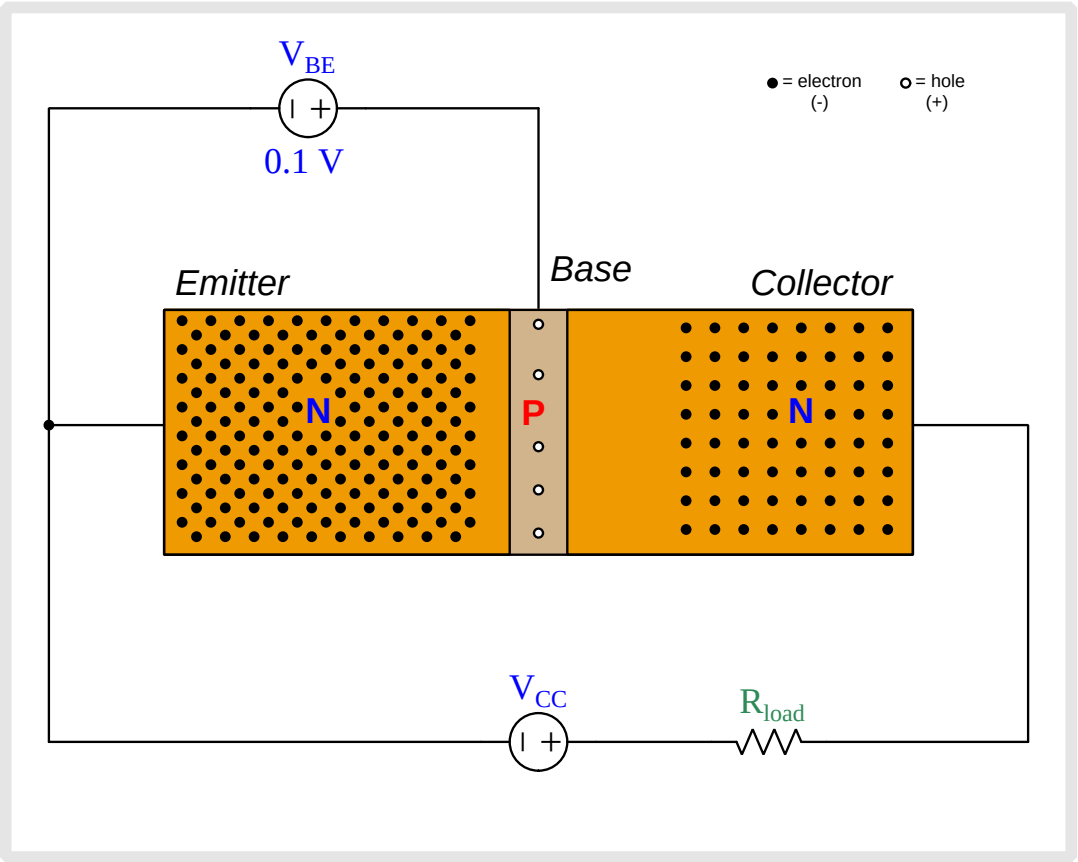

As we begin to apply a base-to-emitter voltage with the base positive and emitter negative, that applied potential difference drives the emitter and base regions' respective charge carriers toward each other, collapsing the base-emitter depletion region.

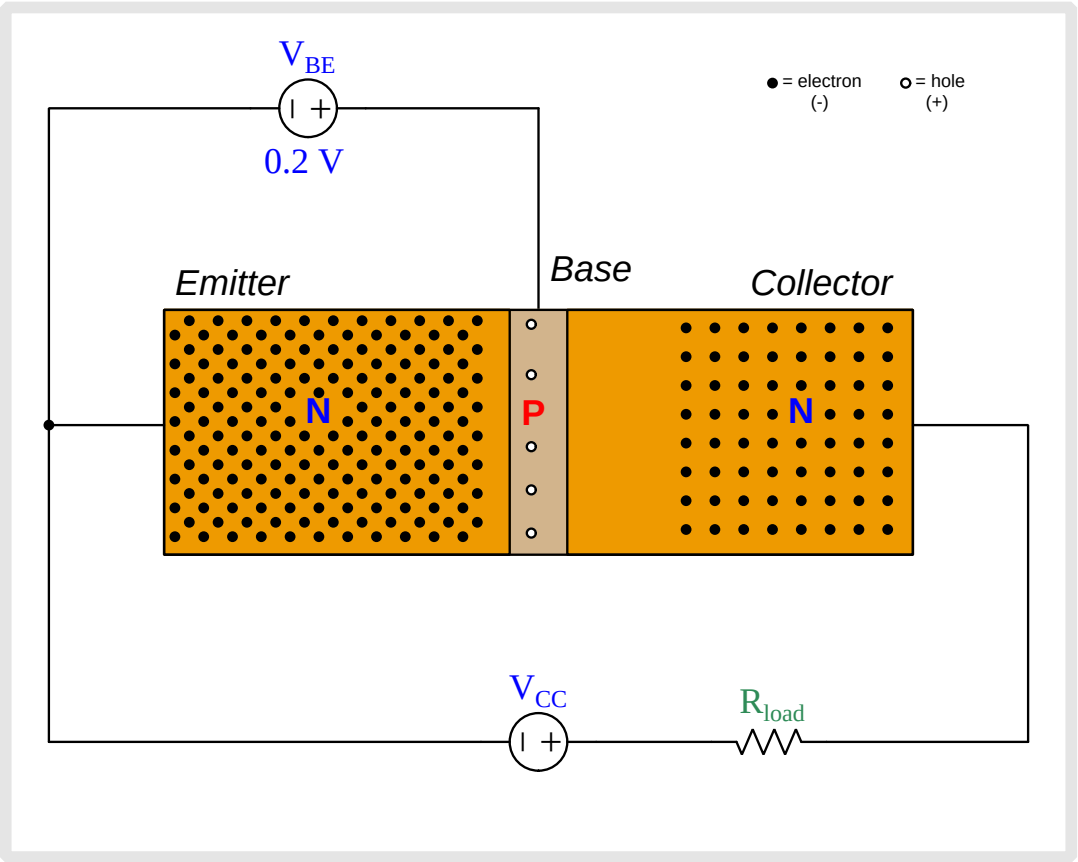

As we begin to apply a base-to-emitter voltage with the base positive and emitter negative, that applied potential difference drives the emitter and base regions' respective charge carriers toward each other, collapsing the base-emitter depletion region.

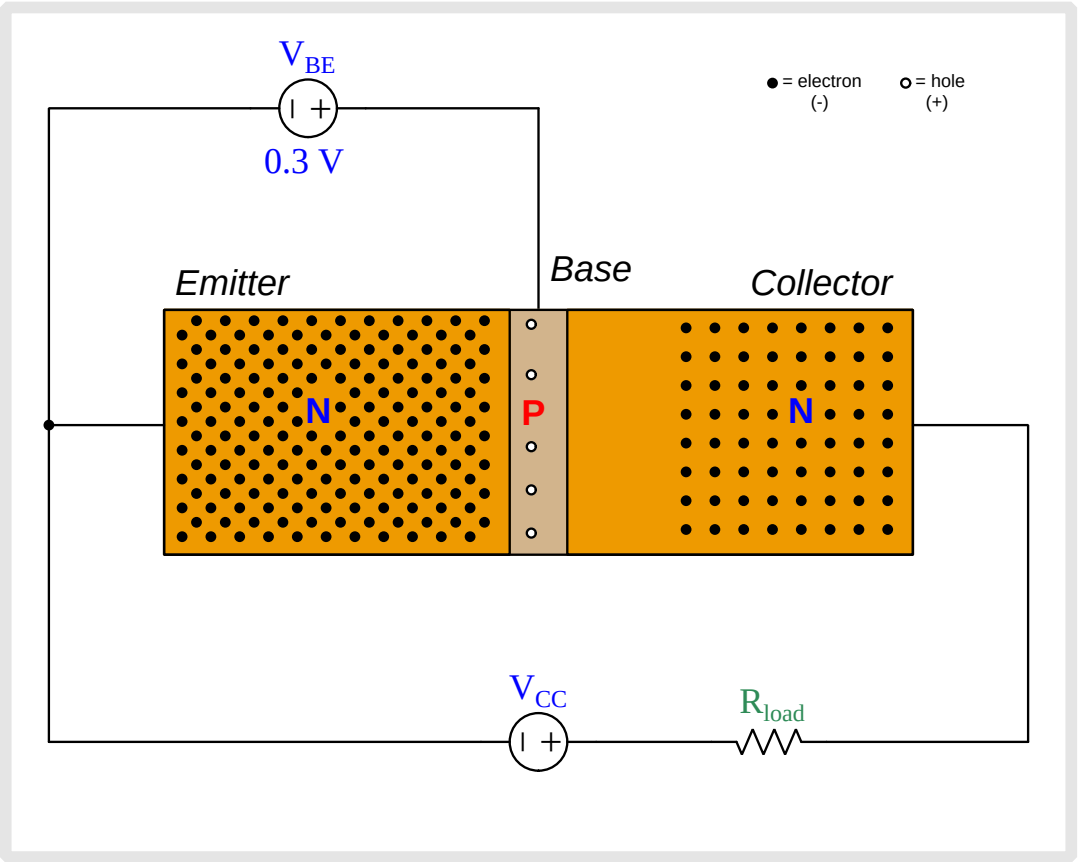

As we begin to apply a base-to-emitter voltage with the base positive and emitter negative, that applied potential difference drives the emitter and base regions' respective charge carriers toward each other, collapsing the base-emitter depletion region.

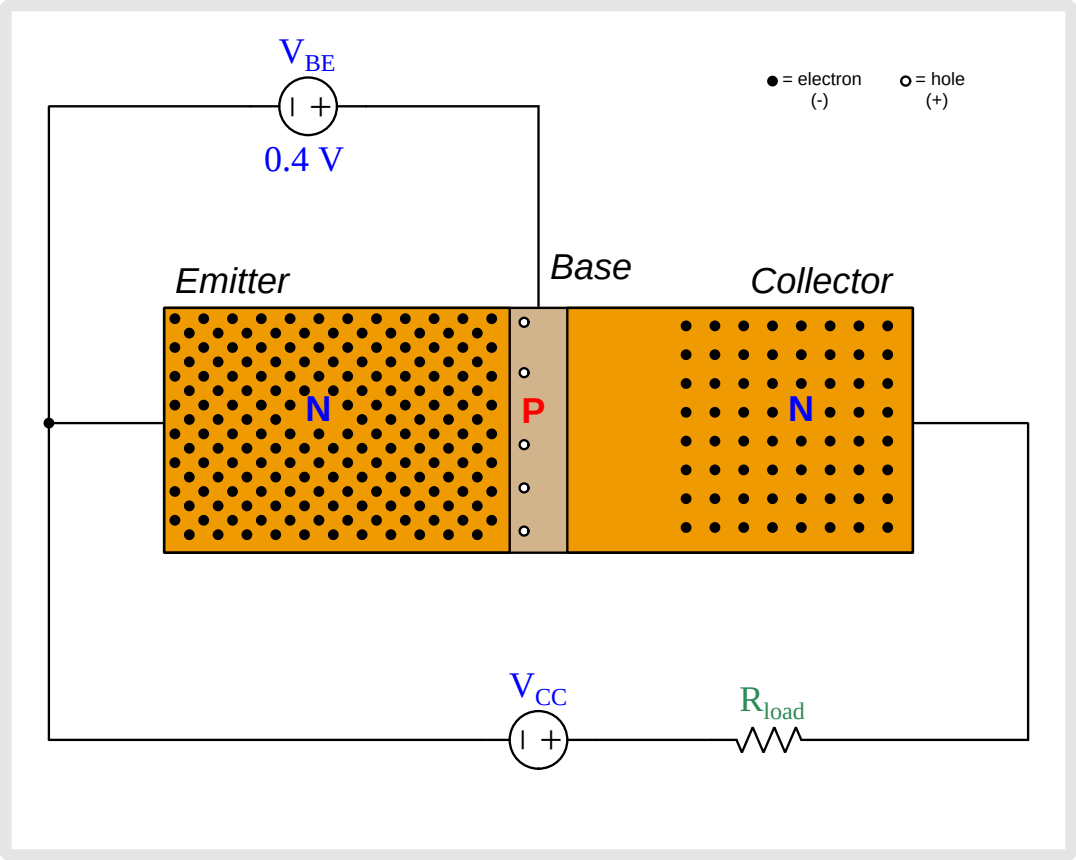

As we begin to apply a base-to-emitter voltage with the base positive and emitter negative, that applied potential difference drives the emitter and base regions' respective charge carriers toward each other, collapsing the base-emitter depletion region.

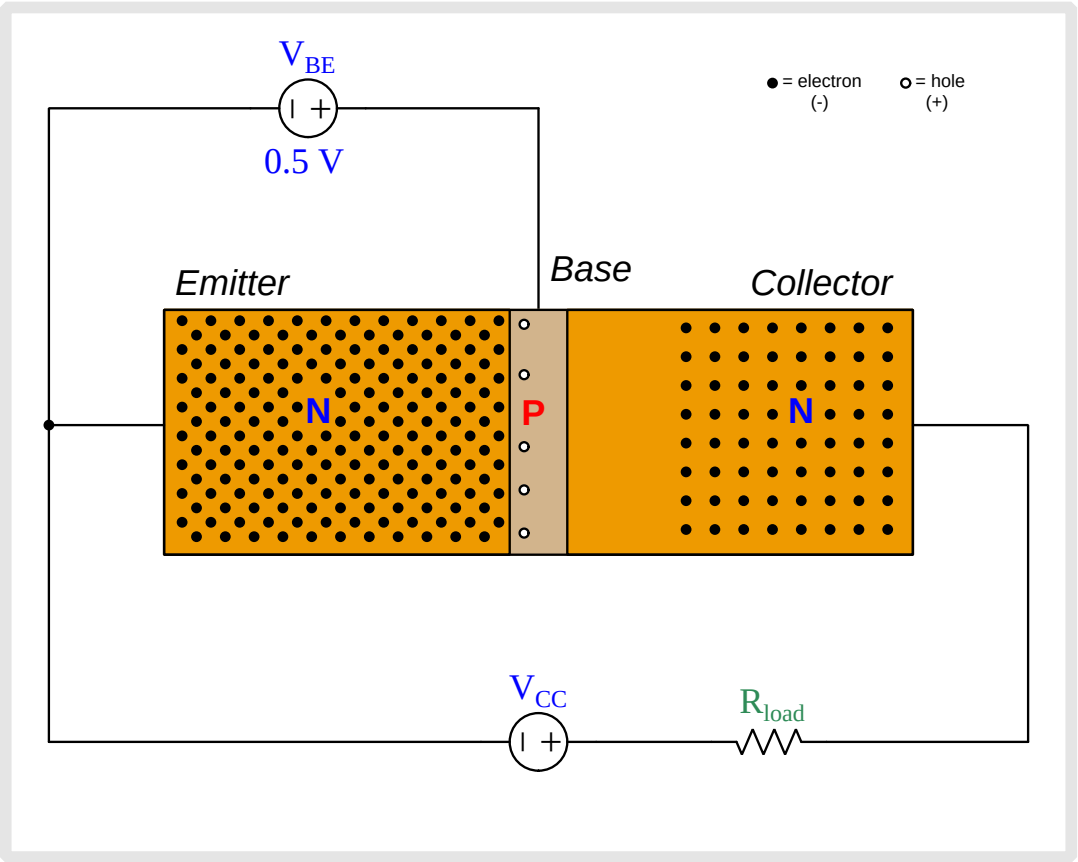

As we begin to apply a base-to-emitter voltage with the base positive and emitter negative, that applied potential difference drives the emitter and base regions' respective charge carriers toward each other, collapsing the base-emitter depletion region.

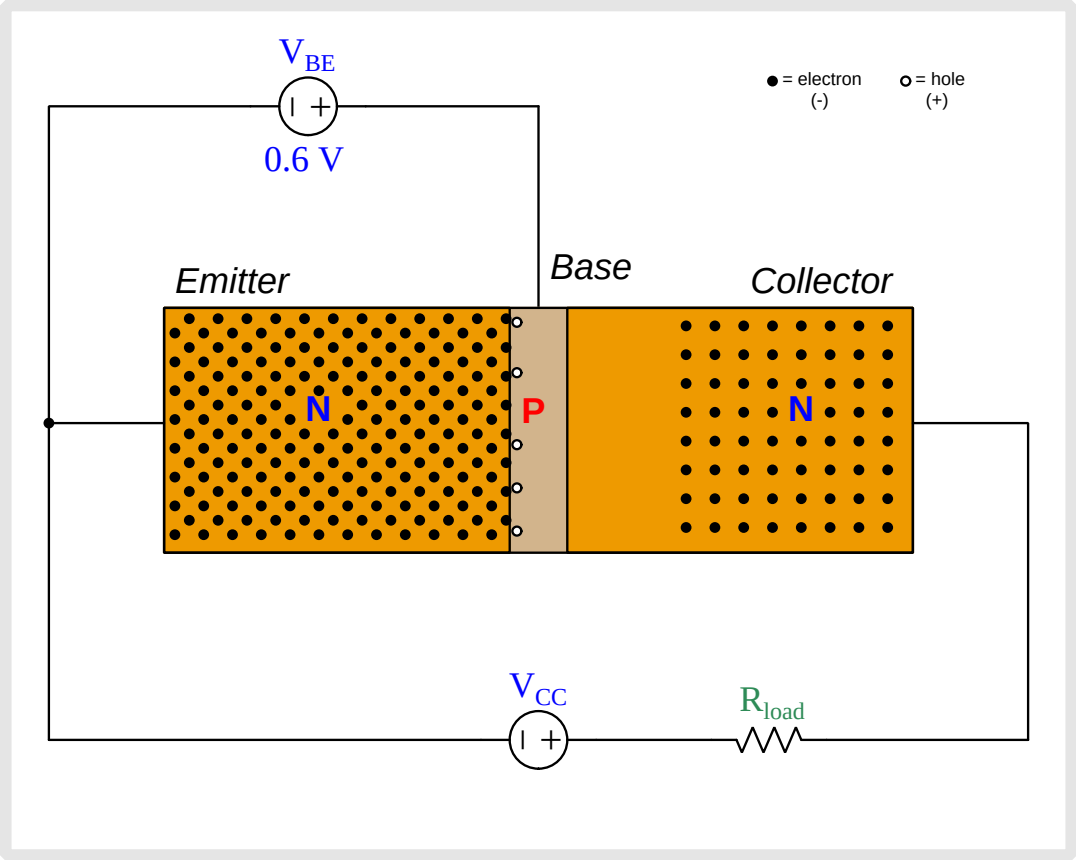

As we begin to apply a base-to-emitter voltage with the base positive and emitter negative, that applied potential difference drives the emitter and base regions' respective charge carriers toward each other, collapsing the base-emitter depletion region.

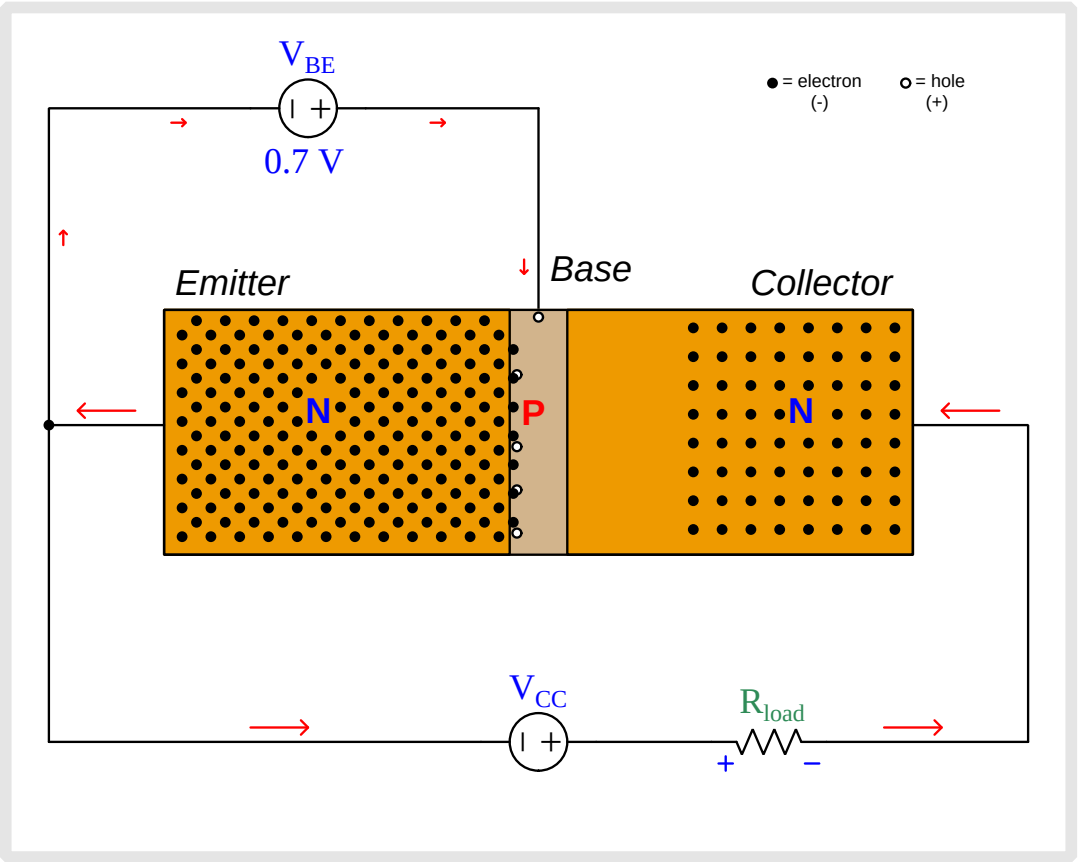

Now there is enough base-to-emitter voltage to fully collapse the base-emitter depletion region which allows base current to flow through the V_{BE} source. New holes enter into the base region from the base terminal to replenish any holes that have "re-combined" with electrons from the emitter region. Most of the emitter-injected electrons do *not* combine with holes, though, because the lightly-doped base region provides few holes to combine with. Therefore, most of the electrons injected into the base region from the emitter region begin to drift toward the collector.

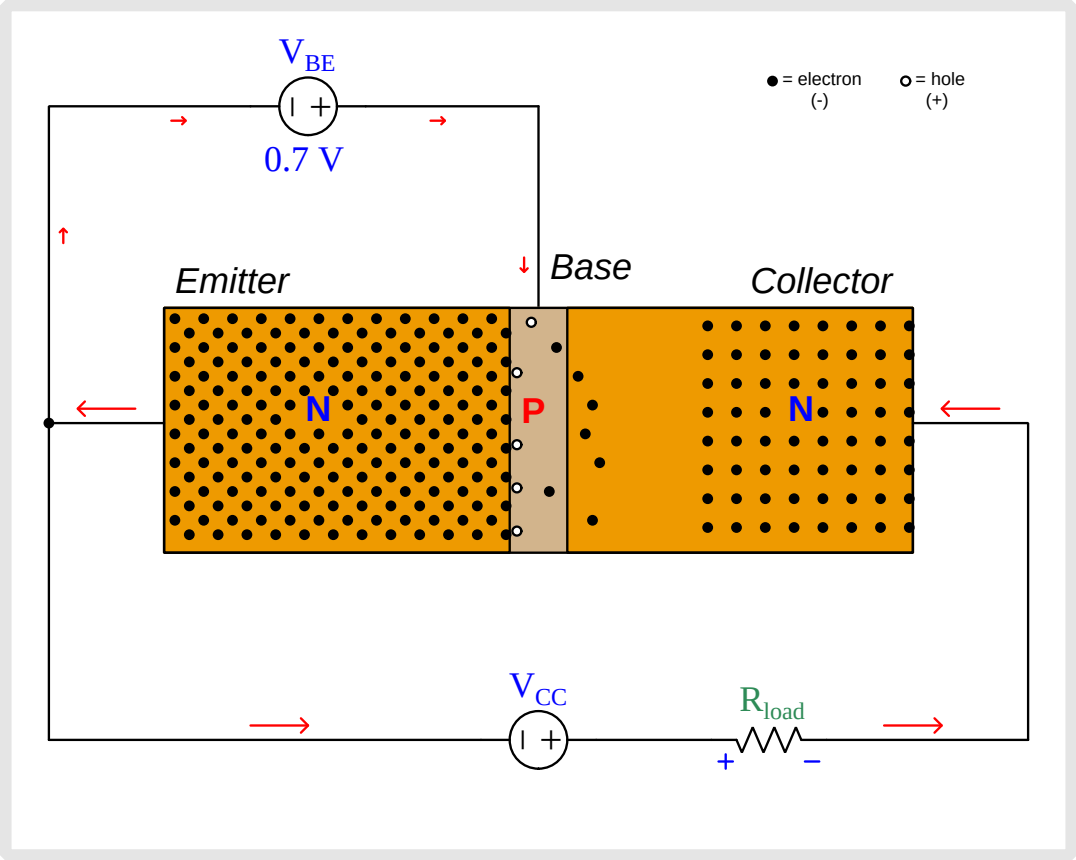

The electric field imposed by V_{CC} across the base-collector PN junction quickly "sweeps" electrons toward the collector terminal, constituting current through the load resistor. A small number of electrons injected into the base "recombined" with holes there to constitute base current rather than collector current.

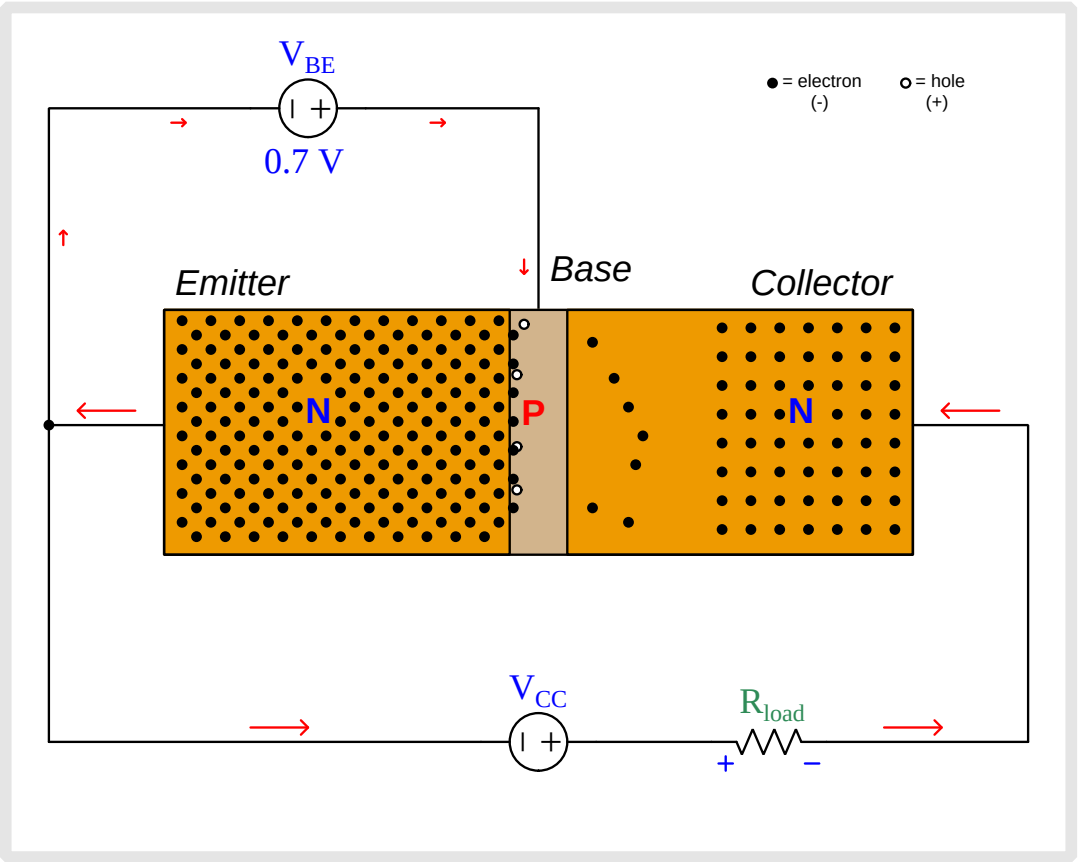

The electric field imposed by V_{CC} across the base-collector PN junction quickly "sweeps" electrons toward the collector terminal, constituting current through the load resistor. A small number of electrons injected into the base "recombined" with holes there to constitute base current rather than collector current.

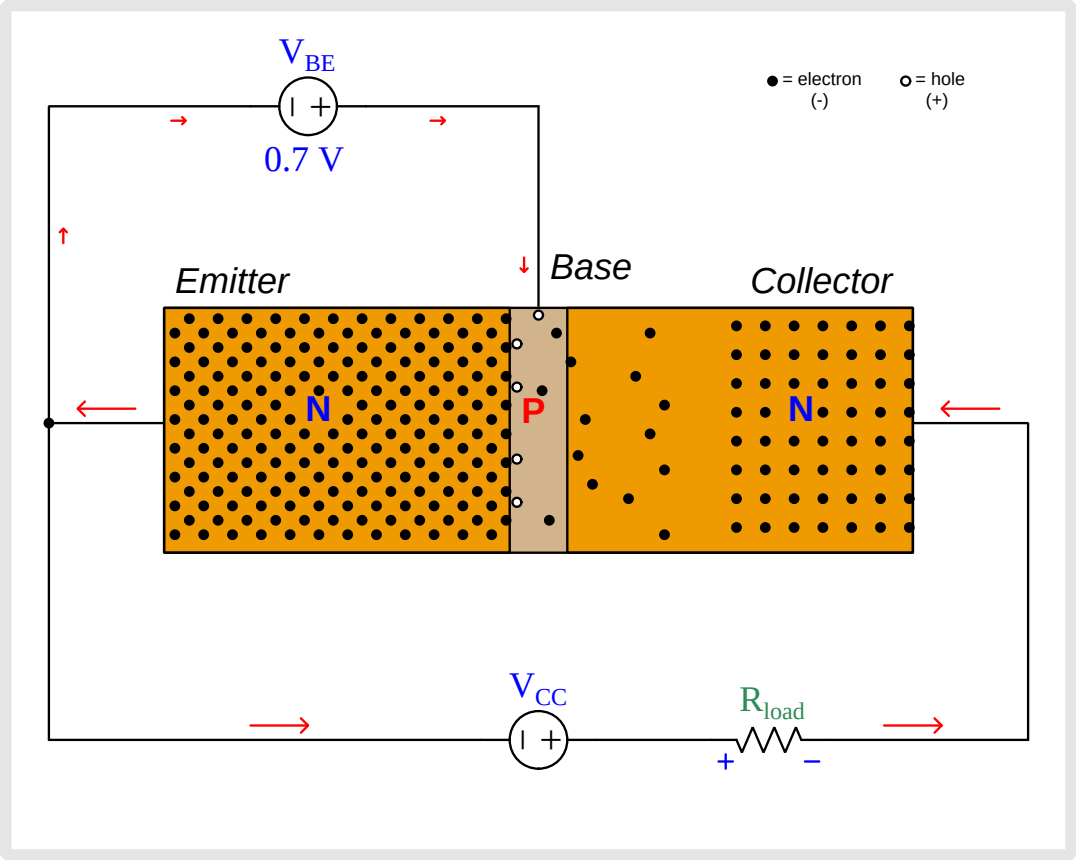

The electric field imposed by V_{CC} across the base-collector PN junction quickly "sweeps" electrons toward the collector terminal, constituting current through the load resistor. A small number of electrons injected into the base "recombined" with holes there to constitute base current rather than collector current.

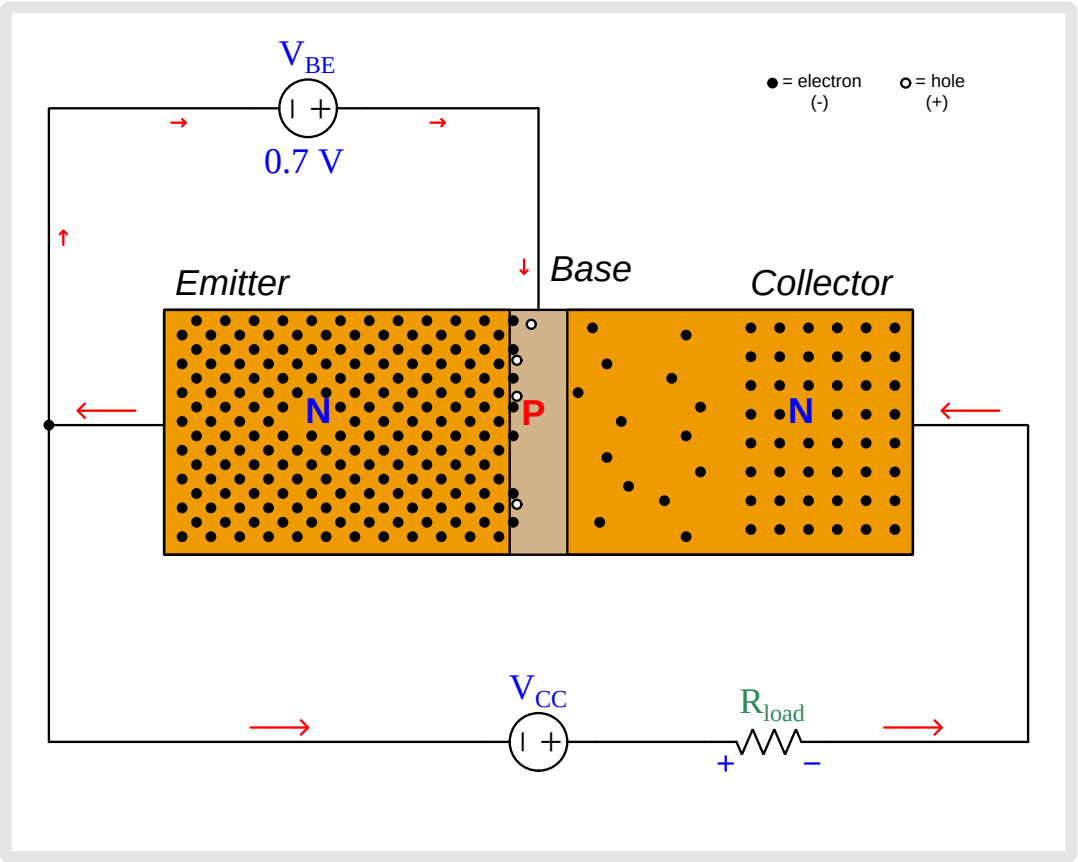

The electric field imposed by V_{CC} across the base-collector PN junction quickly "sweeps" electrons toward the collector terminal, constituting current through the load resistor. A small number of electrons injected into the base "recombined" with holes there to constitute base current rather than collector current.

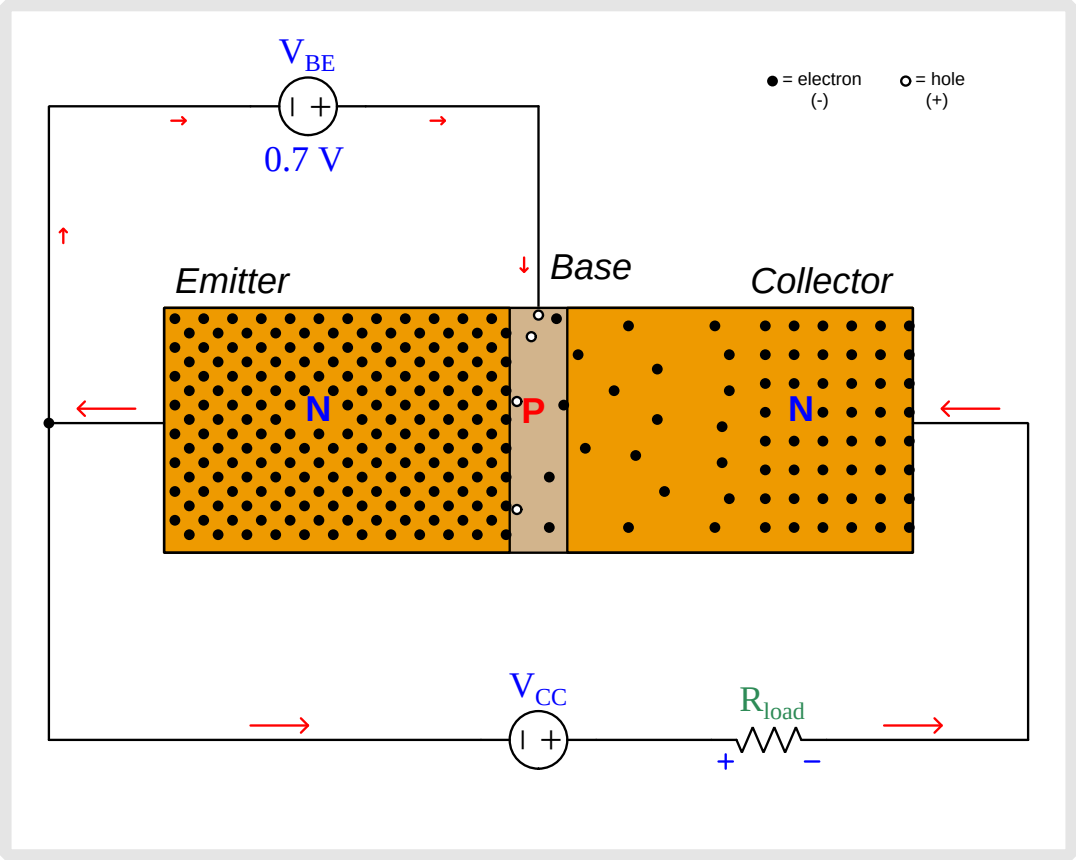

The electric field imposed by V_{CC} across the base-collector PN junction quickly "sweeps" electrons toward the collector terminal, constituting current through the load resistor. A small number of electrons injected into the base "recombined" with holes there to constitute base current rather than collector current.

6.3 Animation of an E-type MOSFET turning on

Enhancement-type (E-type) MOSFET transistors are naturally "off" devices, having no continuous channel of conductive semiconductor material joining source and drain by default the way D-type MOSFETs do. Instead the substrate semiconductor material, which is doped oppositely of the source and drain regions, interrupts that path with two series-opposed PN junctions. However, when a transverse electric field is applied between gate and substrate by an external voltage source, the surface-layer of the substrate in direct contact with the insulating metal-oxide layer converts to the opposite type of semiconductor. For an N-channel MOSFET with N-doped source and drain regions and with P-doped substrate, this means the surface of the P-doped substrate just beneath the gate effectively turns into N-type by the force of the electric field drawing free electrons from source and drain (and a few minority electrons existing within the substrate) to that surface region to form a thin film of free electrons capable of supporting source-drain conduction.

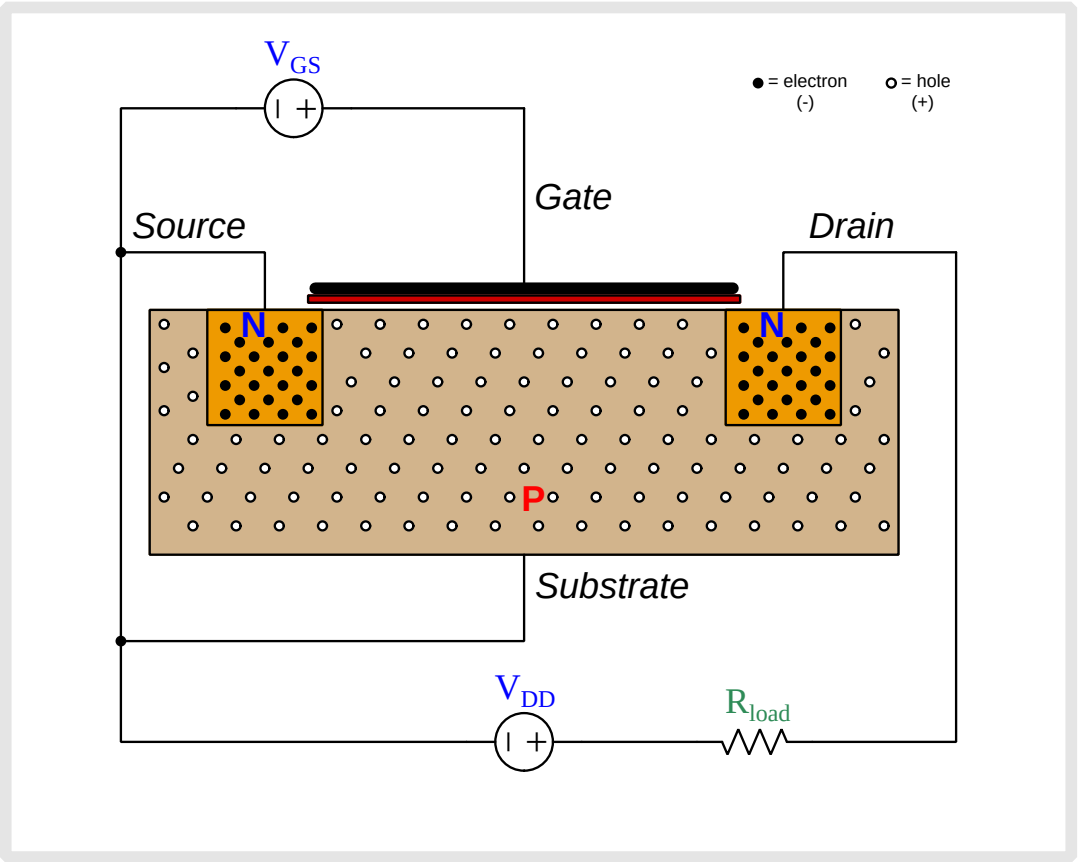

Above we see an N-channel, E-type MOSFET connected to two sources. V_{GS} applies a controlling voltage between the gate and the substrate, with the substrate connected to the source terminal to be one and the same. V_{DD} is the voltage source which will drive current through the load resistor once the transistor begins to conduct.

The N-doped source and drain regions have numerous free electrons (negatively-charged) as mobile charge carriers, while the P-doped substrate has holes (positively-charged). Note that when load current begins to flow, its direction will be shown in *conventional flow notation* which is the same direction in which holes move but is opposite to the direction free electrons move.

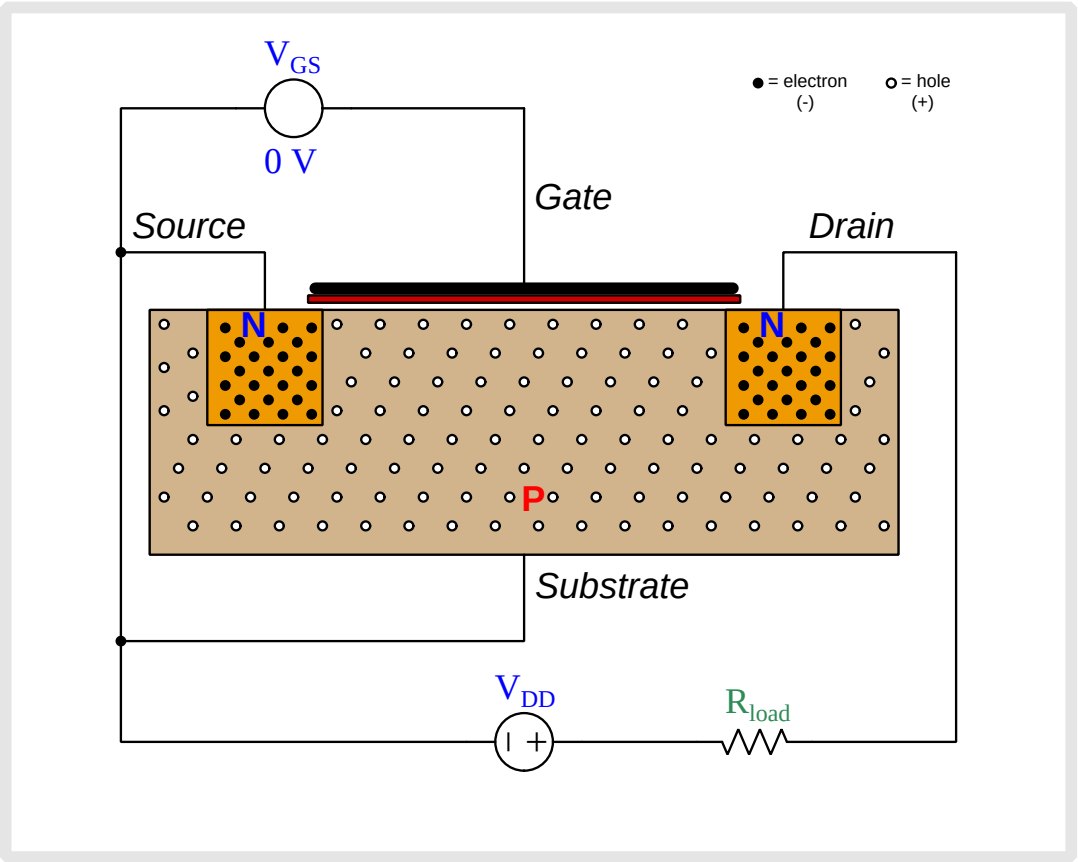
V_{GS}
0 V
● = electron (-)
○ = hole (+)
Gate
Source
Drain
N
N
P
Substrate
V_{DD}
R_{load}

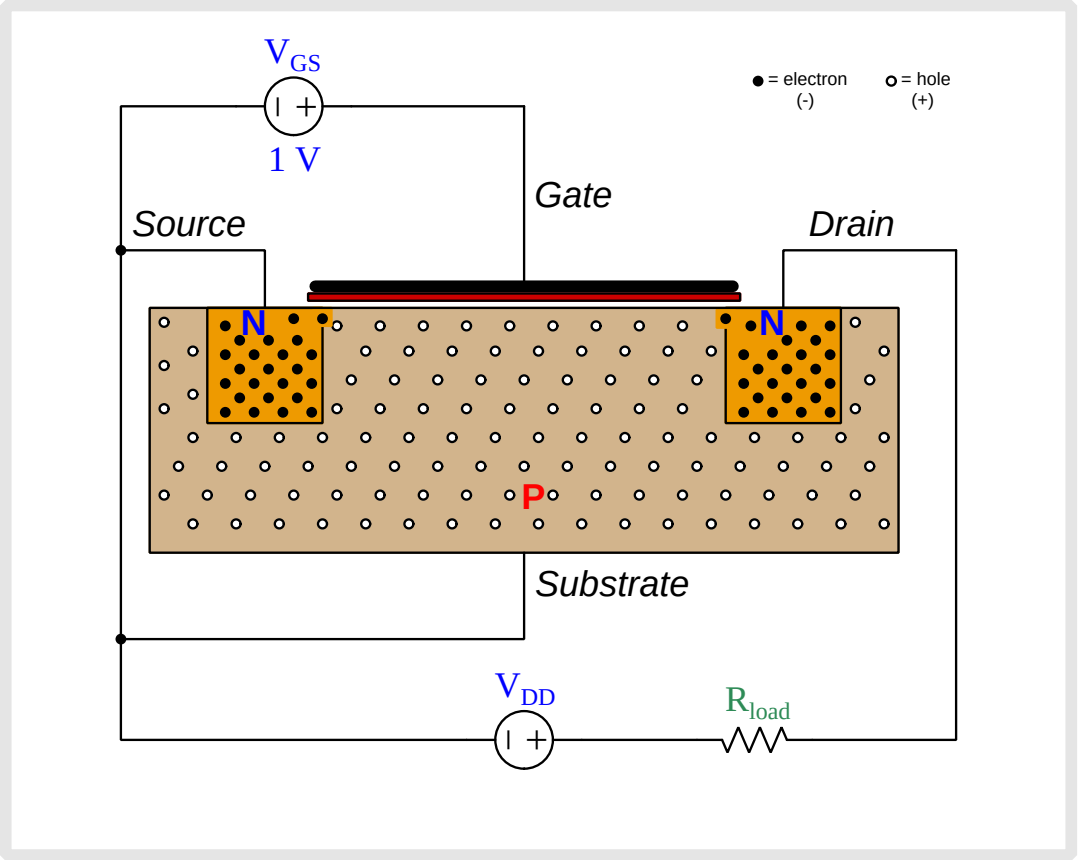

As we begin to apply a gate-to-substrate voltage with the gate positive and substrate negative, negatively-charged free electrons from the source and drain regions are attracted toward the metal gate while positively-charged holes in the substrate are repelled from the gate and driven downward.

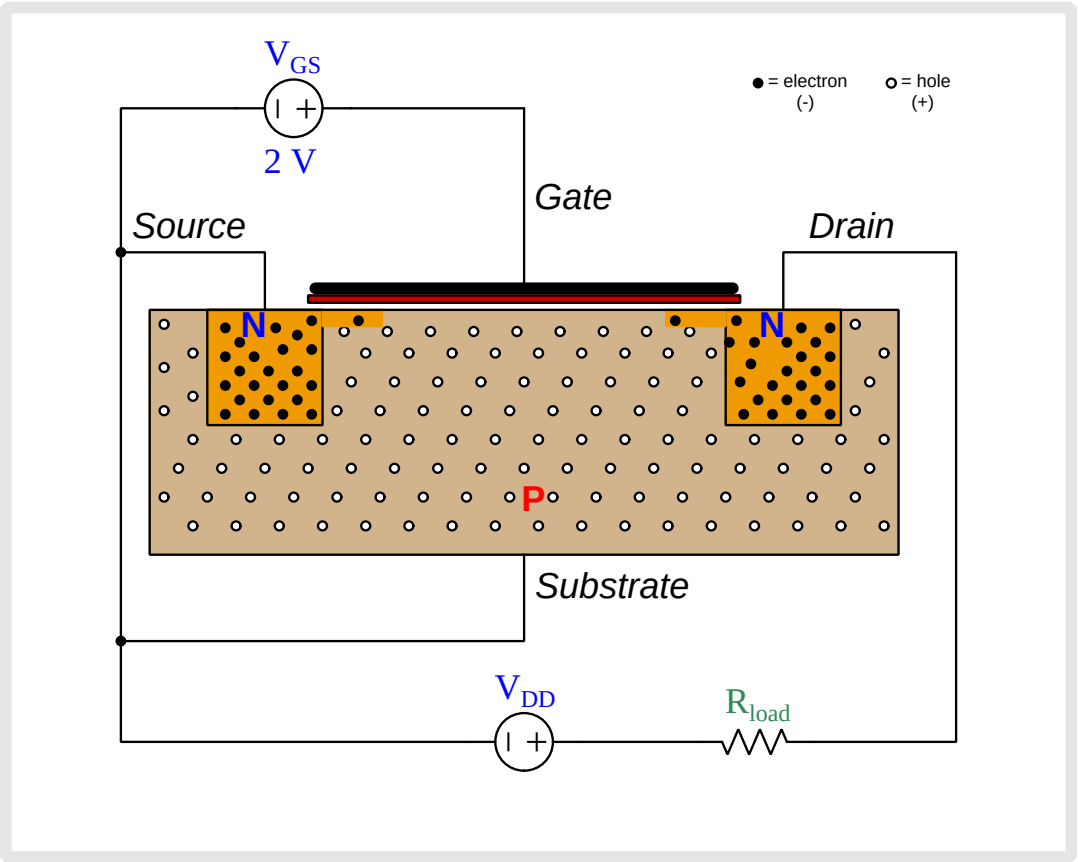

As this V_{GS} voltage increases, a thin film of free electrons begins to form beneath the gate, effectively turning what used to be P-type semiconductor material into N-type. This is called an *inversion layer*.

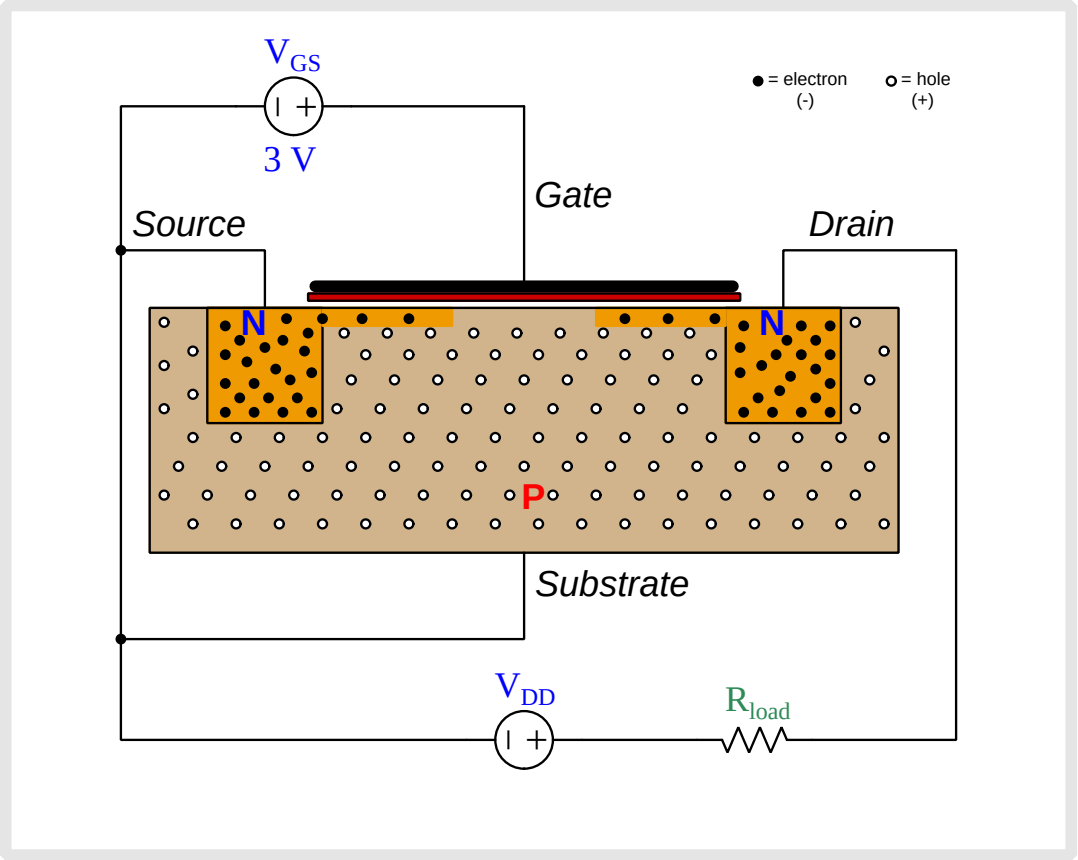

As this V_{GS} voltage increases, a thin film of free electrons begins to form beneath the gate, effectively turning what used to be P-type semiconductor material into N-type. This is called an *inversion layer*.

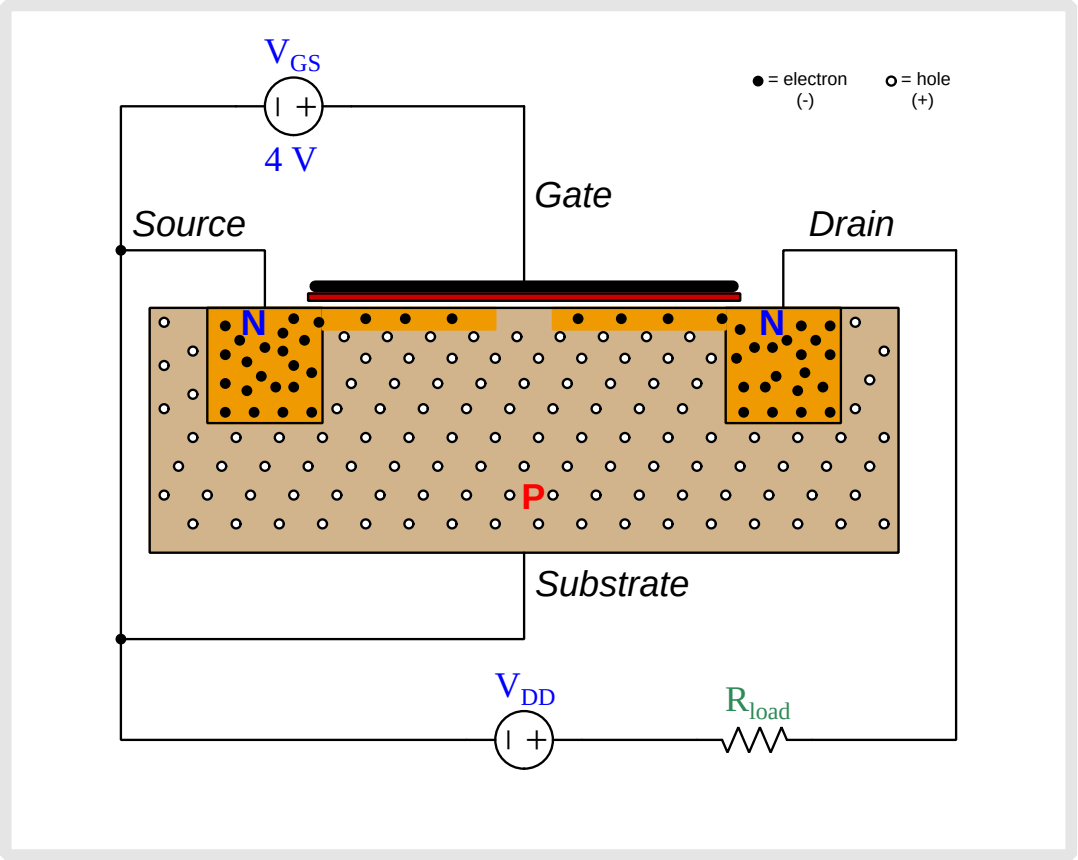

As this V_{GS} voltage increases, a thin film of free electrons begins to form beneath the gate, effectively turning what used to be P-type semiconductor material into N-type. This is called an *inversion layer*.

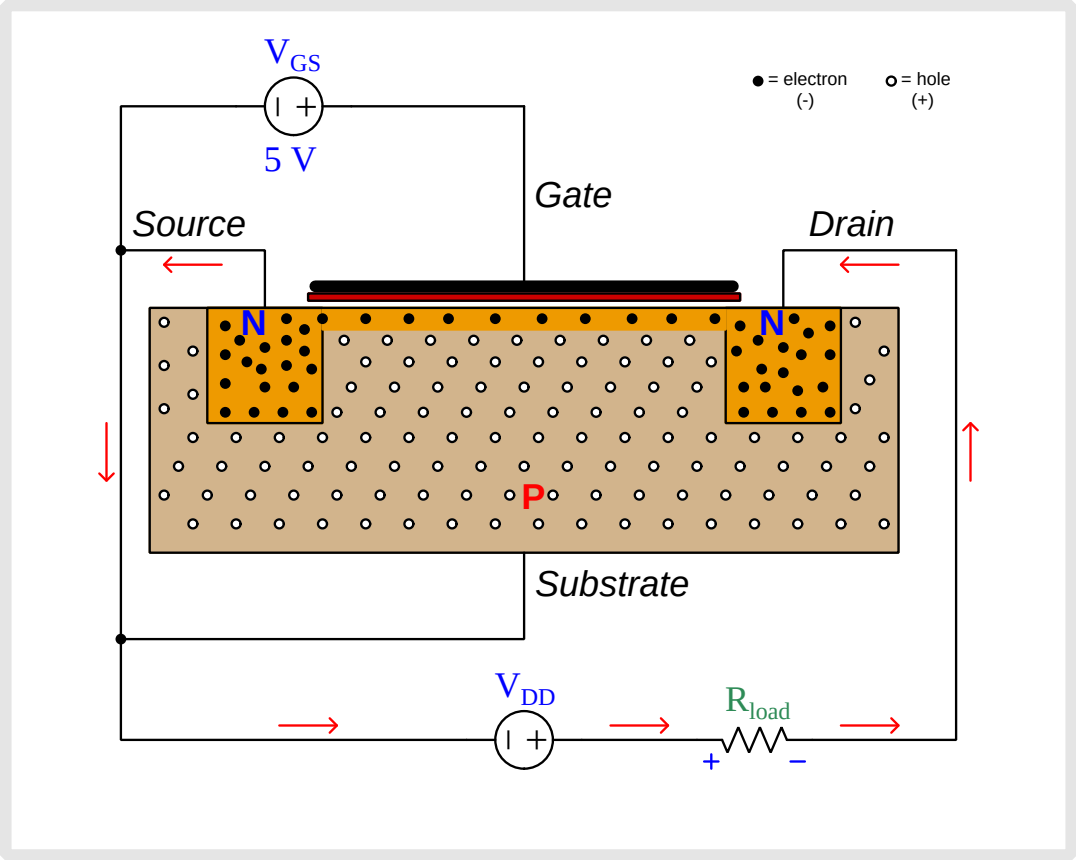

Once V_{GS} reaches the threshold or "turn-on" level, this inversion layer completely connects drain and source to form a new N-type channel permitting conduction between source and drain. The V_{DD} source is now able to promote a current through the MOSFET.

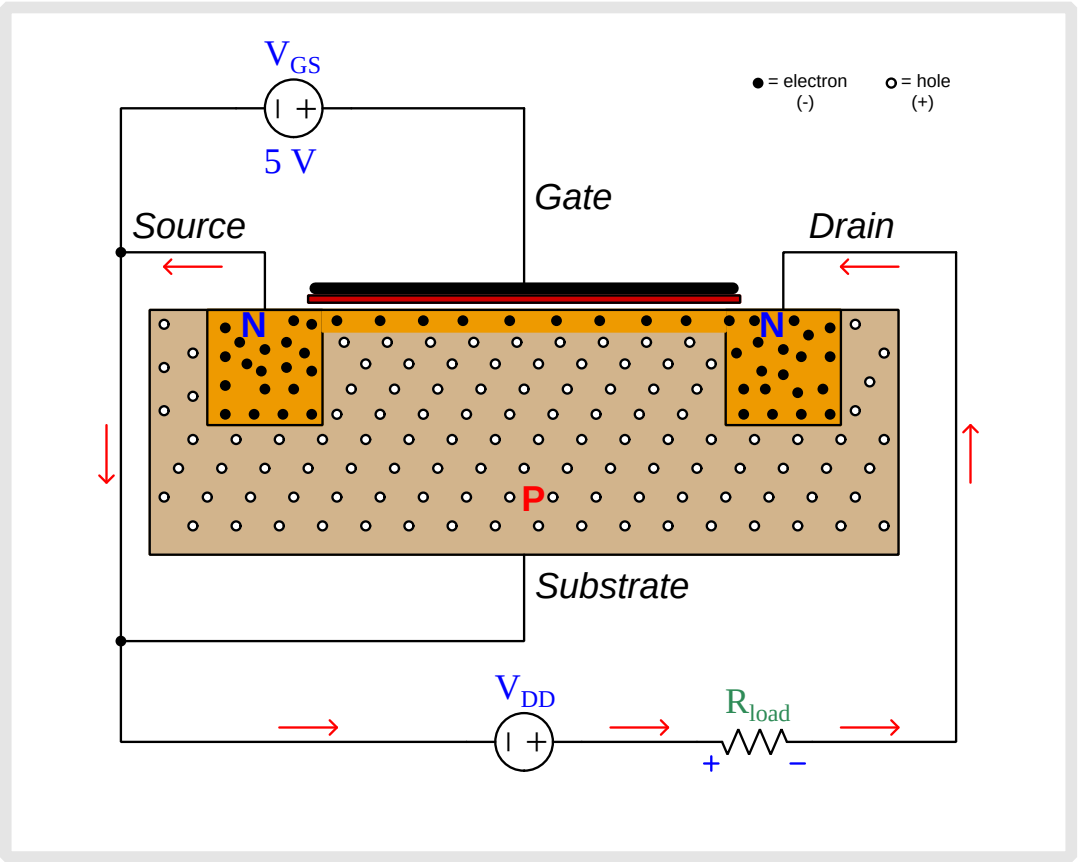

Note how the mobile charge carriers within the inversion layer are electrons, which flow opposite the direction of "conventional flow" as notated by the red arrows. Notice also how new electrons enter the inversion layer from the source region of the MOSFET, which is why that region is so-named (because that is the "source" of fresh charge carriers entering the inversion layer).

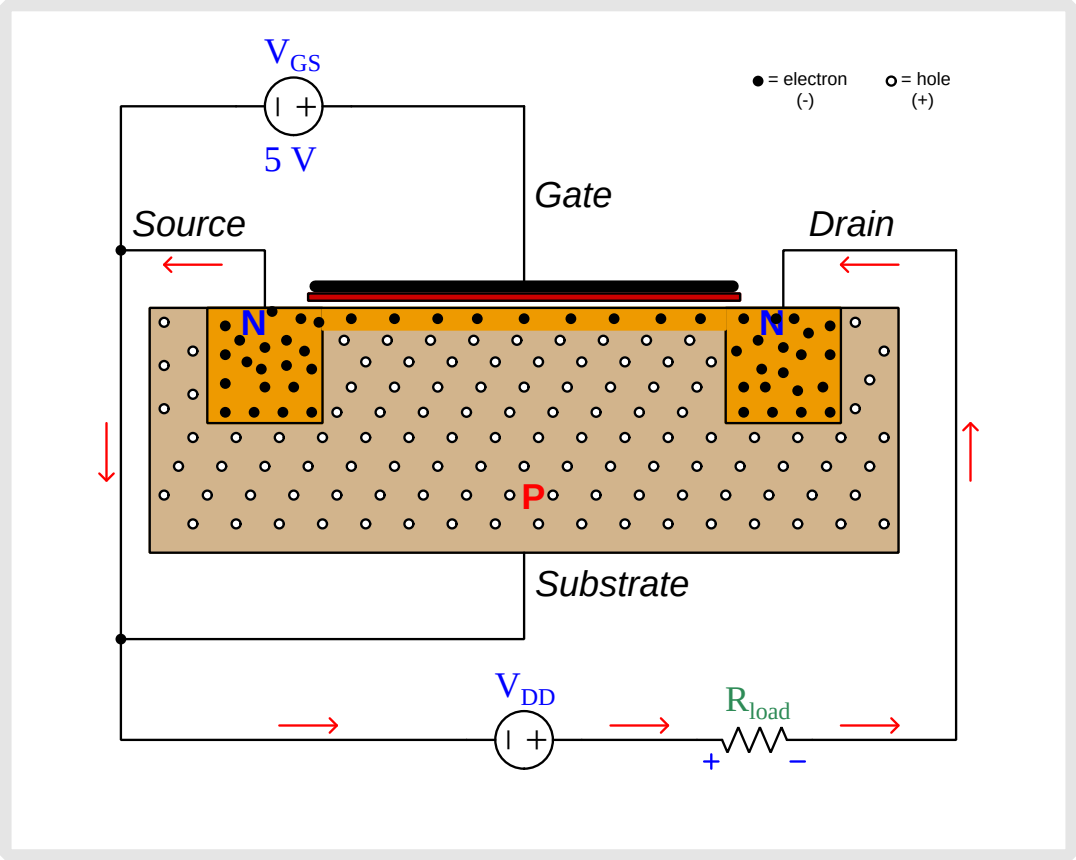

Note how the mobile charge carriers within the inversion layer are electrons, which flow opposite the direction of "conventional flow" as notated by the red arrows. Notice also how new electrons enter the inversion layer from the source region of the MOSFET, which is why that region is so-named (because that is the "source" of fresh charge carriers entering the inversion layer).

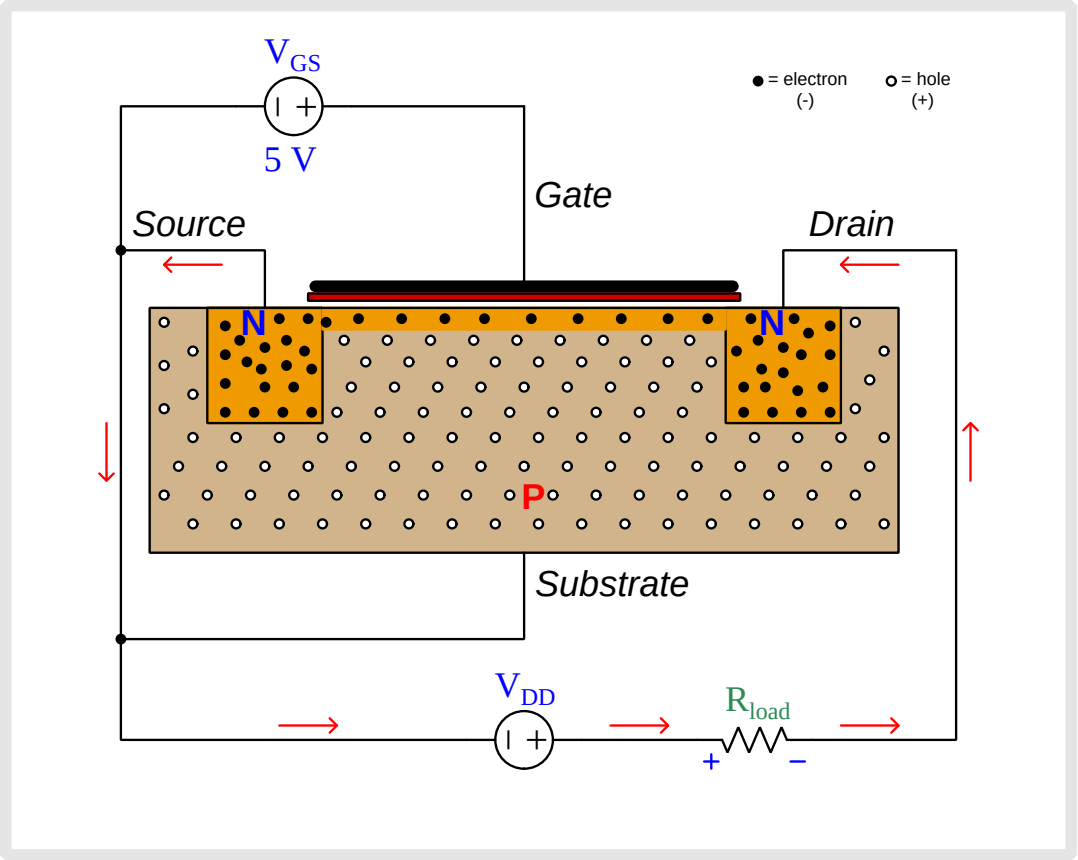

Note how the mobile charge carriers within the inversion layer are electrons, which flow opposite the direction of "conventional flow" as notated by the red arrows. Notice also how new electrons enter the inversion layer from the source region of the MOSFET, which is why that region is so-named (because that is the "source" of fresh charge carriers entering the inversion layer).

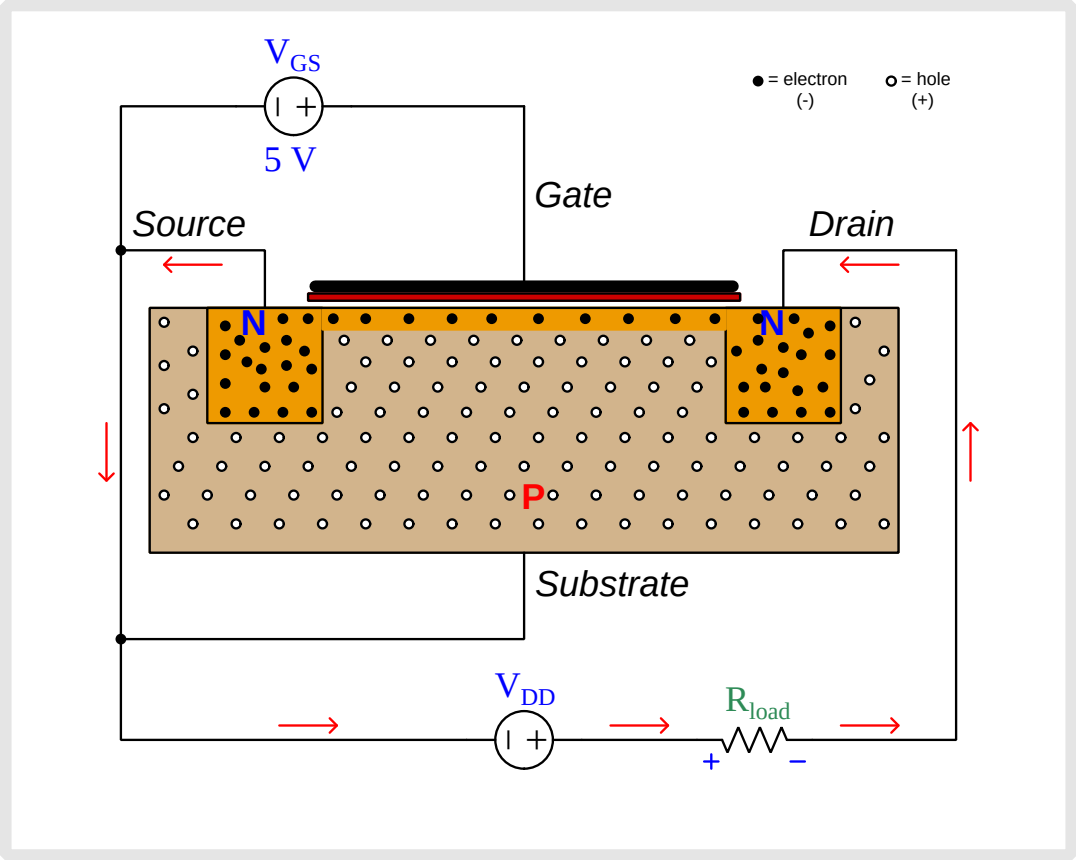

Note how the mobile charge carriers within the inversion layer are electrons, which flow opposite the direction of "conventional flow" as notated by the red arrows. Notice also how new electrons enter the inversion layer from the source region of the MOSFET, which is why that region is so-named (because that is the "source" of fresh charge carriers entering the inversion layer).

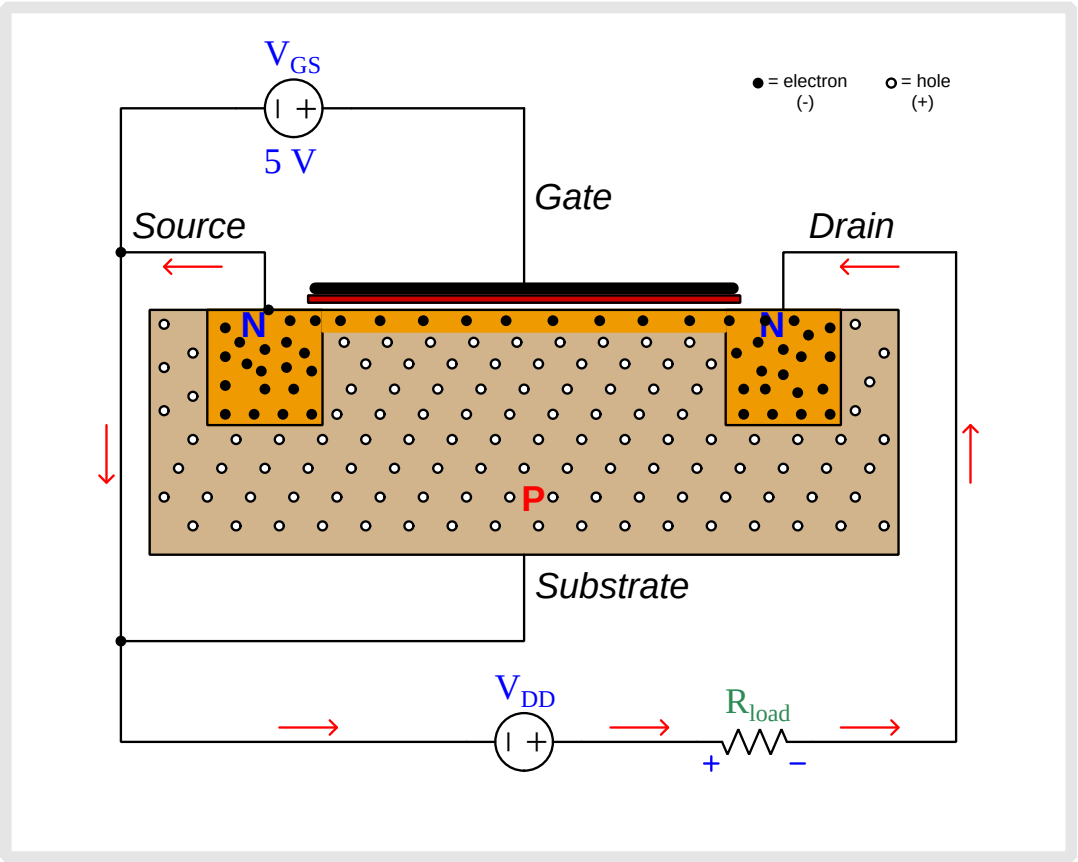

Note how the mobile charge carriers within the inversion layer are electrons, which flow opposite the direction of "conventional flow" as notated by the red arrows. Notice also how new electrons enter the inversion layer from the source region of the MOSFET, which is why that region is so-named (because that is the "source" of fresh charge carriers entering the inversion layer).

Chapter 7

Questions

This learning module, along with all others in the ModEL collection, is designed to be used in an inverted instructional environment where students independently read[1] the tutorials and attempt to answer questions on their own *prior* to the instructor's interaction with them. In place of lecture[2], the instructor engages with students in Socratic-style dialogue, probing and challenging their understanding of the subject matter through inquiry.

Answers are not provided for questions within this chapter, and this is by design. Solved problems may be found in the Tutorial and Derivation chapters, instead. The goal here is *independence*, and this requires students to be challenged in ways where others cannot think for them. Remember that you always have the tools of *experimentation* and *computer simulation* (e.g. SPICE) to explore concepts!

The following lists contain ideas for Socratic-style questions and challenges. Upon inspection, one will notice a strong theme of *metacognition* within these statements: they are designed to foster a regular habit of examining one's own thoughts as a means toward clearer thinking. As such these sample questions are useful both for instructor-led discussions as well as for self-study.

[1] Technical reading is an essential academic skill for any technical practitioner to possess for the simple reason that the most comprehensive, accurate, and useful information to be found for developing technical competence is in textual form. Technical careers in general are characterized by the need for continuous learning to remain current with standards and technology, and therefore any technical practitioner who cannot read well is handicapped in their professional development. An excellent resource for educators on improving students' reading prowess through intentional effort and strategy is the book textitReading For Understanding – How Reading Apprenticeship Improves Disciplinary Learning in Secondary and College Classrooms by Ruth Schoenbach, Cynthia Greenleaf, and Lynn Murphy.

[2] Lecture is popular as a teaching method because it is easy to implement: any reasonably articulate subject matter expert can talk to students, even with little preparation. However, it is also quite problematic. A good lecture always makes complicated concepts seem easier than they are, which is bad for students because it instills a false sense of confidence in their own understanding; reading and re-articulation requires more cognitive effort and serves to verify comprehension. A culture of teaching-by-lecture fosters a debilitating dependence upon direct personal instruction, whereas the challenges of modern life demand independent and critical thought made possible only by gathering information and perspectives from afar. Information presented in a lecture is ephemeral, easily lost to failures of memory and dictation; text is forever, and may be referenced at any time.

General challenges following tutorial reading

- Summarize as much of the text as you can in one paragraph of your own words. A helpful strategy is to explain ideas as you would for an intelligent child: as simple as you can without compromising too much accuracy.
- Simplify a particular section of the text, for example a paragraph or even a single sentence, so as to capture the same fundamental idea in fewer words.
- Where did the text make the most sense to you? What was it about the text's presentation that made it clear?
- Identify where it might be easy for someone to misunderstand the text, and explain why you think it could be confusing.
- Identify any new concept(s) presented in the text, and explain in your own words.
- Identify any familiar concept(s) such as physical laws or principles applied or referenced in the text.
- Devise a proof of concept experiment demonstrating an important principle, physical law, or technical innovation represented in the text.
- Devise an experiment to disprove a plausible misconception.
- Did the text reveal any misconceptions you might have harbored? If so, describe the misconception(s) and the reason(s) why you now know them to be incorrect.
- Describe any useful problem-solving strategies applied in the text.
- Devise a question of your own to challenge a reader's comprehension of the text.

General follow-up challenges for assigned problems

- Identify where any fundamental laws or principles apply to the solution of this problem, especially before applying any mathematical techniques.
- Devise a thought experiment to explore the characteristics of the problem scenario, applying known laws and principles to mentally model its behavior.
- Describe in detail your own strategy for solving this problem. How did you identify and organized the given information? Did you sketch any diagrams to help frame the problem?
- Is there more than one way to solve this problem? Which method seems best to you?
- Show the work you did in solving this problem, even if the solution is incomplete or incorrect.
- What would you say was the most challenging part of this problem, and why was it so?
- Was any important information missing from the problem which you had to research or recall?
- Was there any extraneous information presented within this problem? If so, what was it and why did it not matter?
- Examine someone else's solution to identify where they applied fundamental laws or principles.
- Simplify the problem from its given form and show how to solve this simpler version of it. Examples include eliminating certain variables or conditions, altering values to simpler (usually whole) numbers, applying a limiting case (i.e. altering a variable to some extreme or ultimate value).
- For quantitative problems, identify the real-world meaning of all intermediate calculations: their units of measurement, where they fit into the scenario at hand. Annotate any diagrams or illustrations with these calculated values.
- For quantitative problems, try approaching it qualitatively instead, thinking in terms of "increase" and "decrease" rather than definite values.
- For qualitative problems, try approaching it quantitatively instead, proposing simple numerical values for the variables.
- Were there any assumptions you made while solving this problem? Would your solution change if one of those assumptions were altered?
- Identify where it would be easy for someone to go astray in attempting to solve this problem.
- Formulate your own problem based on what you learned solving this one.

General follow-up challenges for experiments or projects

- In what way(s) was this experiment or project easy to complete?
- Identify some of the challenges you faced in completing this experiment or project.

- Show how thorough documentation assisted in the completion of this experiment or project.
- Which fundamental laws or principles are key to this system's function?
- Identify any way(s) in which one might obtain false or otherwise misleading measurements from test equipment in this system.
- What will happen if (component X) fails (open/shorted/etc.)?
- What would have to occur to make this system unsafe?

7.1 Conceptual reasoning

These questions are designed to stimulate your analytic and synthetic thinking[3]. In a Socratic discussion with your instructor, the goal is for these questions to prompt an extended dialogue where assumptions are revealed, conclusions are tested, and understanding is sharpened. Your instructor may also pose additional questions based on those assigned, in order to further probe and refine your conceptual understanding.

Questions that follow are presented to challenge and probe your understanding of various concepts presented in the tutorial. These questions are intended to serve as a guide for the Socratic dialogue between yourself and the instructor. Your instructor's task is to ensure you have a sound grasp of these concepts, and the questions contained in this document are merely a means to this end. Your instructor may, at his or her discretion, alter or substitute questions for the benefit of tailoring the discussion to each student's needs. The only absolute requirement is that each student is challenged and assessed at a level equal to or greater than that represented by the documented questions.

It is far more important that you convey your *reasoning* than it is to simply convey a correct answer. For this reason, you should refrain from researching other information sources to answer questions. What matters here is that *you* are doing the thinking. If the answer is incorrect, your instructor will work with you to correct it through proper reasoning. A correct answer without an adequate explanation of how you derived that answer is unacceptable, as it does not aid the learning or assessment process.

You will note a conspicuous lack of answers given for these conceptual questions. Unlike standard textbooks where answers to every other question are given somewhere toward the back of the book, here in these learning modules students must rely on other means to check their work. The best way by far is to debate the answers with fellow students and also with the instructor during the Socratic dialogue sessions intended to be used with these learning modules. Reasoning through challenging questions with other people is an excellent tool for developing strong reasoning skills.

Another means of checking your conceptual answers, where applicable, is to use circuit simulation software to explore the effects of changes made to circuits. For example, if one of these conceptual questions challenges you to predict the effects of altering some component parameter in a circuit, you may check the validity of your work by simulating that same parameter change within software and seeing if the results agree.

[3] *Analytical* thinking involves the "disassembly" of an idea into its constituent parts, analogous to dissection. *Synthetic* thinking involves the "assembly" of a new idea comprised of multiple concepts, analogous to construction. Both activities are high-level cognitive skills, extremely important for effective problem-solving, necessitating frequent challenge and regular practice to fully develop.

7.1.1 Reading outline and reflections

"Reading maketh a full man; conference a ready man; and writing an exact man" – Francis Bacon

Francis Bacon's advice is a blueprint for effective education: reading provides the learner with knowledge, writing focuses the learner's thoughts, and critical dialogue equips the learner to confidently communicate and apply their learning. Independent acquisition and application of knowledge is a powerful skill, well worth the effort to cultivate. To this end, students should read these educational resources closely, journal their own reflections on the reading, and discuss in detail their findings with classmates and instructor(s). You should be able to do all of the following after reading any instructional text:

☑ Briefly SUMMARIZE THE TEXT in the form of a journal entry documenting your learning as you progress through the course of study. Share this summary in dialogue with your classmates and instructor. Journaling is an excellent self-test of thorough reading because you cannot clearly express what you have not read or did not comprehend.

☑ Demonstrate ACTIVE READING STRATEGIES, including verbalizing your impressions as you read, simplifying long passages to convey the same ideas using fewer words, annotating text and illustrations with your own interpretations, working through mathematical examples shown in the text, cross-referencing passages with relevant illustrations and/or other passages, identifying problem-solving strategies applied by the author, etc. Technical reading is a special case of problem-solving, and so these strategies work precisely because they help solve any problem: paying attention to your own thoughts (metacognition), eliminating unnecessary complexities, identifying what makes sense, paying close attention to details, drawing connections between separated facts, and noting the successful strategies of others.

☑ Identify IMPORTANT THEMES, especially GENERAL LAWS and PRINCIPLES, expounded in the text and express them in the simplest of terms as though you were teaching an intelligent child. This emphasizes connections between related topics and develops your ability to communicate complex ideas to anyone.

☑ Form YOUR OWN QUESTIONS based on the reading, and then pose them to your instructor and classmates for their consideration. Anticipate both correct and incorrect answers, the incorrect answer(s) assuming one or more plausible misconceptions. This helps you view the subject from different perspectives to grasp it more fully.

☑ Devise EXPERIMENTS to test claims presented in the reading, or to disprove misconceptions. Predict possible outcomes of these experiments, and evaluate their meanings: what result(s) would confirm, and what would constitute disproof? Running mental simulations and evaluating results is essential to scientific and diagnostic reasoning.

☑ Specifically identify any points you found CONFUSING. The reason for doing this is to help diagnose misconceptions and overcome barriers to learning.

7.1.2 Foundational concepts

Correct analysis and diagnosis of electric circuits begins with a proper understanding of some basic concepts. The following is a list of some important concepts referenced in this module's full tutorial. Define each of them in your own words, and be prepared to illustrate each of these concepts with a description of a practical example and/or a live demonstration.

Energy

Conservation of Energy

Annotating diagrams as a problem-solving strategy

Limiting cases as a problem-solving strategy

Electrical source

Electrical load

AC

DC

Linear versus Switching circuits

Diode

Diode principles

Zener diode principles

Capacitance

Inductance

Mutual induction

Transformer

Primary winding

Secondary winding

Efficiency

Second Law of Thermodynamics

Joule's Law

Duty cycle

Pulse-Width Modulation (PWM)

Comparator

Ripple

Majority versus Minority carriers

Sinusoidal decomposition (i.e. Fourier's Theorem)

Harmonic frequency

High-side versus Low-side switching

Faraday's Law of Electromagnetic Induction

Saturation

Filter

Noise

Ground loop

7.1.3 Buck converter regulation

The output voltage of a buck converter is a direct function of the switching transistor's duty cycle. Specifically, $V_{out} = V_{in}\left(\frac{t_{on}}{t_{total}}\right)$. Explain how the following PWM control circuit regulates the output voltage of the buck converter:

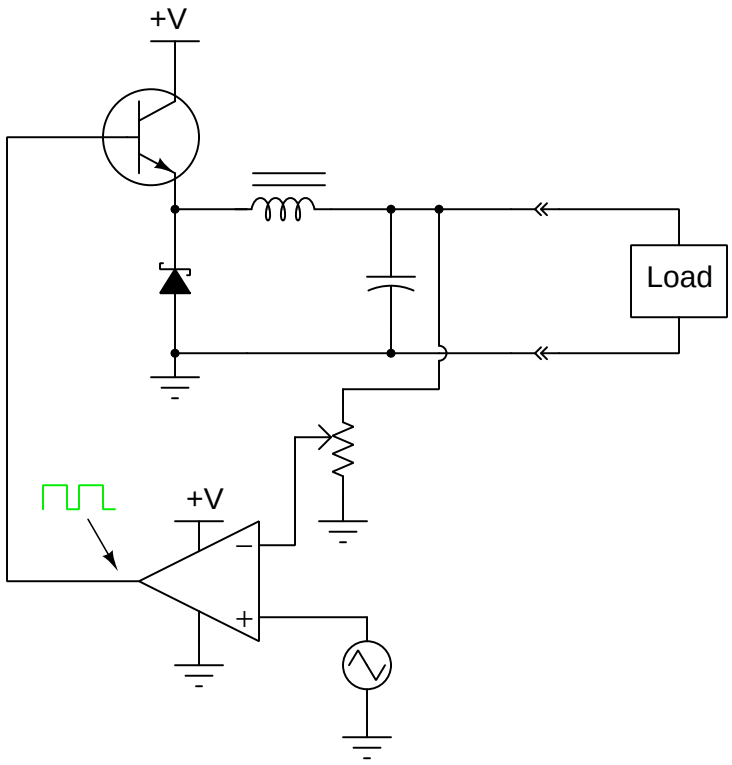

Challenges

- What is the purpose of the potentiometer in this circuit?

7.1.4 24 Volt spotlight

Suppose a friend of yours recently purchased an off-road vehicle. This friend also purchased a military-surplus spotlight, which he thinks would be a great accessory for off-road illumination at night. The only problem is, the spotlight is rated for 24 Volts, while the electrical system in his vehicle is 12 Volt.

Your friend asks you to engineer a solution for powering the 24-Volt spotlight with the 12 Volts available on his vehicle. Of course, you are not allowed to modify the vehicle's electrical system (change it to 24 Volt generator, battery, starter motor, etc.), because it is new and still under warranty. What do you recommend to your friend?

Draw a component-level schematic diagram of your solution to this problem.

Challenges

- Identify some of the information you would need to know about this application before being able to fully specify a design.
- Sketch both synchronous and asynchronous versions of this converter circuit.

7.1.5 Linear vs. switch-mode

Identify which is the linear (brute-force) and which is the switch-mode of these two "wall-wart" power supplies, both outputting DC from an input of AC line voltage:

Challenges

- Identify some of the criteria a designer might consider when choosing a linear power supply versus a switching power supply for their design.

7.1.6 Poor converter efficiency

A student builds a boost converter to step 7 Volts DC up to 10 Volts DC, but when measuring voltage and current for both input and output of this converter in order to calculate P_{in} and P_{out}, they find that the ratio of $\frac{P_{out}}{P_{in}}$ is quite poor: only about 18%. They try swapping different values of load resistance on this circuit, and find that the higher the load resistance, the poorer the energy efficiency.

Explain why this might be, and how it is perfectly normal for a DC-DC converter circuit.

Challenges

- Identify design changes that could improve the energy efficiency of any given DC-DC converter circuit assuming a fixed supply voltage and a fixed load.

7.1.7 Source/load functions in Ćuk converter

Figures 5 and 6 from sheet 2 of Slobodan Ćuk and Robert Middlebrook's US Patent number 4,184,197 ("DC-to-DC Switching Converter") shows the basic topology of this DC-DC converter design:

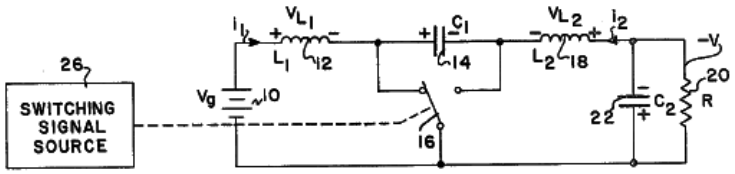

FIG. 5

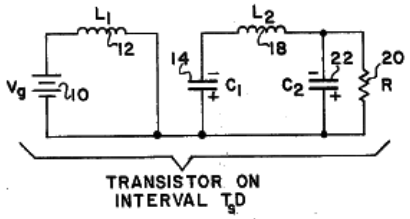

FIG. 5(1)

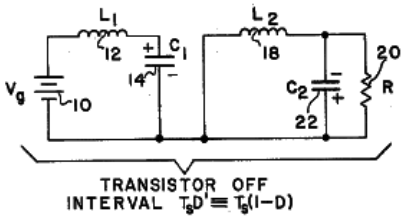

FIG. 5(2)

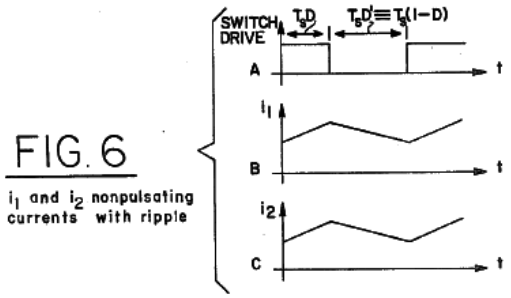

FIG. 6

i_1 and i_2 nonpulsating currents with ripple

Based on the current plots for i_1 and i_2 shown in Figure 6, identify when inductors L_1 and L_2 operate as sources and when they operate as loads.

Challenges

- What is significant about the phrase "i_1 and i_2 nonpulsating currents with ripple" seen in Figure 6?

7.1.8 Transistor biasing

Figure 9 from Slobodan Ćuk and Robert Middlebrook's US Patent number 4,184,197 ("DC-to-DC Switching Converter") shows a set of transistors used to perform power switching in a converter circuit, but the schematic is incomplete:

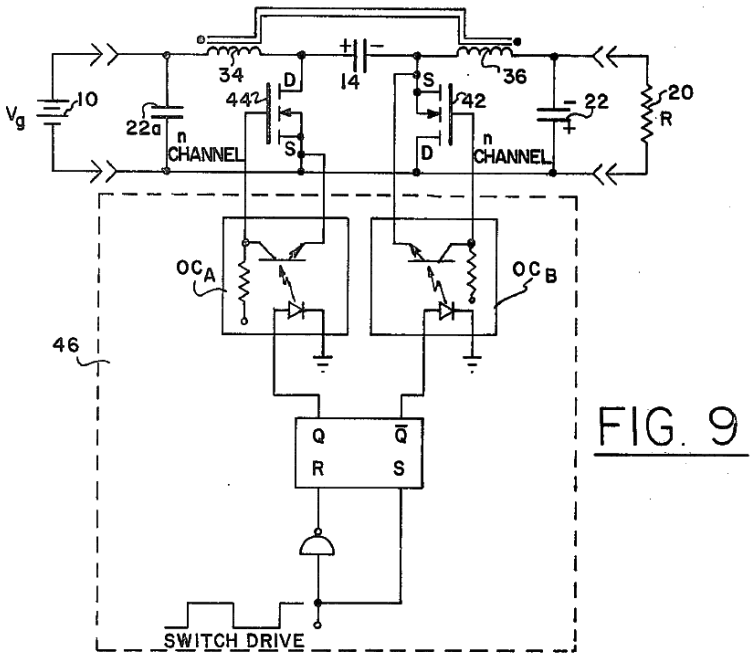

FIG. 9

Each of the optical couplers (OC_A and OC_B) show an unterminated resistor. Where should the unterminated leads of these resistors connect in order to make this a reliably functioning circuit?

Identify the state of each MOSFET when the optocoupler LED is de-energized.

Challenges

- Identify the practical purpose of the SR latch.

7.2 Quantitative reasoning

These questions are designed to stimulate your computational thinking. In a Socratic discussion with your instructor, the goal is for these questions to reveal your mathematical approach(es) to problem-solving so that good technique and sound reasoning may be reinforced. Your instructor may also pose additional questions based on those assigned, in order to observe your problem-solving firsthand.

Mental arithmetic and estimations are strongly encouraged for all calculations, because without these abilities you will be unable to readily detect errors caused by calculator misuse (e.g. keystroke errors).

You will note a conspicuous lack of answers given for these quantitative questions. Unlike standard textbooks where answers to every other question are given somewhere toward the back of the book, here in these learning modules students must rely on other means to check their work. My advice is to use circuit simulation software such as SPICE to check the correctness of quantitative answers. Refer to those learning modules within this collection focusing on SPICE to see worked examples which you may use directly as practice problems for your own study, and/or as templates you may modify to run your own analyses and generate your own practice problems.

Completely worked example problems found in the Tutorial may also serve as "test cases[4]" for gaining proficiency in the use of circuit simulation software, and then once that proficiency is gained you will never need to rely[5] on an answer key!

[4] In other words, set up the circuit simulation software to analyze the same circuit examples found in the Tutorial. If the simulated results match the answers shown in the Tutorial, it confirms the simulation has properly run. If the simulated results disagree with the Tutorial's answers, something has been set up incorrectly in the simulation software. Using every Tutorial as practice in this way will quickly develop proficiency in the use of circuit simulation software.

[5] This approach is perfectly in keeping with the instructional philosophy of these learning modules: *teaching students to be self-sufficient thinkers*. Answer keys can be useful, but it is even more useful to your long-term success to have a set of tools on hand for checking your own work, because once you have left school and are on your own, there will no longer be "answer keys" available for the problems you will have to solve.

7.2.1 Miscellaneous physical constants

Note: constants shown in **bold** type are *exact*, not approximations. Values inside of parentheses show one standard deviation (σ) of uncertainty in the final digits: for example, the magnetic permeability of free space value given as $1.25663706212(19) \times 10^{-6}$ H/m represents a center value (i.e. the location parameter) of $1.25663706212 \times 10^{-6}$ Henrys per meter with one standard deviation of uncertainty equal to $0.0000000000019 \times 10^{-6}$ Henrys per meter.

Avogadro's number (N_A) = **6.02214076** $\times 10^{23}$ **per mole** (mol^{-1})

Boltzmann's constant (k) = **1.380649** $\times 10^{-23}$ **Joules per Kelvin** (J/K)

Electronic charge (e) = **1.602176634** $\times 10^{-19}$ **Coulomb** (C)

Faraday constant (F) = **96,485.33212...** $\times 10^{4}$ **Coulombs per mole** (C/mol)

Magnetic permeability of free space (μ_0) = $1.25663706212(19) \times 10^{-6}$ Henrys per meter (H/m)

Electric permittivity of free space (ϵ_0) = $8.8541878128(13) \times 10^{-12}$ Farads per meter (F/m)

Characteristic impedance of free space (Z_0) = 376.730313668(57) Ohms (Ω)

Gravitational constant (G) = $6.67430(15) \times 10^{-11}$ cubic meters per kilogram-seconds squared ($m^3/kg\text{-}s^2$)

Molar gas constant (R) = **8.314462618... Joules per mole-Kelvin** (J/mol-K) = 0.08205746(14) liters-atmospheres per mole-Kelvin

Planck constant (h) = **6.62607015** $\times 10^{-34}$ **joule-seconds** (J-s)

Stefan-Boltzmann constant (σ) = **5.670374419...** $\times 10^{-8}$ **Watts per square meter-Kelvin**^{4} ($W/m^2 \cdot K^4$)

Speed of light in a vacuum (c) = **299,792,458 meters per second** (m/s) = 186282.4 miles per second (mi/s)

Note: All constants taken from NIST data "Fundamental Physical Constants – Complete Listing", from `http://physics.nist.gov/constants`, National Institute of Standards and Technology (NIST), 2018 CODATA Adjustment.

7.2.2 Introduction to spreadsheets

A powerful computational tool you are encouraged to use in your work is a *spreadsheet*. Available on most personal computers (e.g. Microsoft Excel), *spreadsheet* software performs numerical calculations based on number values and formulae entered into cells of a grid. This grid is typically arranged as lettered columns and numbered rows, with each cell of the grid identified by its column/row coordinates (e.g. cell `B3`, cell `A8`). Each cell may contain a string of text, a number value, or a mathematical formula. The spreadsheet automatically updates the results of all mathematical formulae whenever the entered number values are changed. This means it is possible to set up a spreadsheet to perform a series of calculations on entered data, and those calculations will be re-done by the computer any time the data points are edited in any way.

For example, the following spreadsheet calculates average speed based on entered values of distance traveled and time elapsed:

	A	B	C	D
1	Distance traveled	46.9	Kilometers	
2	Time elapsed	1.18	Hours	
3	Average speed	= B1 / B2	km/h	
4				
5				

Text labels contained in cells `A1` through `A3` and cells `C1` through `C3` exist solely for readability and are not involved in any calculations. Cell `B1` contains a sample distance value while cell `B2` contains a sample time value. The formula for computing speed is contained in cell `B3`. Note how this formula begins with an "equals" symbol (`=`), references the values for distance and speed by lettered column and numbered row coordinates (`B1` and `B2`), and uses a forward slash symbol for division (`/`). The coordinates `B1` and `B2` function as *variables*[6] would in an algebraic formula.

When this spreadsheet is executed, the numerical value `39.74576` will appear in cell `B3` rather than the formula `= B1 / B2`, because 39.74576 is the computed speed value given 46.9 kilometers traveled over a period of 1.18 hours. If a different numerical value for distance is entered into cell `B1` or a different value for time is entered into cell `B2`, cell `B3`'s value will automatically update. All you need to do is set up the given values and any formulae into the spreadsheet, and the computer will do all the calculations for you.

Cell `B3` may be referenced by other formulae in the spreadsheet if desired, since it is a variable just like the given values contained in `B1` and `B2`. This means it is possible to set up an entire chain of calculations, one dependent on the result of another, in order to arrive at a final value. The arrangement of the given data and formulae need not follow any pattern on the grid, which means you may place them anywhere.

[6] Spreadsheets may also provide means to attach text labels to cells for use as variable names (Microsoft Excel simply calls these labels "names"), but for simple spreadsheets such as those shown here it's usually easier just to use the standard coordinate naming for each cell.

Common[7] arithmetic operations available for your use in a spreadsheet include the following:

- Addition (`+`)
- Subtraction (`-`)
- Multiplication (`*`)
- Division (`/`)
- Powers (`^`)
- Square roots (`sqrt()`)
- Logarithms (`ln()` , `log10()`)

Parentheses may be used to ensure[8] proper order of operations within a complex formula. Consider this example of a spreadsheet implementing the *quadratic formula*, used to solve for roots of a polynomial expression in the form of $ax^2 + bx + c$:

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

	A	B
1	x_1	= (-B4 + sqrt((B4^2) - (4*B3*B5))) / (2*B3)
2	x_2	= (-B4 - sqrt((B4^2) - (4*B3*B5))) / (2*B3)
3	a =	9
4	b =	5
5	c =	-2

This example is configured to compute roots[9] of the polynomial $9x^2 + 5x - 2$ because the values of 9, 5, and -2 have been inserted into cells `B3`, `B4`, and `B5`, respectively. Once this spreadsheet has been built, though, it may be used to calculate the roots of *any* second-degree polynomial expression simply by entering the new a, b, and c coefficients into cells `B3` through `B5`. The numerical values appearing in cells `B1` and `B2` will be automatically updated by the computer immediately following any changes made to the coefficients.

[7] Modern spreadsheet software offers a bewildering array of mathematical functions you may use in your computations. I recommend you consult the documentation for your particular spreadsheet for information on operations other than those listed here.

[8] Spreadsheet programs, like text-based programming languages, are designed to follow standard order of operations by default. However, my personal preference is to use parentheses even where strictly unnecessary just to make it clear to any other person viewing the formula what the intended order of operations is.

[9] Reviewing some algebra here, a *root* is a value for x that yields an overall value of zero for the polynomial. For this polynomial ($9x^2 + 5x - 2$) the two roots happen to be $x = 0.269381$ and $x = -0.82494$, with these values displayed in cells `B1` and `B2`, respectively upon execution of the spreadsheet.

Alternatively, one could break up the long quadratic formula into smaller pieces like this:

$$y = \sqrt{b^2 - 4ac} \qquad z = 2a$$

$$x = \frac{-b \pm y}{z}$$

	A	B	C
1	x_1	= (-B4 + C1) / C2	= sqrt((B4^2) - (4*B3*B5))
2	x_2	= (-B4 - C1) / C2	= 2*B3
3	a =	9	
4	b =	5	
5	c =	-2	

Note how the square-root term (y) is calculated in cell C1, and the denominator term (z) in cell C2. This makes the two final formulae (in cells B1 and B2) simpler to interpret. The positioning of all these cells on the grid is completely arbitrary[10] – all that matters is that they properly reference each other in the formulae.

Spreadsheets are particularly useful for situations where the same set of calculations representing a circuit or other system must be repeated for different initial conditions. The power of a spreadsheet is that it automates what would otherwise be a tedious set of calculations. One specific application of this is to simulate the effects of various components within a circuit failing with abnormal values (e.g. a shorted resistor simulated by making its value nearly zero; an open resistor simulated by making its value extremely large). Another application is analyzing the behavior of a circuit design given new components that are out of specification, and/or aging components experiencing drift over time.

[10] My personal preference is to locate all the "given" data in the upper-left cells of the spreadsheet grid (each data point flanked by a sensible name in the cell to the left and units of measurement in the cell to the right as illustrated in the first distance/time spreadsheet example), sometimes coloring them in order to clearly distinguish which cells contain entered data versus which cells contain computed results from formulae. I like to place all formulae in cells below the given data, and try to arrange them in logical order so that anyone examining my spreadsheet will be able to figure out *how* I constructed a solution. This is a general principle I believe all computer programmers should follow: *document and arrange your code to make it easy for other people to learn from it.*

7.2.3 Converter efficiency calculations

A "buck" switching converter operating at 85% efficiency delivers 10 Amperes of current at 5 Volts to a DC load. Calculate the input current if the input voltage is 23 Volts, as well as the amount of power dissipated as heat.

A "boost" switching converter operating at 90% efficiency delivers 50 Volts to a DC load. Calculate the load current if the input voltage is 17 Volts and the input current is 9.3 Amperes, as well as the amount of power dissipated as heat.

A "boost" switching converter operating at 80% efficiency delivers 178 Volts at 1 Ampere to a DC load. Calculate the input voltage if the input current is 11 Amperes, as well as the amount of power dissipated as heat.

A "buck" switching converter delivers 36 Volts at 2.5 Amperes to a DC load. Calculate the energy efficiency of the converter if the input current is 0.8 Amperes and the input voltage is 133 Volts, as well as the amount of power dissipated as heat.

Finally, show how the following block diagram is helpful for tracking the real-world meaning of all given and calculated values for each of these converter scenarios:

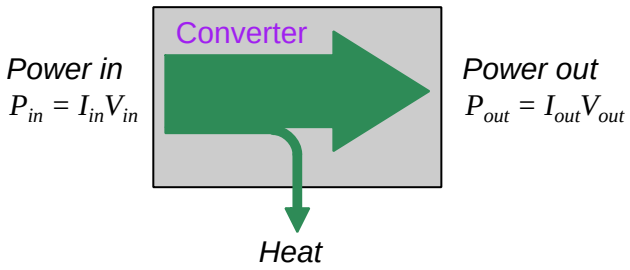

Challenges

- Identify any switched converter styles capable of both boost and buck.
- Identify how the buck converter calculations would differ if it were a linear converter rather than a switching converter.

7.2.4 Solving for duty cycle

The following equations solve for the output voltage of various switching converter circuits, given the switch duty cycle D and the input voltage:

$$V_{out} = D\,V_{in} \qquad \text{(Buck converter circuit)}$$

$$V_{out} = \frac{V_{in}}{1 - D} \qquad \text{(Boost converter circuit)}$$

$$V_{out} = \frac{D\,V_{in}}{1 - D} \qquad \text{(Inverting or Cuk converter circuit)}$$

Manipulate each of these equations to solve for duty cycle (D) in terms of the input voltage (V_{in}) and desired output voltage (V_{out}). Remember that duty cycle is always a quantity between 0 and 1, inclusive.

Challenges

- Identify any factor that could cause a converter's output voltage to deviate from its predicted value given V_{in} and duty cycle.

7.2.5 Dying battery source

Suppose a voltage-regulating boost converter takes energy from a chemical battery to drive a load at a regulated voltage greater than that of the battery. Normally the battery outputs 6 Volts and the load requires a constant 9.5 Volts, hence the need for this converter to be voltage-regulated.

Calculate the battery's current delivered to the converter if the load resistance is 8 Ω and the converter's efficiency is 87%.

Suppose the battery begins to die, and its voltage sags from 6 Volts down to 5.5 Volts, 5 Volts, and 4.5 Volts over time. Calculate the battery current at each of those successively lower voltages.

Based on your current calculations for the dying battery, how do you think this system will react to a battery as it gradually depletes of all stored energy?

Suppose the 6-Volt battery has an Ampere-Hour rating of *40,000 milliAmpere-hours*. How long would it be able to sustain its proper output voltage while powering a load drawing a constant current of 3 Amperes?

Challenges

- Would a forward converter be a better choice for a battery-powered switching supply, since it injects a small amount of recharging current back through the battery at every PWM cycle?

7.2.6 Instantaneous analyses of DC-DC converter circuits

Apply foundational principles of circuits (e.g. Ohm's Law, Kirchhoff's Laws, etc.) to determine unknown voltages and currents in these switch-mode converter circuits, assuming each of these circuits has been operating normally for some time prior to the specified moment.

Circuit #1

Determine all component voltages and currents at an instant in time when I_L is 2.5 Amperes, V_C is 10.2 Volts, and the transistor is off:

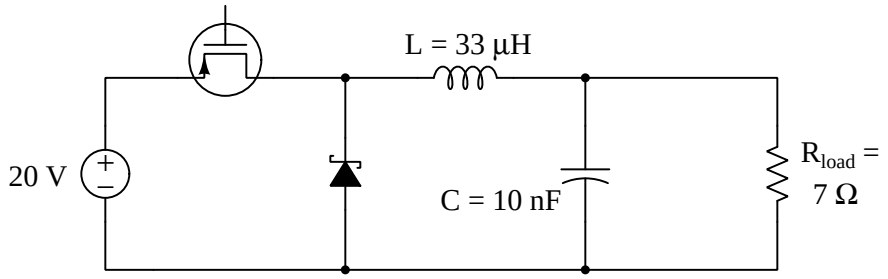

Circuit #2

Determine all component voltages and currents at an instant in time when I_L is 1.7 Amperes, V_C is 8.5 Volts, and the transistor is on:

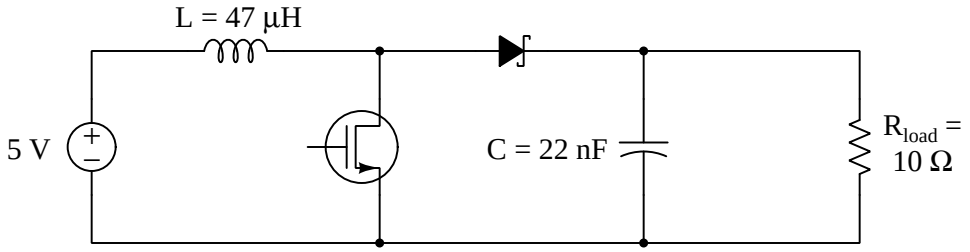

Circuit #3

Determine all component voltages and currents at an instant in time when I_Q is 2.1 Amperes (transistor on):

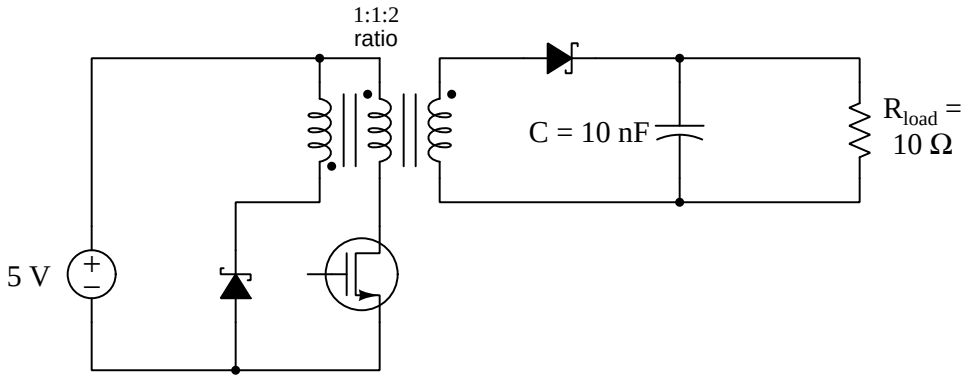

Challenges

- Assuming the use of MOSFETs, how might you incorporate $R_{DS(on)}$ into these calculations?

7.3 Diagnostic reasoning

These questions are designed to stimulate your deductive and inductive thinking, where you must apply general principles to specific scenarios (deductive) and also derive conclusions about the failed circuit from specific details (inductive). In a Socratic discussion with your instructor, the goal is for these questions to reinforce your recall and use of general circuit principles and also challenge your ability to integrate multiple symptoms into a sensible explanation of what's wrong in a circuit. Your instructor may also pose additional questions based on those assigned, in order to further challenge and sharpen your diagnostic abilities.

As always, your goal is to fully *explain* your analysis of each problem. Simply obtaining a correct answer is not good enough – you must also demonstrate sound reasoning in order to successfully complete the assignment. Your instructor's responsibility is to probe and challenge your understanding of the relevant principles and analytical processes in order to ensure you have a strong foundation upon which to build further understanding.

You will note a conspicuous lack of answers given for these diagnostic questions. Unlike standard textbooks where answers to every other question are given somewhere toward the back of the book, here in these learning modules students must rely on other means to check their work. The best way by far is to debate the answers with fellow students and also with the instructor during the Socratic dialogue sessions intended to be used with these learning modules. Reasoning through challenging questions with other people is an excellent tool for developing strong reasoning skills.

Another means of checking your diagnostic answers, where applicable, is to use circuit simulation software to explore the effects of faults placed in circuits. For example, if one of these diagnostic questions requires that you predict the effect of an open or a short in a circuit, you may check the validity of your work by simulating that same fault (substituting a very high resistance in place of that component for an open, and substituting a very low resistance for a short) within software and seeing if the results agree.

7.3.1 Faults in a buck converter

Predict how the operation of this buck converter circuit will be affected as a result of the following faults. Consider each fault independently (i.e. one at a time, no multiple faults):

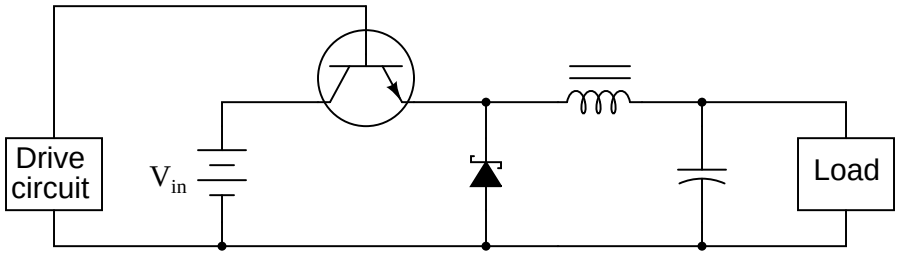

- Drive circuit fails with a constant "low" (0 Volts) output signal
- Drive circuit fails with a constant "high" (+V) output signal
- Diode fails shorted
- Inductor fails open
- Capacitor fails shorted

For each of these conditions, explain *why* the resulting effects will occur.

Challenges

- Identify the drive signal voltage requirements. As it pulses high and low, how high must it go, and how low must it go?
- For each fault condition, identify the effect on the converter's energy efficiency.
- For each fault condition, identify the pathway(s) of any fault current(s).

7.3.2 Faults in a boost converter

Predict how the operation of this boost converter circuit will be affected as a result of the following faults. Consider each fault independently (i.e. one at a time, no multiple faults):

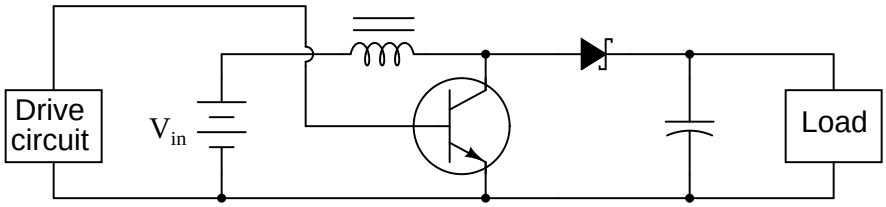

- Drive circuit fails with a constant "low" (0 Volts) output signal
- Drive circuit fails with a constant "high" (+V) output signal
- Diode fails shorted
- Inductor fails open
- Capacitor fails shorted

For each of these conditions, explain *why* the resulting effects will occur.

Challenges

- Identify the drive signal voltage requirements. As it pulses high and low, how high must it go, and how low must it go?
- For each fault condition, identify the effect on the converter's energy efficiency.
- For each fault condition, identify the pathway(s) of any fault current(s).

7.3.3 Identifying possible root-cause faults

Suppose that after months of faithful service, the transistor in this DC-DC converter circuit happens to fail:

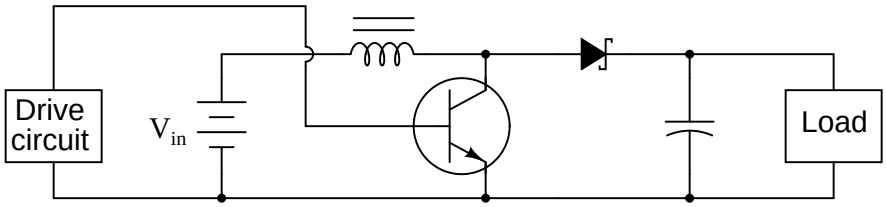

A technician de-solders the failed transistor from the printed circuit board, then re-solders a new one in its place. However, the new transistor fails immediately when the converter is energized. Clearly, the failed transistor is not the *root cause*, but something else is wrong which is causing the transistor to fail.

Identify possible root-cause faults which could explain why this transistor fails prematurely.

Challenges

- How might we have determined the transistor was not the root cause without going through the trouble of replacing it?

Appendix A

Problem-Solving Strategies

The ability to solve complex problems is arguably one of the most valuable skills one can possess, and this skill is particularly important in any science-based discipline.

- Study principles, not procedures. Don't be satisfied with merely knowing how to compute solutions – learn *why* those solutions work.
- Identify what it is you need to solve, identify all relevant data, identify all units of measurement, identify any general principles or formulae linking the given information to the solution, and then identify any "missing pieces" to a solution. Annotate all diagrams with this data.
- Sketch a diagram to help visualize the problem. When building a real system, always devise a plan for that system and analyze its function *before* constructing it.
- Follow the units of measurement and meaning of every calculation. If you are ever performing mathematical calculations as part of a problem-solving procedure, and you find yourself unable to apply each and every intermediate result to some aspect of the problem, it means you don't understand what you are doing. Properly done, every mathematical result should have practical meaning for the problem, and not just be an abstract number. You should be able to identify the proper units of measurement for each and every calculated result, and show where that result fits into the problem.
- Perform "thought experiments" to explore the effects of different conditions for theoretical problems. When troubleshooting real systems, perform *diagnostic tests* rather than visually inspecting for faults, the best diagnostic test being the one giving you the most information about the nature and/or location of the fault with the fewest steps.
- Simplify the problem until the solution becomes obvious, and then use that obvious case as a model to follow in solving the more complex version of the problem.
- Check for exceptions to see if your solution is incorrect or incomplete. A good solution will work for *all* known conditions and criteria. A good example of this is the process of testing scientific hypotheses: the task of a scientist is not to find support for a new idea, but rather to *challenge* that new idea to see if it holds up under a battery of tests. The philosophical

principle of *reductio ad absurdum* (i.e. disproving a general idea by finding a specific case where it fails) is useful here.

- Work "backward" from a hypothetical solution to a new set of given conditions.
- Add quantities to problems that are qualitative in nature, because sometimes a little math helps illuminate the scenario.
- Sketch graphs illustrating how variables relate to each other. These may be quantitative (i.e. with realistic number values) or qualitative (i.e. simply showing increases and decreases).
- Treat quantitative problems as qualitative in order to discern the relative magnitudes and/or directions of change of the relevant variables. For example, try determining what happens if a certain variable were to increase or decrease before attempting to precisely calculate quantities: how will each of the dependent variables respond, by increasing, decreasing, or remaining the same as before?
- Consider limiting cases. This works especially well for qualitative problems where you need to determine which direction a variable will change. Take the given condition and magnify that condition to an extreme degree as a way of simplifying the direction of the system's response.
- Check your work. This means regularly testing your conclusions to see if they make sense. This does *not* mean repeating the same steps originally used to obtain the conclusion(s), but rather to use some other means to check validity. Simply repeating procedures often leads to *repeating the same errors* if any were made, which is why alternative paths are better.

Appendix B

Instructional philosophy

B.1 First principles of learning

- **Anyone can learn anything** given appropriate time, effort, resources, challenges, encouragement, and expectations. Dedicating time and investing effort are the student's responsibility; providing resources, challenges, and encouragement are the teacher's responsibility; high expectations are a responsibility shared by both student and teacher.

- **Transfer is not automatic.** The human mind has a natural tendency to compartmentalize information, which means the process of taking knowledge learned in one context and applying it to another usually does not come easy and therefore should never be taken for granted.

- **Learning is iterative.** The human mind rarely learns anything perfectly on the first attempt. Anticipate mistakes and plan for multiple tries to achieve full understanding, using the lessons of those mistakes as feedback to guide future attempts.

- **Information is absorbed, but understanding is created.** Facts and procedures may be memorized easily enough by repeated exposure, but the ability to reliably apply principles to novel scenarios only comes through intense personal effort. This effort is fundamentally creative in nature: casting new concepts in one's own words, running mental experiments to perceive patterns and trends, etc. These acts of making knowledge "one's own" need not be perfect in order to be effective, as the value lies in the activity and not necessarily the finished product. Depth of understanding and permanence of memory are both directly related to mental effort during the learning process – i.e. *easy learning is shallow and fleeting.*

- **Education trumps training.** There is no such thing as an entirely isolated subject, as all fields of knowledge are connected. Training is narrowly-focused and task-oriented. Education is broad-based and principle-oriented. When preparing for a life-long technical career, education beats training every time.

- **Character matters.** Poor habits are more destructive than deficits of knowledge or skill. This is especially true in collective endeavors, where a team's ability to function depends on trust between its members. Simply put, no one wants an untrustworthy person on their team. An essential component of education then, is character development.

- **People learn to be responsible by bearing responsibility.** An irresponsible person is someone who has never *had* to be responsible for anything that mattered enough to them. Just as anyone can learn anything, anyone can become responsible if the personal cost of irresponsibility becomes high enough.

- **What gets measured, gets done.** Accurate and relevant assessment of learning is key to ensuring all students learn. Therefore, it is imperative to measure what matters.

- **Failure is nothing to fear.** Every human being fails, and fails in multiple ways at multiple times. Eventual success only happens when we don't stop trying.

B.2 Proven strategies for instructors

- Assume every student is capable of learning anything they desire given the proper conditions. Treat them as capable adults by granting real responsibility and avoiding artificial incentives such as merit or demerit points.
- Create a consistent culture of high expectations across the entire program of study. Demonstrate and encourage patience, persistence, and a healthy sense of self-skepticism. Anticipate and de-stigmatize error. Teach respect for the capabilities of others as well as respect for one's own fallibility.
- Replace lecture with "inverted" instruction, where students first encounter new concepts through reading and then spend class time in Socratic dialogue with the instructor exploring those concepts and solving problems individually. There is a world of difference between observing someone solve a problem versus actually solving a problem yourself, and so the point of this form of instruction is to place students in a position where they *cannot* passively observe.
- Require students to read extensively, write about what they learn, and dialogue with you and their peers to sharpen their understanding. Apply Francis Bacon's advice that "reading maketh a full man; conference a ready man; and writing an exact man". These are complementary activities helping students expand their confidence and abilities.
- Use artificial intelligence (AI) to challenge student understanding rather than merely provide information. Find productive ways for AI to critique students' clarity of thought and of expression, for example by employing AI as a Socratic-style interlocutor or as a reviewer of students' journals. Properly applied, AI has the ability to expand student access to critical review well outside the bounds of their instructor's reach.
- Build frequent and rapid feedback into the learning process so that students know at all times how well they are learning, to identify problems early and fix them before they grow. Model the intellectual habit of self-assessing and self-correcting your own understanding (i.e. a cognitive *feedback loop*), encouraging students to do the same.
- Use "mastery" as the standard for every assessment, which means the exam or experiment or project must be done with 100% competence in order to pass. Provide students with multiple opportunity for re-tries (different versions of the assessment every time).
- Require students to devise their own hypotheses and procedures on all experiments, so that the process is truly a scientific one. Have students assess their proposed experimental procedures for risk and devise mitigations for those risks. Let nothing be pre-designed about students' experiments other than a stated task (i.e. what principle the experiment shall test) at the start and a set of demonstrable knowledge and skill objectives at the end.
- Have students build as much of their lab equipment as possible: building power sources, building test assemblies[1], and building complete working systems (no kits!). In order to provide

[1] In the program I teach, every student builds their own "Development Board" consisting of a metal chassis with DIN rail, terminal blocks, and an AC-DC power supply of their own making which functions as a portable lab environment they can use at school as well as take home.

this same "ground-up" experience for every new student, this means either previous students take their creations with them, or the systems get disassembled in preparation for the new students, or the systems grow and evolve with each new student group.

- Incorporate external accountability for you and for your students, continuously improving the curriculum and your instructional methods based on proven results. Have students regularly network with active professionals through participation in advisory committee meetings, service projects, tours, jobshadows, internships, etc. Practical suggestions include requiring students to design and build projects for external clients (e.g. community groups, businesses, different departments within the institution), and also requiring students attend all technical advisory committee meetings and to dialogue with the industry representatives at those meetings.
- Repeatedly explore difficult-to-learn concepts across multiple courses, so that students have multiple opportunities to build their understanding.
- Relate all new concepts, whenever possible, to previous concepts and to relevant physical laws. Challenge each and every student, every day, to *reason* from concept to concept and to explain the logical connections between. Challenge students to verify their conclusions by multiple approaches (e.g. double-checking their work using different methods). Ask *"Why?"* often.
- Maintain detailed records on each student's performance and share these records privately with them. These records should include academic performance as well as professionally relevant behavioral tendencies.
- Hold mandatory "check-in" meetings between all program faculty and each new student during their first term. Offer these to all other students as an option, except for any students continuing to manifest unprofessional behaviors, poor academic performance, or who have some other need for a face-to-face meeting with faculty.
- Address problems while they are small, before they grow larger. This is equally true for tutoring technical concepts as it is for helping students build professional habits.
- Build rigorous quality control into the curriculum to ensure every student masters every important concept, and that the mastery is retained over time. This includes (1) review questions added to every exam to re-assess knowledge taught in previous terms, (2) cumulative exams at the end of every term to re-assess all important concepts back to the very beginning of the program, and (3) review assessments in practical (hands-on) coursework to ensure critically-important skills were indeed taught and are still retained. What you will find by doing this is that it actually boosts retention of students by ensuring that important knowledge gets taught and is retained over long spans of time. In the absence of such quality control, student learning and retention tends to be spotty and this contributes to drop-out and failure rates later in their education.
- Finally, *never rush learning.* Education is not a race. Give your students ample time to digest complex ideas, as you continually remind yourself of just how long it took you to achieve mastery! Long-term retention and the consistently correct application of concepts are always the result of *focused effort over long periods of time* which means there are no shortcuts to learning.

B.3 Proven strategies for students

The single most important piece of advice I have for any student of any subject is to take responsibility for your own development in all areas of life including mental development. Expecting others in your life to entirely guide your own development is a recipe for disappointment. This is just as true for students enrolled in formal learning institutions as it is for auto-didacts pursuing learning entirely on their own. Learning to think in new ways is key to being able to gainfully use information, to make informed decisions about your life, and to best serve those you care about. With this in mind, I offer the following advice to students:

- **Approach all learning as valuable**. No matter what course you take, no matter who you learn from, no matter the subject, there is something useful in every learning experience. If you don't see the value of every new experience, you are not looking closely enough!
- **Continually challenge yourself**. Let other people take shortcuts and find easy answers to easy problems. The purpose of education is to stretch your mind, in order to shape it into a more powerful tool. This doesn't come by taking the path of least resistance. An excellent analogy for an empowering education is productive physical exercise: becoming stronger, more flexible, and more persistent only comes through intense personal effort.
- **Master the use of language**. This includes reading extensively, writing every day, listening closely, and speaking articulately. To a great extent language channels and empowers thought, so the better you are at wielding language the better you will be at grasping abstract concepts and articulating them not only for your benefit but for others as well.
- **Do not limit yourself to the resources given to you**. Read books that are not on the reading list. Run experiments that aren't assigned to you. Form study groups outside of class. Take an entrepreneurial approach to your own education, as though it were a business you were building for your future benefit.
- **Express and share what you learn**. Take every opportunity to teach what you have learned to others, as this will not only help them but will also strengthen your own understanding[2].
- Realize that **no one can give you understanding**, just as no one can give you physical fitness. These both must be *built*.
- **Above all, recognize that learning is hard work, and that a certain level of frustration is unavoidable**. There are times when you will struggle to grasp some of these concepts, and that struggle is a natural thing. Take heart that it will yield with persistent and varied[3] effort, and never give up! That concepts don't immediately come to you is not a sign of something wrong, but rather of something right: that you have found a worthy challenge!

[2] On a personal note, I was surprised to learn just how much my own understanding of electronics and related subjects was strengthened by becoming a teacher. When you are tasked every day with helping other people grasp complex topics, it catalyzes your own learning by giving you powerful incentives to study, to articulate your thoughts, and to reflect deeply on the process of learning.

[3] As the old saying goes, "Insanity is trying the same thing over and over again, expecting different results." If you find yourself stumped by something in the text, you should attempt a different approach. Alter the thought experiment, change the mathematical parameters, do whatever you can to see the problem in a slightly different light, and then the solution will often present itself more readily.

B.4 Design of these learning modules

"The unexamined circuit is not worth energizing" – Socrates (if he had taught electricity)

These learning modules, although useful for self-study, were designed to be used in a formal learning environment where a subject-matter expert challenges students to digest the content and exercise their critical thinking abilities in the answering of questions and in the construction and testing of working circuits. Every effort has been made to embed the following instructional and assessment philosophies within:

- The first goal of education is to enhance clear and independent thought, in order that every student reach their fullest potential in a highly complex and inter-dependent world. Robust reasoning is *always* more important than particulars of any subject matter, because its application is universal.
- Literacy is fundamental to independent learning and thought because text continues to be the most efficient way to communicate complex ideas over space and time. Those who cannot read with ease are limited in their ability to acquire knowledge and perspective.
- Articulate communication is fundamental to work that is complex and interdisciplinary.
- Faulty assumptions and poor reasoning are best corrected through challenge, not presentation. The rhetorical technique of *reductio ad absurdum* (disproving an assertion by exposing an absurdity) works well to discipline student's minds, not only to correct the problem at hand but also to learn how to detect and correct future errors.
- Important principles should be repeatedly explored and widely applied throughout a course of study, not only to reinforce their importance and help ensure their mastery, but also to showcase the interconnectedness and utility of knowledge.

These learning modules were expressly designed to be used in an "inverted" teaching environment[4] where students first read the introductory and tutorial chapters on their own, then individually attempt to answer the questions and construct working circuits according to the experiment and project guidelines. The instructor never lectures, but instead meets regularly with each individual student to review their progress, answer questions, identify misconceptions, and challenge the student to new depths of understanding through further questioning. Regular meetings between instructor and student should resemble a Socratic[5] dialogue, where questions serve as scalpels to dissect topics and expose assumptions. The student passes each module only after consistently demonstrating their ability to logically analyze and correctly apply all major concepts in each question or project/experiment. The instructor must be vigilant in probing each student's understanding to ensure they are truly *reasoning* and not just *memorizing*. This is why "Challenge" points appear throughout, as prompts for students to think deeper about topics and as starting points for instructor queries. Sometimes these challenge points require additional knowledge that hasn't been covered in the series to answer in full. This is okay, as the major purpose of the Challenges is to stimulate analysis and synthesis on the part of each student.

The instructor must possess enough mastery of the subject matter and awareness of students' reasoning to generate their own follow-up questions to practically any student response. Even completely correct answers given by the student should be challenged by the instructor for the purpose of having students practice articulating their thoughts and defending their reasoning. Conceptual errors committed by the student should be exposed and corrected not by direct instruction, but rather by reducing the errors to an absurdity[6] through well-chosen questions and thought experiments posed by the instructor. Becoming proficient at this style of instruction requires time and dedication, but the positive effects on critical thinking for both student and instructor are spectacular.

An inspection of these learning modules reveals certain unique characteristics. One of these is a bias toward thorough explanations in the tutorial chapters. Without a live instructor to explain concepts and applications to students, the text itself must fulfill this role. This philosophy results in lengthier explanations than what you might typically find in a textbook, each step of the reasoning process fully explained, including footnotes addressing common questions and concerns students raise while learning these concepts. Each tutorial seeks to not only explain each major concept in sufficient detail, but also to explain the logic of each concept and how each may be developed

[4] In a traditional teaching environment, students first encounter new information via *lecture* from an expert, and then independently apply that information via *homework*. In an "inverted" course of study, students first encounter new information via *homework*, and then independently apply that information under the scrutiny of an expert. The expert's role in lecture is to simply *explain*, but the expert's role in an inverted session is to *challenge*, *critique*, and if necessary *explain* where gaps in understanding still exist.

[5] Socrates is a figure in ancient Greek philosophy famous for his unflinching style of questioning. Although he authored no texts, he appears as a character in Plato's many writings. The essence of Socratic philosophy is to leave no question unexamined and no point of view unchallenged. While purists may argue a topic such as electric circuits is too narrow for a true Socratic-style dialogue, I would argue that the essential thought processes involved with scientific reasoning on *any* topic are not far removed from the Socratic ideal, and that students of electricity and electronics would do very well to challenge assumptions, pose thought experiments, identify fallacies, and otherwise employ the arsenal of critical thinking skills modeled by Socrates.

[6] This rhetorical technique is known by the Latin phrase *reductio ad absurdum*. The concept is to expose errors by counter-example, since only one solid counter-example is necessary to disprove a universal claim. As an example of this, consider the common misconception among beginning students of electricity that voltage cannot exist without current. One way to apply *reductio ad absurdum* to this statement is to ask how much current passes through a fully-charged battery connected to nothing (i.e. a clear example of voltage existing without current).

from "first principles". Again, this reflects the goal of developing clear and independent thought in students' minds, by showing how clear and logical thought was used to forge each concept. Students benefit from witnessing a model of clear thinking in action, and these tutorials strive to be just that.

Another feature of these learning modules is that they do not treat topics in isolation. Rather, important concepts are introduced early in the series, and appear repeatedly as stepping-stones toward other concepts in subsequent modules. This helps to avoid the "compartmentalization" of knowledge, demonstrating the inter-connectedness of concepts and simultaneously reinforcing them. Each module is fairly complete in itself, reserving the beginning of its tutorial to a review of foundational concepts.

To high standards of education,

Tony R. Kuphaldt

Appendix C

Tools used

I am indebted to the developers of many open-source software applications in the creation of these learning modules. The following is a list of these applications with some commentary on each.

You will notice a theme common to many of these applications: a bias toward *code*. Although I am by no means an expert programmer in any computer language, I understand and appreciate the flexibility offered by code-based applications where the user (you) enters commands into a plain ASCII text file, which the software then reads and processes to create the final output. Code-based computer applications are by their very nature *extensible*, while WYSIWYG (What You See Is What You Get) applications are generally limited to whatever user interface the developer makes for you.

The `GNU/Linux` computer operating system

> There is so much to be said about Linus Torvalds' `Linux` and Richard Stallman's `GNU` project. First, to credit just these two individuals is to fail to do justice to the *mob* of passionate volunteers who contributed to make this amazing software a reality. I first learned of `Linux` back in 1996, and have been using this operating system on my personal computers almost exclusively since then. It is *free*, it is completely *configurable*, and it permits the continued use of highly efficient `Unix` applications and scripting languages (e.g. shell scripts, Makefiles, `sed`, `awk`) developed over many decades. `Linux` not only provided me with a powerful computing platform, but its open design served to inspire my life's work of creating open-source educational resources.

Bram Moolenaar's `Vim` text editor

> Writing code for any code-based computer application requires a *text editor*, which may be thought of as a word processor strictly limited to outputting plain-ASCII text files. Many good text editors exist, and one's choice of text editor seems to be a deeply personal matter within the programming world. I prefer `Vim` because it operates very similarly to `vi` which is ubiquitous on `Unix`/`Linux` operating systems, and because it may be entirely operated via keyboard (i.e. no mouse required) which makes it fast to use.

Donald Knuth's TEX typesetting system

Developed in the late 1970's and early 1980's by computer scientist extraordinaire Donald Knuth to typeset his multi-volume magnum opus *The Art of Computer Programming*, this software allows the production of formatted text for screen-viewing or paper printing, all by writing plain-text code to describe how the formatted text is supposed to appear. TEX is not just a markup language for documents, but it is also a Turing-complete programming language in and of itself, allowing useful algorithms to be created to control the production of documents. Simply put, TEX *is a programmer's approach to word processing*. Since TEX is controlled by code written in a plain-text file, this means anyone may read that plain-text file to see exactly how the document was created. This openness afforded by the code-based nature of TEX makes it relatively easy to learn how other people have created their own TEX documents. By contrast, examining a beautiful document created in a conventional WYSIWYG word processor such as Microsoft `Word` suggests nothing to the reader about *how* that document was created, or what the user might do to create something similar. As Mr. Knuth himself once quipped, conventional word processing applications should be called WYSIAYG (What You See Is *All* You Get).

Leslie Lamport's LATEX extensions to TEX

Like all true programming languages, TEX is inherently extensible. So, years after the release of TEX to the public, Leslie Lamport decided to create a massive extension allowing easier compilation of book-length documents. The result was LATEX, which is the markup language used to create all ModEL module documents. You could say that TEX is to LATEX as `C` is to `C++`. This means it is permissible to use any and all TEX commands within LATEX source code, and it all still works. Some of the features offered by LATEX that would be challenging to implement in TEX include automatic index and table-of-content creation.

Tim Edwards' `Xcircuit` drafting program

This wonderful program is what I use to create all the schematic diagrams and illustrations (but not photographic images or mathematical plots) throughout the ModEL project. It natively outputs PostScript format which is a true vector graphic format (this is why the images do not pixellate when you zoom in for a closer view), and it is so simple to use that I have never had to read the manual! Object libraries are easy to create for `Xcircuit`, being plain-text files using PostScript programming conventions. Over the years I have collected a large set of object libraries useful for drawing electrical and electronic schematics, pictorial diagrams, and other technical illustrations.

`Gimp` graphic image manipulation program

Essentially an open-source clone of Adobe's `PhotoShop`, I use `Gimp` to resize, crop, and convert file formats for all of the photographic images appearing in the ModEL modules. Although `Gimp` does offer its own scripting language (called `Script-Fu`), I have never had occasion to use it. Thus, my utilization of `Gimp` to merely crop, resize, and convert graphic images is akin to using a sword to slice bread.

`SPICE` circuit simulation program

`SPICE` is to circuit analysis as TEX is to document creation: it is a form of markup language designed to describe a certain object to be processed in plain-ASCII text. When the plain-text "source file" is compiled by the software, it outputs the final result. More modern circuit analysis tools certainly exist, but I prefer `SPICE` for the following reasons: it is *free*, it is *fast*, it is *reliable*, and it is a fantastic tool for *teaching* students of electricity and electronics how to write simple code. I happen to use rather old versions of `SPICE`, version 2g6 being my "go to" application when I only require text-based output. `NGSPICE` (version 26), which is based on Berkeley `SPICE` version 3f5, is used when I require graphical output for such things as time-domain waveforms and Bode plots. In all `SPICE` example netlists I strive to use coding conventions compatible with all `SPICE` versions.

Andrew D. Hwang's `ePiX` mathematical visualization programming library

This amazing project is a `C++` library you may link to any `C/C++` code for the purpose of generating PostScript graphic images of mathematical functions. As a completely free and open-source project, it does all the plotting I would otherwise use a Computer Algebra System (CAS) such as `Mathematica` or `Maple` to do. It should be said that `ePiX` is *not* a Computer Algebra System like `Mathematica` or `Maple`, but merely a mathematical *visualization* tool. In other words, it won't determine integrals for you (you'll have to implement that in your own `C/C++` code!), but it can graph the results, and it does so beautifully. What I really admire about `ePiX` is that it is a `C++` programming library, which means it builds on the existing power and toolset available with that programming language. Mr. Hwang could have probably developed his own stand-alone application for mathematical plotting, but by creating a `C++` library to do the same thing he accomplished something much greater.

`gnuplot` mathematical visualization software

Another open-source tool for mathematical visualization is `gnuplot`. Interestingly, this tool is *not* part of Richard Stallman's GNU project, its name being a coincidence. For this reason the authors prefer "gnu" *not* be capitalized at all to avoid confusion. This is a much "lighter-weight" alternative to a spreadsheet for plotting tabular data, and the fact that it easily outputs directly to an X11 console or a file in a number of different graphical formats (including PostScript) is very helpful. I typically set my `gnuplot` output format to default (X11 on my Linux PC) for quick viewing while I'm developing a visualization, then switch to PostScript file export once the visual is ready to include in the document(s) I'm writing. As with my use of `Gimp` to do rudimentary image editing, my use of `gnuplot` only scratches the surface of its capabilities, but the important points are that it's *free* and that it *works well*.

`Python` programming language

Both Python and C++ find extensive use in these modules as instructional aids and exercises, but I'm listing Python here as a *tool* for myself because I use it almost daily as a *calculator*. If you open a Python interpreter console and type `from math import *` you can type mathematical expressions and have it return results just as you would on a hand calculator. Complex-number (i.e. *phasor*) arithmetic is similarly supported if you include the complex-math library (`from cmath import *`). Examples of this are shown in the Programming References chapter (if included) in each module. Of course, being a fully-featured programming language, Python also supports conditionals, loops, and other structures useful for calculation of quantities. Also, running in a console environment where all entries and returned values show as text in a chronologically-ordered list makes it easy to copy-and-paste those calculations to document exactly how they were performed.

Appendix D

Creative Commons License

Creative Commons Attribution 4.0 International Public License

By exercising the Licensed Rights (defined below), You accept and agree to be bound by the terms and conditions of this Creative Commons Attribution 4.0 International Public License ("Public License"). To the extent this Public License may be interpreted as a contract, You are granted the Licensed Rights in consideration of Your acceptance of these terms and conditions, and the Licensor grants You such rights in consideration of benefits the Licensor receives from making the Licensed Material available under these terms and conditions.

Section 1 – Definitions.

a. **Adapted Material** means material subject to Copyright and Similar Rights that is derived from or based upon the Licensed Material and in which the Licensed Material is translated, altered, arranged, transformed, or otherwise modified in a manner requiring permission under the Copyright and Similar Rights held by the Licensor. For purposes of this Public License, where the Licensed Material is a musical work, performance, or sound recording, Adapted Material is always produced where the Licensed Material is synched in timed relation with a moving image.

b. **Adapter's License** means the license You apply to Your Copyright and Similar Rights in Your contributions to Adapted Material in accordance with the terms and conditions of this Public License.

c. **Copyright and Similar Rights** means copyright and/or similar rights closely related to copyright including, without limitation, performance, broadcast, sound recording, and Sui Generis Database Rights, without regard to how the rights are labeled or categorized. For purposes of this Public License, the rights specified in Section 2(b)(1)-(2) are not Copyright and Similar Rights.

d. **Effective Technological Measures** means those measures that, in the absence of proper authority, may not be circumvented under laws fulfilling obligations under Article 11 of the WIPO Copyright Treaty adopted on December 20, 1996, and/or similar international agreements.

e. **Exceptions and Limitations** means fair use, fair dealing, and/or any other exception or

limitation to Copyright and Similar Rights that applies to Your use of the Licensed Material.

f. **Licensed Material** means the artistic or literary work, database, or other material to which the Licensor applied this Public License.

g. **Licensed Rights** means the rights granted to You subject to the terms and conditions of this Public License, which are limited to all Copyright and Similar Rights that apply to Your use of the Licensed Material and that the Licensor has authority to license.

h. **Licensor** means the individual(s) or entity(ies) granting rights under this Public License.

i. **Share** means to provide material to the public by any means or process that requires permission under the Licensed Rights, such as reproduction, public display, public performance, distribution, dissemination, communication, or importation, and to make material available to the public including in ways that members of the public may access the material from a place and at a time individually chosen by them.

j. **Sui Generis Database Rights** means rights other than copyright resulting from Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the legal protection of databases, as amended and/or succeeded, as well as other essentially equivalent rights anywhere in the world.

k. **You** means the individual or entity exercising the Licensed Rights under this Public License. **Your** has a corresponding meaning.

Section 2 – Scope.

a. License grant.

1. Subject to the terms and conditions of this Public License, the Licensor hereby grants You a worldwide, royalty-free, non-sublicensable, non-exclusive, irrevocable license to exercise the Licensed Rights in the Licensed Material to:

A. reproduce and Share the Licensed Material, in whole or in part; and

B. produce, reproduce, and Share Adapted Material.

2. Exceptions and Limitations. For the avoidance of doubt, where Exceptions and Limitations apply to Your use, this Public License does not apply, and You do not need to comply with its terms and conditions.

3. Term. The term of this Public License is specified in Section 6(a).

4. Media and formats; technical modifications allowed. The Licensor authorizes You to exercise the Licensed Rights in all media and formats whether now known or hereafter created, and to make technical modifications necessary to do so. The Licensor waives and/or agrees not to assert any right or authority to forbid You from making technical modifications necessary to exercise the Licensed Rights, including technical modifications necessary to circumvent Effective Technological Measures.

For purposes of this Public License, simply making modifications authorized by this Section 2(a)(4) never produces Adapted Material.

5. Downstream recipients.

A. Offer from the Licensor – Licensed Material. Every recipient of the Licensed Material automatically receives an offer from the Licensor to exercise the Licensed Rights under the terms and conditions of this Public License.

B. No downstream restrictions. You may not offer or impose any additional or different terms or conditions on, or apply any Effective Technological Measures to, the Licensed Material if doing so restricts exercise of the Licensed Rights by any recipient of the Licensed Material.

6. No endorsement. Nothing in this Public License constitutes or may be construed as permission to assert or imply that You are, or that Your use of the Licensed Material is, connected with, or sponsored, endorsed, or granted official status by, the Licensor or others designated to receive attribution as provided in Section 3(a)(1)(A)(i).

b. Other rights.

1. Moral rights, such as the right of integrity, are not licensed under this Public License, nor are publicity, privacy, and/or other similar personality rights; however, to the extent possible, the Licensor waives and/or agrees not to assert any such rights held by the Licensor to the limited extent necessary to allow You to exercise the Licensed Rights, but not otherwise.

2. Patent and trademark rights are not licensed under this Public License.

3. To the extent possible, the Licensor waives any right to collect royalties from You for the exercise of the Licensed Rights, whether directly or through a collecting society under any voluntary or waivable statutory or compulsory licensing scheme. In all other cases the Licensor expressly reserves any right to collect such royalties.

Section 3 – License Conditions.

Your exercise of the Licensed Rights is expressly made subject to the following conditions.

a. Attribution.

1. If You Share the Licensed Material (including in modified form), You must:

A. retain the following if it is supplied by the Licensor with the Licensed Material:

i. identification of the creator(s) of the Licensed Material and any others designated to receive attribution, in any reasonable manner requested by the Licensor (including by pseudonym if designated);

ii. a copyright notice;

iii. a notice that refers to this Public License;

iv. a notice that refers to the disclaimer of warranties;

v. a URI or hyperlink to the Licensed Material to the extent reasonably practicable;

B. indicate if You modified the Licensed Material and retain an indication of any previous modifications; and

C. indicate the Licensed Material is licensed under this Public License, and include the text of, or the URI or hyperlink to, this Public License.

2. You may satisfy the conditions in Section 3(a)(1) in any reasonable manner based on the medium, means, and context in which You Share the Licensed Material. For example, it may be reasonable to satisfy the conditions by providing a URI or hyperlink to a resource that includes the required information.

3. If requested by the Licensor, You must remove any of the information required by Section 3(a)(1)(A) to the extent reasonably practicable.

4. If You Share Adapted Material You produce, the Adapter's License You apply must not prevent recipients of the Adapted Material from complying with this Public License.

Section 4 – Sui Generis Database Rights.

Where the Licensed Rights include Sui Generis Database Rights that apply to Your use of the Licensed Material:

a. for the avoidance of doubt, Section 2(a)(1) grants You the right to extract, reuse, reproduce, and Share all or a substantial portion of the contents of the database;

b. if You include all or a substantial portion of the database contents in a database in which You have Sui Generis Database Rights, then the database in which You have Sui Generis Database Rights (but not its individual contents) is Adapted Material; and

c. You must comply with the conditions in Section 3(a) if You Share all or a substantial portion of the contents of the database.

For the avoidance of doubt, this Section 4 supplements and does not replace Your obligations under this Public License where the Licensed Rights include other Copyright and Similar Rights.

Section 5 – Disclaimer of Warranties and Limitation of Liability.

a. Unless otherwise separately undertaken by the Licensor, to the extent possible, the Licensor offers the Licensed Material as-is and as-available, and makes no representations or warranties of any kind concerning the Licensed Material, whether express, implied, statutory, or other. This includes, without limitation, warranties of title, merchantability, fitness for a particular purpose, non-infringement, absence of latent or other defects, accuracy, or the presence or absence of errors,

whether or not known or discoverable. Where disclaimers of warranties are not allowed in full or in part, this disclaimer may not apply to You.

b. To the extent possible, in no event will the Licensor be liable to You on any legal theory (including, without limitation, negligence) or otherwise for any direct, special, indirect, incidental, consequential, punitive, exemplary, or other losses, costs, expenses, or damages arising out of this Public License or use of the Licensed Material, even if the Licensor has been advised of the possibility of such losses, costs, expenses, or damages. Where a limitation of liability is not allowed in full or in part, this limitation may not apply to You.

c. The disclaimer of warranties and limitation of liability provided above shall be interpreted in a manner that, to the extent possible, most closely approximates an absolute disclaimer and waiver of all liability.

Section 6 – Term and Termination.

a. This Public License applies for the term of the Copyright and Similar Rights licensed here. However, if You fail to comply with this Public License, then Your rights under this Public License terminate automatically.

b. Where Your right to use the Licensed Material has terminated under Section 6(a), it reinstates:

1. automatically as of the date the violation is cured, provided it is cured within 30 days of Your discovery of the violation; or

2. upon express reinstatement by the Licensor.

For the avoidance of doubt, this Section 6(b) does not affect any right the Licensor may have to seek remedies for Your violations of this Public License.

c. For the avoidance of doubt, the Licensor may also offer the Licensed Material under separate terms or conditions or stop distributing the Licensed Material at any time; however, doing so will not terminate this Public License.

d. Sections 1, 5, 6, 7, and 8 survive termination of this Public License.

Section 7 – Other Terms and Conditions.

a. The Licensor shall not be bound by any additional or different terms or conditions communicated by You unless expressly agreed.

b. Any arrangements, understandings, or agreements regarding the Licensed Material not stated herein are separate from and independent of the terms and conditions of this Public License.

Section 8 – Interpretation.

a. For the avoidance of doubt, this Public License does not, and shall not be interpreted to, reduce, limit, restrict, or impose conditions on any use of the Licensed Material that could lawfully

be made without permission under this Public License.

b. To the extent possible, if any provision of this Public License is deemed unenforceable, it shall be automatically reformed to the minimum extent necessary to make it enforceable. If the provision cannot be reformed, it shall be severed from this Public License without affecting the enforceability of the remaining terms and conditions.

c. No term or condition of this Public License will be waived and no failure to comply consented to unless expressly agreed to by the Licensor.

d. Nothing in this Public License constitutes or may be interpreted as a limitation upon, or waiver of, any privileges and immunities that apply to the Licensor or You, including from the legal processes of any jurisdiction or authority.

Creative Commons is not a party to its public licenses. Notwithstanding, Creative Commons may elect to apply one of its public licenses to material it publishes and in those instances will be considered the "Licensor." Except for the limited purpose of indicating that material is shared under a Creative Commons public license or as otherwise permitted by the Creative Commons policies published at `creativecommons.org/policies`, Creative Commons does not authorize the use of the trademark "Creative Commons" or any other trademark or logo of Creative Commons without its prior written consent including, without limitation, in connection with any unauthorized modifications to any of its public licenses or any other arrangements, understandings, or agreements concerning use of licensed material. For the avoidance of doubt, this paragraph does not form part of the public licenses.

Creative Commons may be contacted at `creativecommons.org`.

Appendix E

References

Basso, Christophe P., *Switch-Mode Power Supplies – SPICE Simulations and Practical Designs*, McGraw-Hill Companies Incorporated, New York, NY, 2008.

Ćuk, Slobodan M., and Middlebrook, Robert D. *US Patent 4,184,197*, "DC-to-DC Switching Converter", application 28 September 1977, patent granted 15 January 1980.

Hughes, Thomas P., *Networks of Power – Electrification in Western Society, 1880-1930*, The Johns Hopkins University Press, Baltimore, MD, 1983.

"LNK362-364 energy-efficient, low power off-line switcher IC", datasheet revision I, Power Integrations, San Jose, CA, December 2021.

Maniktala, Sanjaya, *Switching Power Supplies A-Z*, Elsevier, London, 2012.

Niknejad, A.M. and Dunsmore, Joel; "EECS 142 Laboratory #1 – High Frequency Passive Components", University of California, Berkeley, CA, 11 February 2013.

"Power Supply Topologies", document SLUW001A, Texas Instruments Incorporated, Dallas, TX, 2004.

"Power Supply Topologies Quick Reference Guide", document SLYU032, Texas Instruments Incorporated, Dallas, TX, 2016.

"The Chelsea Electricity Supply Company's Station and Sub-Stations", *Western Electrician*, Volume VI, Numbers 1-26, pp. 342-343, Electrician Publishing Company, Chicago, IL, 21 June 1890.

Garcia, David; Valencia, Jose; Hausman, Greg; "US Department of Energy's 2016 transformer efficiency changes and impact on inrush current for low voltage dry-type transformers", whitepaper WP009002EN / Z18069, Eaton, Cleveland, OH, 2016.

Zehendner, Markus and Ulmann, Matthias, *Power Topologies Handbook*, Texas Instruments Incorporated, Dallas, TX, 2017.

Appendix F

Version history

This is a list showing all significant additions, corrections, and other edits made to this learning module. Each entry is referenced by calendar date in reverse chronological order (newest version first), which appears on the front cover of every learning module for easy reference. Any contributors to this open-source document are listed here as well.

11 October 2026 – added an analogy to the "Continuous and discontinuous converter operation" section of the Tutorial chapter, explaining inductor and capacitor energy storage in terms of a cup being alternately filled and emptied.

17 September 2026 – added a new section to the Tutorial chapter on continuous versus discontinuous converter operation.

12 May 2026 – added some content to the "Integrated circuit power converters" section of the Tutorial, specifically on voltage regulation via feedback. Also corrected some spelling errors.

20 April 2026 – corrected a couple of errors courtesy of Wyatt Gleason. One of these was a simple typographical error, and the other was conceptual: stating that load voltage in the inverting converter circuit would be the series sum of inductor and source voltage when in fact it was all due to the inductor's voltage.

6 October 2025 – added a new Tutorial section showing micro-photographs of power converter ICs from zeptobars.com

6-7 November 2024 – added "on" and "off" MOSFET state labels to the timing diagrams in image_4101. Also added a sub-question regarding Ampere-hour calculations to the "Dying battery source" Quantitative Reasoning question. Also corrected a typo courtesy of Ethan Kelley. Also added some more descriptive text and changed the Challenge question for the "Instantaneous analyses of DC-DC converter circuits" Quantitative Reasoning question.

4 November 2024 – added two new animations to the module: one showing a BJT turning on and the other showing an E-type MOSFET turning on.

25 October 2024 – divided the Introduction chapter into sections, one with recommendations for students, one with a listing of challenging concepts, and one with recommendations for instructors.

23 September 2024 – grammatical edits to a couple of the questions in the Introduction chapter, courtesy of Nick Brown.

10 April 2024 – re-formatted some of instructor-only images for better viewability, as well as edited image_6780 to correct a typo.

14 November 2023 – added a new Case Tutorial section on empirically determining signal rates of change.

9 November 2023 – added a instructor note on "Identifying possible root-cause faults" Diagnostic Reasoning question, and also added an Introduction chapter question.

3 July 2023 – added a Quantitative Reasoning question requiring careful KVL/KCL analysis of a buck converter circuit at a particular moment in time.

26 April 2023 – added a Conceptual Reasoning question focusing on energy efficiency for a lightly-loaded DC-DC converter.

11-12 April 2023 – added content to the Tutorial chapter regarding disadvantages of simple PWM-based step-down conversion, combining linear and switch-mode power converter circuits, as well as added some instructor notes to questions. Also corrected some typographical errors and added questions to the Introduction chapter.

28 November 2022 – placed questions at the top of the itemized list in the Introduction chapter prompting students to devise experiments related to the tutorial content.

28 December 2021 – added Technical Reference sections on capacitor parasitics and inductor parasitics.

27 November 2021 – added a Diagnostic Reasoning question on root-causes for a failed transistor in a converter circuit.

11 November 2021 – minor correction (forgot a right-parentheses symbol) in the Tutorial, and also elaborated on Ćuk converter circuits using low-side switching elements.

22 October 2021 – added technical reference sections on capacitor and inductor parameters, types. Also added comments about inductor parameters needed for a DC-to-DC converter circuit.

18 September 2021 – added an example contrasting energy efficiency for a linear versus switch-mode voltage regulator to the first section of the Tutorial chapter.

13 May 2021 – added review_10.latex to the Tutorial to give students a refresher on transistor operation.

9 May 2021 – commented out or deleted empty chapters.

3 May 2021 – clarified in the Tutorial that the duty cycle formulae for various converter circuits is ideal and only applies to very light- or no-load conditions.

18 April 2021 – changed all lower-case Greek letter "phi" symbols (ϕ) to upper-case (Φ).

9 April 2021 – minor edit to the Tutorial section on forward converter circuits, discussing how the reset current is able to slightly re-charge a chemical battery functioning as the source.

6-7 April 2021 – added content to the Introduction chapter and to the Conceptual Reasoning section, as well as minor edits to other questions and to the Tutorial.

2 April 2021 – added a section to the Tutorial on charge pump circuits, and corrected a minor typographical error in the Ćuk converter patent Historical Reference section.

18 March 2021 – corrected one instance of "volts" that should have been capitalized "Volts".

1 January 2021 – added discussion of energy efficiency to the Tutorial, as well as added comments to the "Battery substations" Historical Reference section.

12 November 2020 – added another Case Tutorial section, this one showing a boost converter.

10-11 November 2020 – corrected and added some instructor notes, as well as made minor additions to the Introduction chapter. Also, corrected some details in the Inverting and Ćuk converter circuit diagrams.

8 November 2020 – added labels to several components in several schematic diagrams that were unlabeled.

25 October 2020 – added Case Tutorial chapter with a section showing an elementary buck converter circuit suitable for breadboard demonstration.

14 October 2020 – significantly edited the Introduction chapter to make it more suitable as a pre-study guide and to provide cues useful to instructors leading "inverted" teaching sessions. Also, added some index entries.

27 July 2020 – added another question to the Conceptual Reasoning section.

17 July 2020 – added another question to the Conceptual Reasoning section.

9-13 July 2020 – added content to Tutorial and questions.

7 July 2020 – added Historical Reference to the Chelsea Electricity Supply Company's battery-based substations.

14 November 2019 – document first created.

Index

Ćuk converter, 43

Adding quantities to a qualitative problem, 180
Air-core inductor, 85
Alternating current (AC), 21
Amplification, 26
Annotating diagrams, 33, 36, 40, 42, 45, 48, 50, 58, 60, 61, 179

B-H curve, 85
Base, 27
Bipolar, 27
Bipolar junction transistor, 26
BJT, 26
Boost charge pump, 60
Boost converter, 39
Brute-force power supply, 63
Buck charge pump, 61
Buck converter, 35
Buck-boost converter, 41

Capacitance, 23
Capacitor, electrolytic, 63
CCM, 52
Charge carrier injection, BJT, 27
Charge pump, 57
Charge pump, boost, 60
Charge pump, buck, 61
Charge pump, inverting, 58
Checking for exceptions, 180
Checking your work, 180
Class 1 ceramic dielectrics, 82
Class 2 ceramic dielectrics, 82
Code, computer, 189
Coefficient, temperature, 82
Coercivity, 86
Collector, 27
Conservation of Energy, 21–23, 33
Continuous mode, 51
Converter, Ćuk, 43
Converter, boost, 39
Converter, buck, 35
Converter, buck-boost, 41
Converter, flyback, 47
Converter, forward, 49
Converter, inverting, 41
Core reset, 50
Coupling, 57
Curie temperature, 84
Current probe, 15

DC-DC converter, 21
DCM, 52
Department of Energy, 21
Dielectric, 23
Dimensional analysis, 179
Diode, Zener, 20
Direct current (DC), 20
Discontinuous mode, 51
Displacement current, 15
Distance versus speed, 17
Drain, 28
Duty cycle, 31

Eddy current, 92
Edwards, Tim, 190
Efficiency, 20, 21, 62
Electric field, 23, 57
Electrical isolation, 47
Electrical noise, 32, 62
Electrolytic capacitor, 63
Electromagnetic induction, 85, 86
Emitter, 27
Equivalent Series Resistance, 88, 91

Equivalent Series Resistance, capacitor, 63
ESR, 57, 63, 88, 91

Faraday's Law of Electromagnetic Induction, 10, 50
Feedback, negative, 34, 66
Ferrite, 86
FET, 26
Field, electric, 23
Field, magnetic, 23
Field-effect transistor, 26
Filter, low-pass, 38, 41, 44, 62
Flyback converter, 47
Flying capacitor, 57
Forward converter, 49
Fourier's Theorem, 32, 62

Gate, 28
Graph values to solve a problem, 180
Greenleaf, Cynthia, 151
Ground loop, 47

Hard magnetic material, 85
High-side switching, 41, 42
How to teach with these modules, 187
Hwang, Andrew D., 191
Hysteresis curve, 85

Identify given data, 179
Identify relevant principles, 179
Inductance, 23
Induction, 85, 86
Injection, charge carriers in BJT, 27
Insulated-gate field effect transistor, 28
Intermediate results, 179
Inversion layer, 142
Inverted instruction, 187
Inverting charge pump, 58
Inverting converter, 41
Isolation, electrical, 47

JFET, 28
Joule's Law, 21, 68, 85
Junction field-effect transistor, 28

Knuth, Donald, 190

Laminated core, 86
Lamport, Leslie, 190
Limiting cases, 3, 36, 40, 42, 180
Linear voltage regulator, 20
Load, 24
Low-pass filter, 38, 41, 44, 62
Low-side switching, 41

Magnetic field, 23, 57
Magnetic saturation, 50, 85, 86
Magnetically "hard" versus "soft" materials, 85
Magnetostriction, 85
Maxwell, James Clerk, 67
Metacognition, 156
Metal-oxide-semiconductor field-effect transistor, 28
Mode, continuous, 51
Mode, discontinuous, 51
Moolenaar, Bram, 189
MOSFET, 28
Murphy, Lynn, 151
Mutual inductance, 47
Mylar dielectric, 81

Negative feedback, 34, 66
Noise, electrical, 32, 62

Open-source, 189

Parasitic effect, 15, 57, 88, 91
Piezoelectricity, 82
Pinch-off voltage, 26
Pitch, 15
Power supply, switch-mode, 21
Problem-solving: annotate diagrams, 33, 36, 40, 42, 45, 48, 50, 58, 60, 61, 179
Problem-solving: check for exceptions, 180
Problem-solving: checking work, 180
Problem-solving: dimensional analysis, 179
Problem-solving: graph values, 180
Problem-solving: identify given data, 179
Problem-solving: identify relevant principles, 179
Problem-solving: interpret intermediate results, 179
Problem-solving: limiting cases, 3, 36, 40, 42, 180
Problem-solving: qualitative to quantitative, 180

Problem-solving: quantitative to qualitative, 180
Problem-solving: reductio ad absurdum, 180
Problem-solving: simplify the system, 179
Problem-solving: thought experiment, 179
Problem-solving: track units of measurement, 179
Problem-solving: visually represent the system, 179
Problem-solving: work in reverse, 180
Pulse width modulation (PWM), 31

Qualitatively approaching a quantitative problem, 180

Reading Apprenticeship, 151
Reductio ad absurdum, 180, 186, 187
Regulator, voltage, 34
Remanence, 86
Reset, 50
Reverse power flow, 59
Ripple, 33
Ripple rejection, 64
Ripple voltage, 63
Rise over run, 15
Roll-off, filter network, 62

Saturation, magnetic, 50, 85, 86
Schoenbach, Ruth, 151
Scientific method, 156
Self-inductance, 47
Simplifying a system, 179
Sinusoidal decomposition, 36
Skin effect, 85
Slope, 15
SMPS, 21
Soakage, capacitor, 88, 90
Socrates, 186
Socratic dialogue, 187
Soft magnetic material, 85
Source, 24, 28
Speed versus distance, 17
SPICE, 151
Stallman, Richard, 189
Switching converter, 21

Tangent line, 15
Tank circuit, 91
Temperature coefficient of capacitance, 82
Thought experiment, 179
Torvalds, Linus, 189
Transistor, 20
Turn-on voltage, 26

Units of measurement, 179

Visualizing a system, 179
Voltage regulator, 34
Voltage, pinch-off, 26
Voltage, turn-on, 26

Work in reverse to solve a problem, 180
WYSIWYG, 189, 190

Zener diode, 20