Modular Electronics Learning (ModEL) project

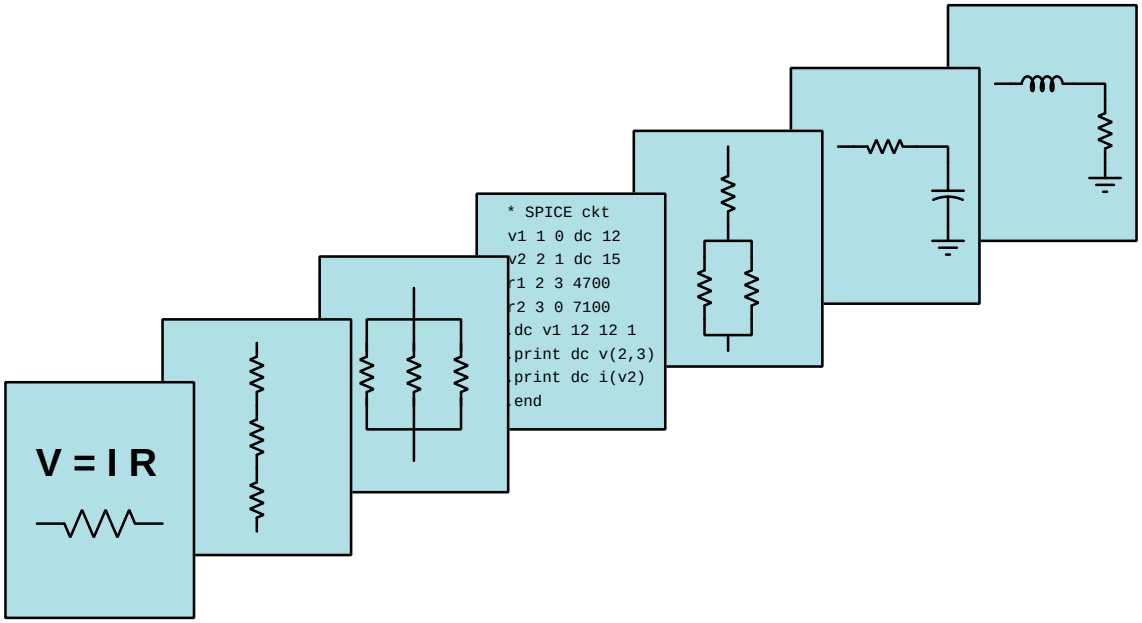

Thyristors

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

Last update = 24 July 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 thyristors . . . 5
- 1.3 Recommendations for instructors . . . 6

2 Case Tutorial **7**
- 2.1 Example: measuring signal rates of change . . . 8

3 Tutorial **13**
- 3.1 Hysteresis . . . 14
- 3.2 Shockley diodes and DIAC thyristors . . . 15
- 3.3 Critical rate of voltage rise . . . 18
- 3.4 SCR and GTO thyristors . . . 20
- 3.5 TRIAC thyristors . . . 25
- 3.6 Thyristor gate sensitivity . . . 28

4 Historical References **29**
- 4.1 Thyratron inverter patent . . . 30

5 Derivations and Technical References **33**
- 5.1 Lightning bolts, relaxation oscillators, and thyratron tubes . . . 34
- 5.2 Latch-up in integrated circuits . . . 38

6 Questions **49**
- 6.1 Conceptual reasoning . . . 53
  - 6.1.1 Reading outline and reflections . . . 54
  - 6.1.2 Foundational concepts . . . 55
  - 6.1.3 SCR versus TRIAC circuits . . . 56
  - 6.1.4 Turning an SCR on and off . . . 57
  - 6.1.5 SCR voltage drop . . . 57
  - 6.1.6 Sensitive-gate thyristors . . . 57
  - 6.1.7 SCR-based amplifier . . . 58
  - 6.1.8 Different thyristor types . . . 58
  - 6.1.9 UJT test circuit . . . 59

6.2 Quantitative reasoning . . . 60
6.2.1 Miscellaneous physical constants . . . 61
6.2.2 Introduction to spreadsheets . . . 62
6.2.3 Voltage rate-of-change from oscillographs . . . 65
6.2.4 2N6403 specifications . . . 66
6.2.5 2N6344 specifications . . . 67
6.2.6 UJT trigger voltage . . . 68
6.2.7 UJT standoff ratio . . . 69
6.3 Diagnostic reasoning . . . 70
6.3.1 Identifying SCR terminals . . . 71
6.3.2 Faults in a UJT circuit . . . 72

A Problem-Solving Strategies **73**

B Instructional philosophy **75**
B.1 First principles of learning . . . 76
B.2 Proven strategies for instructors . . . 77
B.3 Proven strategies for students . . . 79
B.4 Design of these learning modules . . . 80

C Tools used **83**

D Creative Commons License **87**

E References **95**

F Version history **97**

Index **99**

Chapter 1

Introduction

1.1 Recommendations for students

An important classification of semiconductor device is the *thyristor*, characterized by its hysteretic behavior: once turned on, it tends to remain on; once turned off, it tends to remain off. These devices are popularly used for electronic power control, because their tendency to "latch" in the conductive state results in quick switching and low voltage drop, both of which minimize power dissipation in the device. Some thyristors are passive while others come equipped with "gate" terminals for triggering by external signals. Gated thyristors are the most versatile, the two predominant types being *silicon-controlled rectifiers* (*SCR*s) and *TRIAC*s, the former being DC-only and the latter able to switch AC. Thyristors find widespread application as high-power control devices, rivaling transistors in their ability to quickly turn on and to conduct large currents with minimal power dissipation. This module explores the basic behavior of several types of thyristors as well as power-control applications.

Important concepts related to thyristors include positive **feedback**, **hysteresis**, **PN junctions**, **bipolar junction transistor behavior**, **breakdown voltage**, **breakover**, **low-current dropout**, **holding current**, **crowbar** circuits, properties of **series** versus **parallel** networks, **capacitor** behavior, **snubber** circuits, **active** versus **passive** components, **rectification**, **commutation**, **extracting gate current**, and operational **quadrants**.

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

- How might an experiment be designed and conducted to measure the holding current for a thyristor? 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 maximum rate of voltage rise for a thyristor? What hypothesis (i.e. prediction) might you pose for that experiment, and what result(s) would either support or disprove that hypothesis?
- What are some practical applications of thyristors?
- What is "hysteresis" and what are some practical applications of this concept?

- What conditions are necessary to turn on a bipolar junction transistor?
- What conditions are necessary or sufficient to turn on a thyristor?
- What can a transistor do that a thyristor cannot?
- What can a thyristor do that a transistor cannot?
- What does it mean to "fire" a thyristor?
- What does it mean to "latch" a thyristor?
- What does it mean to "commutate" a thyristor?
- How may a thyristor be *accidently* fired?
- How may a thyristor be tested?
- What is *breakdown voltage*, and why is this important to the operation of a thyristor?
- What is a *thought experiment*, and how was this problem-solving technique applied in the tutorial? Hint: it was used more than once!
- What conditions are necessary or sufficient to turn off a thyristor?
- How does a "crowbar" help protect voltage-sensitive components?
- How does the rate-of-change of a variable differ from the instantaneous value of that same variable; e.g. Volts versus Volts-per-second, or kilometers versus kilometers-per-hour?
- Show how to calculate the voltage rate-of-change described in the Tutorial from the information found on the oscillograph.
- Why are thyristors difficult to turn off?
- How does the presence of a capacitor help commutation in a dual-thyristor power control circuit?
- What are some practical applications of thyristors?

1.2 Challenging concepts related to thyristors

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

- **Hysteresis** – whenever a component or system tends to remain in its last state, we refer to that lagging tendency as *hysteresis*. There are many physical manifestations of hysteresis including *magnetic hysteresis* (where a magnetic material tends to remain magnetized after the magnetizing influence is removed), *thyristor devices* (electronic devices that tend to remain "on" even after the triggering stimulus is removed), and physical mechanisms having *friction* (where the mechanism's position tends to stay put after the influencing force is removed).
- **Rates of change** – thyristors are liable to triggering when subjected to excessive voltage rates-of-change, which requires one to be aware of how rapidly voltage varies and not just the value of that voltage. This is the first hurdle in calculus: to comprehend what a rate of change is, and it is far from obvious. A helpful strategy here is to relate electrical rates of change to other, more common, rates of change such as *speed* (i.e. the rate-of-change of physical position over time). Also helpful is to relate rates of change to the *slope* of a line or curve plotted on a graph, showing how *rise/fall over run* describes not only the phenomenon of a rate-of-change but also the appropriate units of measurement.

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** – Distinguish between SCR and TRIAC behavior

  Assessment – Compare and contrast the operation of two circuits, one using an SCR and the other a TRIAC to control power to a load; e.g. pose problems in the form of the "SCR versus TRIAC circuits" Conceptual Reasoning question.

- **Outcome** – Quantify voltage rates of change from an oscillograph

  Assessment – Identify $\frac{dV}{dt}$ at various marked points of an oscillograph given the image as well as the oscilloscope's vertical sensitivity and horizontal timebase settings; e.g. pose problems in the form of the "???" Quantitative Reasoning question.

- **Outcome** – Independent research

  Assessment – Locate SCR datasheets and properly interpret some of the information contained in those documents including holding current, critical rate of voltage rise, maximum current rating, maximum voltage rating, etc.

  Assessment – Locate TRIAC datasheets and properly interpret some of the information contained in those documents including holding current, critical rate of voltage rise, maximum current rating, maximum voltage rating, operating quadrant chart, 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: 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"[1] 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[2]. 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*.

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

[2] 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 ($\frac{dV}{dt}$):

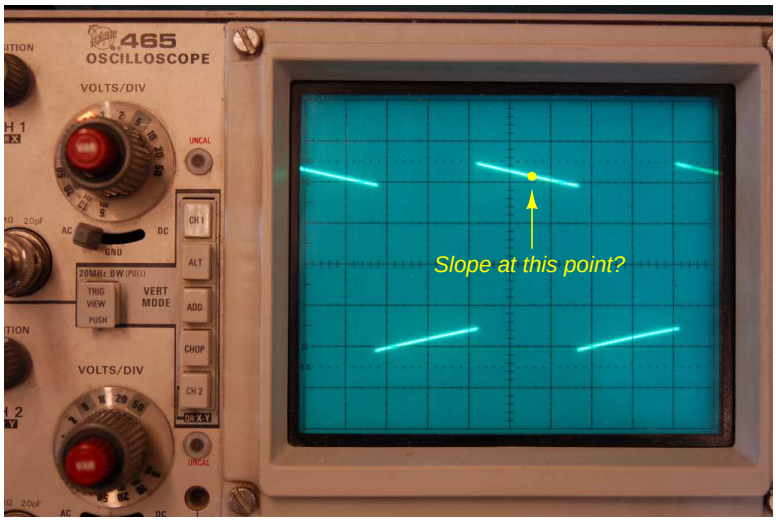

Next we see a tangent line drawn to match the slope of the waveform at the specified location, with convenient points for fall/run measurements taken on that line against the oscilloscope grid's major divisions. In this example, the oscilloscope's vertical sensitivity has been set for 0.5 Volts per division,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

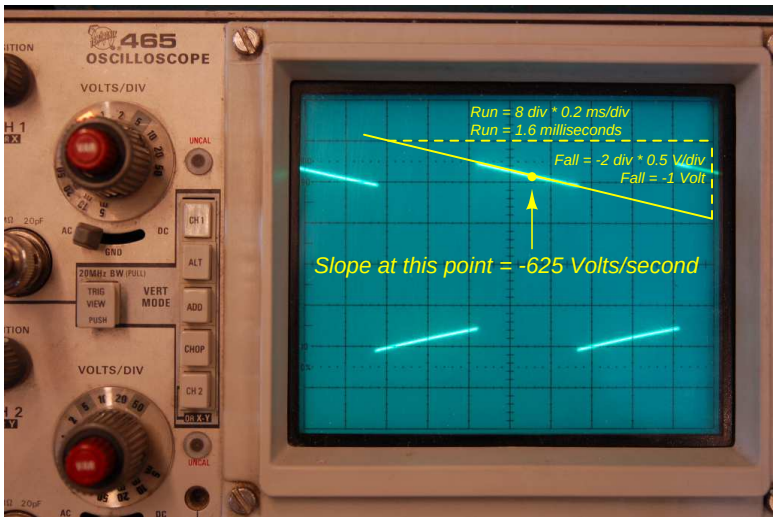

As we can see, the tangent line falls 2 vertical divisions (-1 Volt) over a timespan of 8 horizontal divisions (1.6 milliseconds), yielding a $\frac{dV}{dt}$ quotient of -625 Volts per second, which may also be expressed as -0.625 Volts per millisecond. The negative sign is important, as it distinguishes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

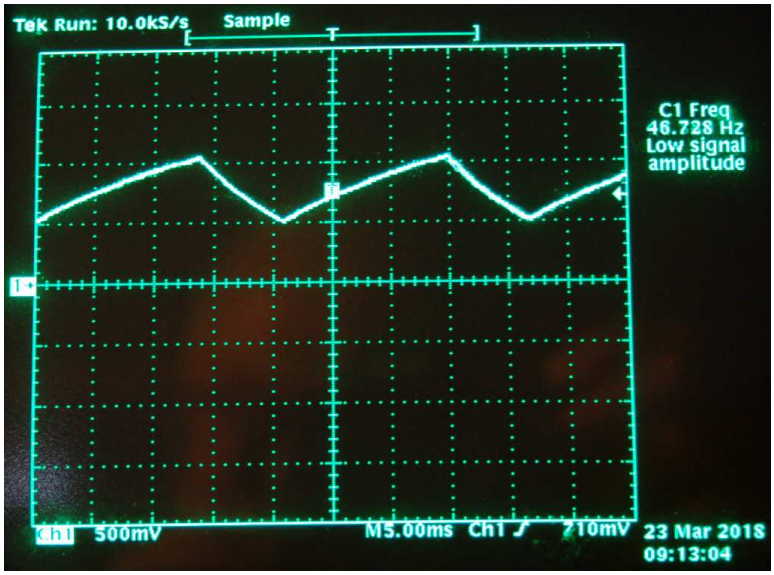

Approximating the slope for each rising section of this wave, we count one division of rise over 2.8 divisions of run, or 500 milliVolts rise over 14 milliseconds of run. This is a rate-of-rise of +35.7 Volts per second.

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:

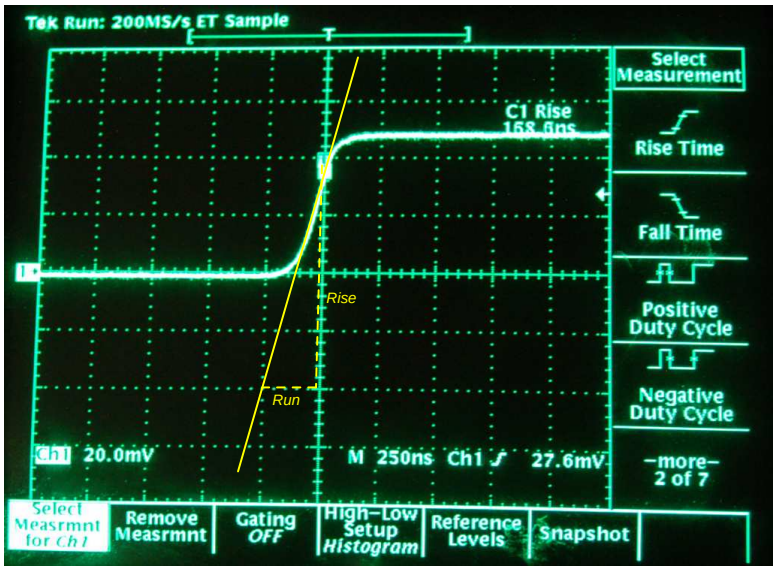

The tangent line overlaid on this screenshot for the purpose of measuring the pulse edge's rate-of-change rises approximately 4 vertical divisions over a run of 1 division, which is 80 milliVolts of rise over 250 nanoseconds of run. The pulse edge's rate-of-change, therefore, is approximately +320,000 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 Hysteresis

Hysteresis is that property of any component or system which fails to return to its original state after removal of some stimulus. A very simple example of hysteresis is the mechanical action of a toggle switch: when the lever is pushed, it flips to one position and will remain there even after the lever is no longer being pushed. The only way for the toggle switch to return to its original position is if you apply force in the opposite direction. Contrast this against the action of a spring-loaded pushbutton switch, which actuates when pressed and immediately de-actuates when you stop pressing it.

Most electronic devices are non-hysteretic: that is to say, they return to their default states upon removal of their control signals. Transistors, both bipolar and field-effect, are non-hysteretic because they always return to their default states when their controlling signals (e.g. base current for a BJT, gate-source voltage for a FET) cease. However, a class of semiconductor devices exist intentionally designed to be hysteretic, called *thyristors*. Once a thyristor has been "triggered" into its *on* state, it will continue to conduct current even after the triggering signal ceases. Like a toggle switch, a thyristor remains on after being turned on, and remains off after being turned off. As such, thyristors are only useful as on/off switching devices, unable to "throttle" between these two extreme states as transistors can.

Thyristors are constructed of PN semiconductor junctions just like BJTs, and achieve their hysteretic action by means of *positive feedback*. This is where a portion of their controlled current is directed in such a way that it maintains forward-bias in one of the PN junctions to hold the other junctions in their conductive states. In other words, once triggered "on" by an external signal, the thyristor's own controlled current functions as a controlling signal to keep it on. Once forced "off" by some external circumstance, the thyristor's lack of controlled current will also mean it lacks that self-driving stimulus resulting from a portion of the controlled current to hold itself in the "on" state, and so it remains off.

This positive-feedback action not only produces the hysteresis effect, but it also results in the thyristor having a very low voltage drop while conducting. The more current it conducts, the more it turns itself on, resulting in more controlled current, and so on until the device utterly saturates in the "on" state. Thyristors, therefore, are ideal for high-power switching applications where significant voltage drop would lead to excessive heat dissipation. Unlike a power *transistor* which could be damaged by its own heat dissipation if not "driven" to saturation by an appropriately large controlling signal, a power *thyristor* only needs a small driving signal to turn on, after which it very quickly forces itself into saturation and dissipates minimal heat. In fact, this positive feedback effect is so strong that it is often more difficult to make a thyristor turn back *off* than it is to trigger it on!

3.2 Shockley diodes and DIAC thyristors

Perhaps the most basic form of thyristor is the *four-layer diode*, also known as a *PNPN diode*, or a *Shockley diode* after its inventor, William Shockley. This is not to be confused with a *Schottky diode*, that two-layer metal-semiconductor device known for its high switching speed. A crude illustration of the Shockley diode, often seen in textbooks, is a four-layer sandwich of P-N-P-N semiconductor material:

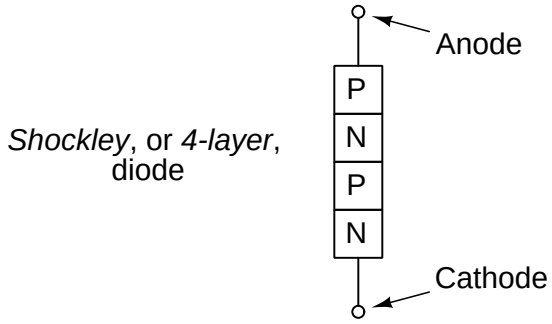

This rendering of a Shockley diode, though, does little to explain how it functions. It would appear to be nothing more than a bipolar junction transistor with an additional layer and corresponding junction. A more useful illustration represents the four-layer semiconductor structure as a pair of interconnected three-layer structures, essentially a PNP and NPN transistor pair:

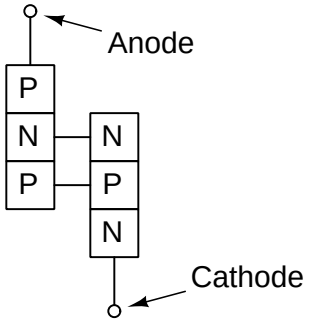

A schematic representation of the Shockley diode clarifies which terminals are emitter and collector between the two transistors:

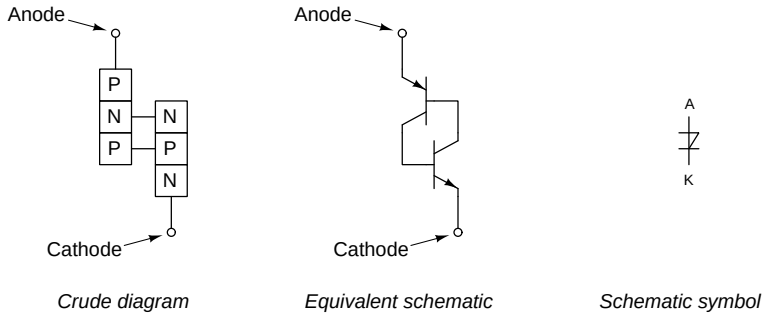

Crude diagram *Equivalent schematic* *Schematic symbol*

A test circuit will help us perform a *thought experiment* on the Shockley diode:

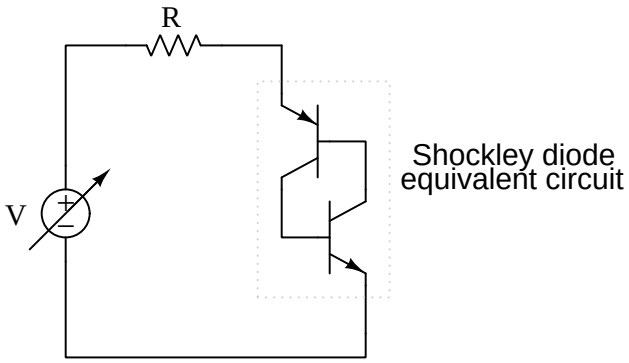

If we imagine the variable-voltage source V being adjusted from zero upwards, it would appear the Shockley diode never conducts current. Its constituent BJTs are, of course, *normally-off* devices, requiring base current before any collector current may flow. However, since the base of each BJT connects to the collector of the other, it would seem both transistors are essentially "waiting" for the other to turn on, and since neither of them are able, they both remain off.

However, this is what a pair of *ideal* BJTs would do. In practice, reverse-biased PN junctions – such as the base-collector junctions of the two transistors – have *breakdown voltage* limits, beyond which some current will pass "backwards" through the junction. If the applied DC voltage exceeds this limit for either of the two equivalent transistors within the Shockley diode, as soon as one of the reverse-biased base-collector junctions breaks down it will pass a small amount of current to the base of the other transistor, properly turning that transistor on. Once turned on, that other transistor allows base current through the first, properly turning it on as well, at which point the two BJTs mutually driving one another into saturation and hold each other in their "on" states even if the applied voltage were decreased below the initial breakdown point. The Shockley diode will remain conducting current until the source voltage decreases nearly to zero, at which point current ceases and the two BJTs return to their "off" states.

So we see that a Shockley diode's behavior is hysteretic: *once off it tends to remain off, and once on it tends to remain on.* It is also unidirectional, with the anode needing to be positive and the cathode negative in order to conduct current. If we parallel another Shockley diode (facing the reverse direction), we will have a device known as a *DIAC* which is bidirectional:

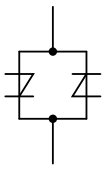

DIAC equivalent circuit

DIAC schematic symbol

Some special terms exist for all PNPN-type thyristors. First is the term used to describe its "on" state: *latched*. The word "latch" is reminiscent of a door lock mechanism, which tends to keep the door closed once it has been pushed shut. The term *firing* refers to the initiation of a latched state. In order to get a Shockley diode to latch, the applied voltage must be increased until *breakover* is attained. Despite the fact that this action is best described in terms of transistor break*down*, the term break*over* is used instead because the end result is a pair of transistors in mutual saturation rather than destruction as would be the case with a normal transistor. A latched Shockley diode is re-set back into its nonconducting state by reducing current through it until *low-current dropout* occurs. The minimum amount of anode-cathode current necessary to guarantee the thyristor remains in its conductive state is called the *holding current*.

Shockley diodes and DIACs alike are useful as *triggering* and as *crowbar* devices. As a triggering device, it passes no current until the impressed voltage exceeds the breakover value at which point it turns on fully, this sudden onset of current useful as a controlling impulse signal for some other circuit component. As a crowbar device, it is connected in parallel with the power terminals of voltage-sensitive components, drawing no current unless and until a condition of excessive voltage occurs, at which point it fires and "shorts out" the offending over-voltage[1] and in so doing protects all the other parallel-connected components.

[1] The colorful term "crowbar" evokes the image of a metal pry-bar intentionally dropped across two bare conductors for the purpose of causing an intentional short-circuit. While no doubt a desperate measure, presumably the consequences of allowing an over-voltage condition would be worse for the other paralleled components than the consequences for the crowbar or any other components in the fault-current path. If the notion of an electronic "crowbar" strikes you as humorous, you are in good company: I actually laughed out loud when I first learned there was such a thing.

3.3 Critical rate of voltage rise

PNPN-type thyristors may fire by means other than breakover voltage: excessive *voltage rate of rise*[2] ($\frac{dV}{dt}$) may cause thyristors to fire prematurely. This is when the applied voltage across the thyristor increases at a high rate of change, causing latching due to inherent junction capacitances within all non-conducting PN junctions. Recall that current "through" a capacitance is proportional to the applied voltage's rate-of-change ($I = C\frac{dV}{dt}$). If the applied voltage between anode and cathode rises too quickly, the intrinsic capacitances of the reverse-biased PN junctions will permit enough current during that time to activate the transistor pair, firing the thyristor.

Measuring a voltage signal's *rate-of-change* is best done using an oscilloscope which graphs voltage on the vertical axis and time on the horizontal. The example below shows a signal captured with the sensitivity set for 100 milliVolts per division (vertical square) and the timebase set for 20 nanoseconds per division (horizontal square):

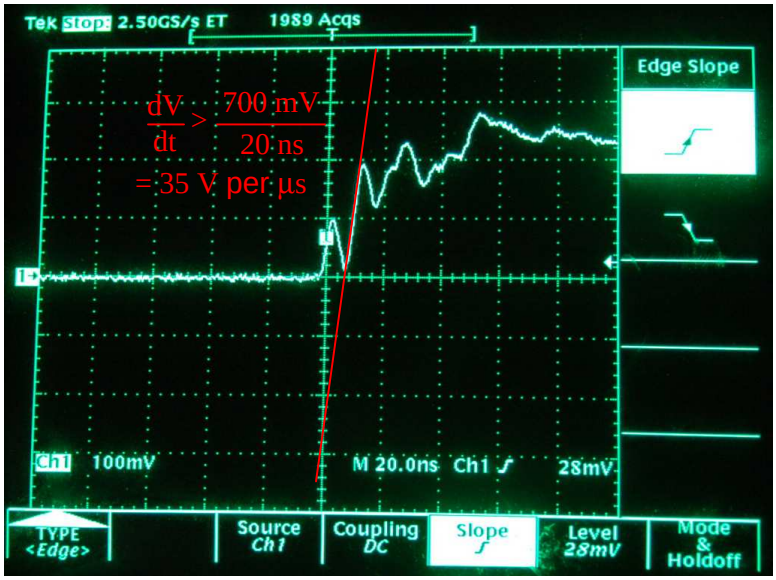

This signal's highest peak voltage is slightly less than 300 milliVolts (i.e. just under three divisions of height), yet its *rate of change* during the rising interval marked by the red diagonal line[3] superimposed on the oscillograph is steeper than *35 million Volts per second* (i.e. more than 700 milliVolts of rise over less than 20 nanoseconds of time)! 35 million Volts per second, more commonly expressed as 35 V/μs, is an extremely high rate of change which could in principle fire a thyristor even though the peak voltage of < 300 mV is too low to cause breakover.

[2] The rate-of-change of any variable is a measure of how *quickly* its value changes over time. You may think of rate-of-change as being akin to the *speed* of a variable, versus its actual value at any given time. In fact, distance (x) and velocity ($\frac{dx}{dt}$) represent this same distinction: the former representing the length of some physical span, and the latter representing how rapidly that length varies over time. In the case of thyristors, the rate of the applied voltage's change over time is relevant because a high rate of "speed" for this voltage is capable of triggering the thyristor even in cases where the actual voltage value is significantly less than the breakover voltage.

[3] This red line is merely a visual aid for determining the signal's rate-of-change over that rising interval, allowing us to more easily discern its slope, its rise exceeding seven vertical divisions over a horizontal run of one division.

A performance parameter specified in thyristor datasheets is the device's *critical rate of voltage rise*, which is the steepest rate-of-change of applied voltage it may tolerate without firing. Since it is usually undesirable for thyristors to fire by steep rates-of-change, it is important to have mitigation techniques for this possibility. One common strategy is to block fast voltage rises from reaching the thyristor using series inductors and/or parallel resistor-capacitor networks called *snubbers*:

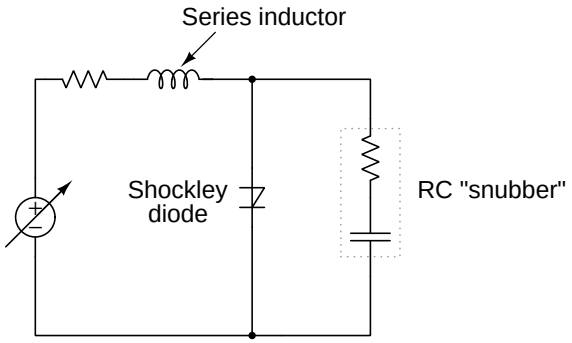

Both the series inductor and the parallel resistor-capacitor "snubber" circuit help minimize the Shockley diode's exposure to excessively rising voltages.

The inductor's task is to limit high rates-of-change of current by dropping voltage (as a load) if ever the source's current rises or falls rapidly ($V = L\frac{dI}{dt}$). The capacitor's task is to directly limit high rates-of-change of voltage across the thyristor by passing current (as a load) if the thyristor's terminal voltage rises rapidly ($I = C\frac{dV}{dt}$). Acting together, these two devices help shield the thyristor from experiencing too steep of a voltage rate-of-change.

Another mechanism of false firing for thyristors is excessive temperature, causing greater "leakage" current when the device is supposed to be in its "off" state.

3.4 SCR and GTO thyristors

As interesting and useful as Shockley diodes and DIACs are, though, they would be even more interesting and more useful if we could *control* them with weak electrical signals as we can transistors. We call such controllable components *active*, in contrast with *passive* components such as resistors, inductors, capacitors, and diodes. It was the desire for "active" behavior that led to the development of vacuum tubes and transistors, and ushered in the era of truly *electronic* circuits. So it is here, where we would like to be able to exert control over a thyristor's state using a relatively small electrical signal, rather than only have it passively respond to impressed voltage.

A simple modification of the Shockley diode yields this valuable active characteristic: attaching a third terminal to the PNPN structure which we will label the *gate*. The following illustration shows the crude diagram, equivalent schematic, and symbol for the resulting device, called a *silicon-controlled rectifier* or *SCR*:

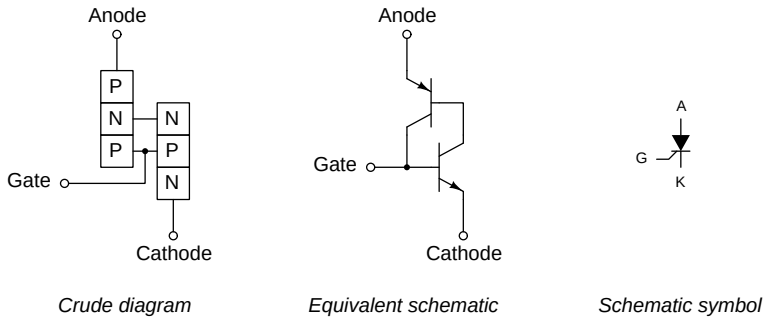

Applying a small voltage signal between the gate (+) and cathode (−) terminals "triggers" the SCR into its conductive state without requiring a large anode-cathode applied voltage, enabling us to drive it "on" at will. Once triggered "on" the SCR will remain on so long as sufficient current passes from anode to cathode to sustain the two internal transistors' mutual conductance. A very simple demonstration circuit proves this hysteretic behavior:

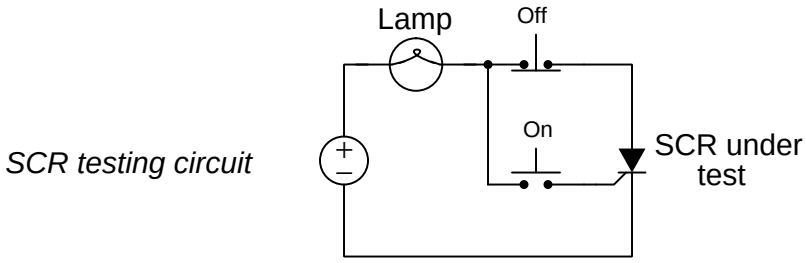

Momentarily pressing the normally-open "On" pushbutton switch fires the SCR and energizes the lamp. Once latched, the lamp remains illuminated until the normally-closed "Off" pushbutton is pressed.

This circuit not only proves the principle of SCR latching, but it also serves as a convenient means of *testing* an SCR for proper function. Given a source voltage and load current within the SCR's ratings, the circuit has every reason to function as previously described. An important design

element of this test circuit is load current, as even a good SCR may fail to latch if the load current is less than the SCR's rated *holding current*. A lamp with too much filament resistance may not draw enough current to keep an SCR latched when gate current ceases, thus giving the false impression of a bad (unlatchable) SCR in the test circuit. Holding current values for different SCRs should be available from the manufacturers. Typical holding current values range from 1 milliAmperes to 50 milliAmperes or more for larger units.

SCRs are, after all, controlled *rectifiers*, and as such they are more commonly found in AC-powered circuits than DC-powered circuits. If we replace the DC voltage source in our test circuit with an AC voltage source (and add a diode in series with the "on" switch so that the gate-cathode PN junction of the SCR is never reverse-biased), we will notice different behavior:

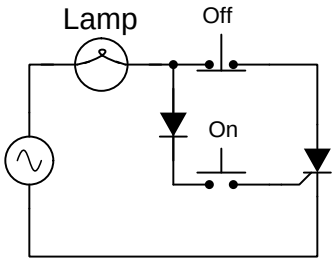

In this circuit, the "Off" pushbutton switch serves no useful purpose, because the lamp de-energizes every time the "On" pushbutton is released. At first this may seem strange given the clear latching behavior of the DC-sourced test circuit, but it makes sense once we carefully consider the *alternating* nature of the AC source. Twice per cycle, the AC naturally reverses direction, and at these points in time the SCR "drops out" and requires re-triggering to turn on again.

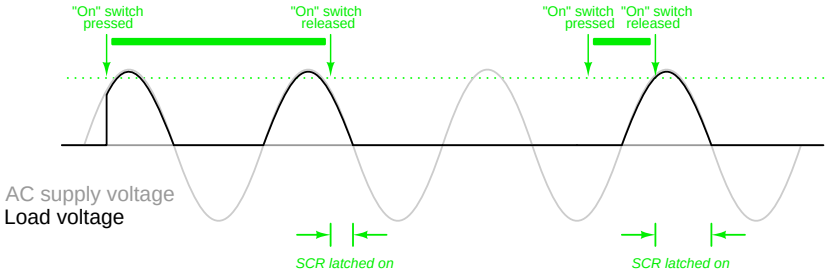

For certain high-voltage applications, special thyristors are manufactured with *optical* triggering capability, called *Light-Triggered Thyristors* (*LTT*s). Instead of relying on an electrical pulse at a gate terminal to trigger, an optical pulse signal guided to the core of the semiconductor structure of the thyristor via a glass fiber-optic cable directly excites one of the PN junctions. Glass fiber-optic cables, of course, are entirely non-conducting to electricity, and so this design of thyristor provides complete electrical isolation between the high-voltage power circuit and the triggering control circuit (using a light source such as an LED or a semiconductor laser to send optical pulses to the LTT).

Thyristor *triggering* (i.e. turning on) tends to be easier[4] than thyristor *commutation* (i.e. turning

[4] In fact, a bare SCR is so easily triggered that typically gate currents in the *microAmpere* range are sufficient.

off). This is due to the positive feedback design of PNPN-type thyristors, where the controlled (load) current acts to hold the device in its "on" state. Since thyristors such as SCRs are often used in high-current applications, the amount of current available through the load to latch the thyristor into its "on" state is normally quite large, and must be either diverted or halted by some other means in order for the thyristor to turn off. In the example circuit shown with an AC source, the natural cessation of current every time the AC source crosses zero is called *line commutation* or *natural commutation*: the alternating nature of the source solves the problem of commutation for us.

Forcing a thyristor to turn off when full load current is still present is a more challenging matter. SCRs, for example, are extremely difficult if not impossible to commutate using a reverse-polarity pulse signal between gate and cathode (sometimes called an *extracting gate current*; i.e. a current exiting the gate terminal, diverting enough of the current from the upper equivalent transistor to cause the lower equivalent transistor to turn off). Specially-designed thyristors called *Gate Turn-Off*, or *GTO*, thyristors are manufactured for the purpose of being able to turn both on and off via gate current, but even the GTO which is designed to be easier to commutate via the gate terminal than a regular SCR requires negative gate currents as large as one-fifth to one-third the anode (load) current! A simple DC-powered demonstration circuit for a GTO thyristor appears here:

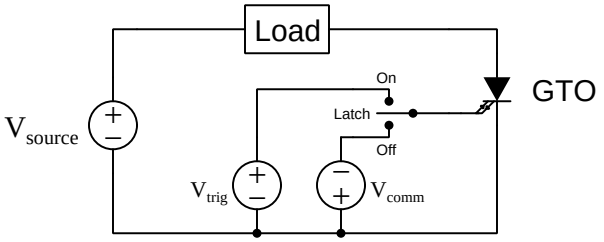

This can be problematic, as small amounts of "noise" current at the gate of such an SCR may unintentionally trigger it. Many SCRs are therefore manufactured to have *de-sensitized* gates requiring more substantial current to trigger, in order to avoid false triggering.

When two SCRs are used to exclusively switch DC power between two loads, a common method of forced thyristor commutation uses the stored energy within a capacitor to force one SCR to commutate when the other is triggered:

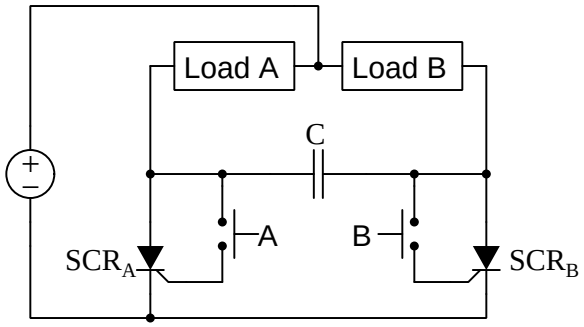

A thought experiment reveals how this circuit functions. Imagine SCR$_A$ conducting and SCR$_B$ non-conducting, energizing load A while maintaining load B in the de-energized state. In this condition, capacitor C drops load A's full voltage, positive on right and negative on left. If someone presses pushbutton B, SCR$_B$ will trigger on and cause the capacitor (acting as a voltage source) to deliver a surge of current from left to right. This surge diverts load A's current through the capacitor and down to ground through SCR$_B$ while impressing a reverse bias across SCR$_A$ (+ on cathode and − on anode). The surge of current through SCR$_B$ secures it in the "on" state while the reverse bias across SCR$_A$ forces it off. Load A's current decays to zero as the capacitor accumulates a voltage (+ on left, − on right), while Load B's current remains steady. Load B will remain energized and load A de-energized until someone presses pushbutton A.

SCRs find application in many high-power electronic circuits because once triggered they will force themselves into a condition of saturation which means low voltage drop and therefore low power dissipation when conducting high values of current. An example of an SCR built for very high-current applications (up to 500 Amperes!) appears in the photograph below:

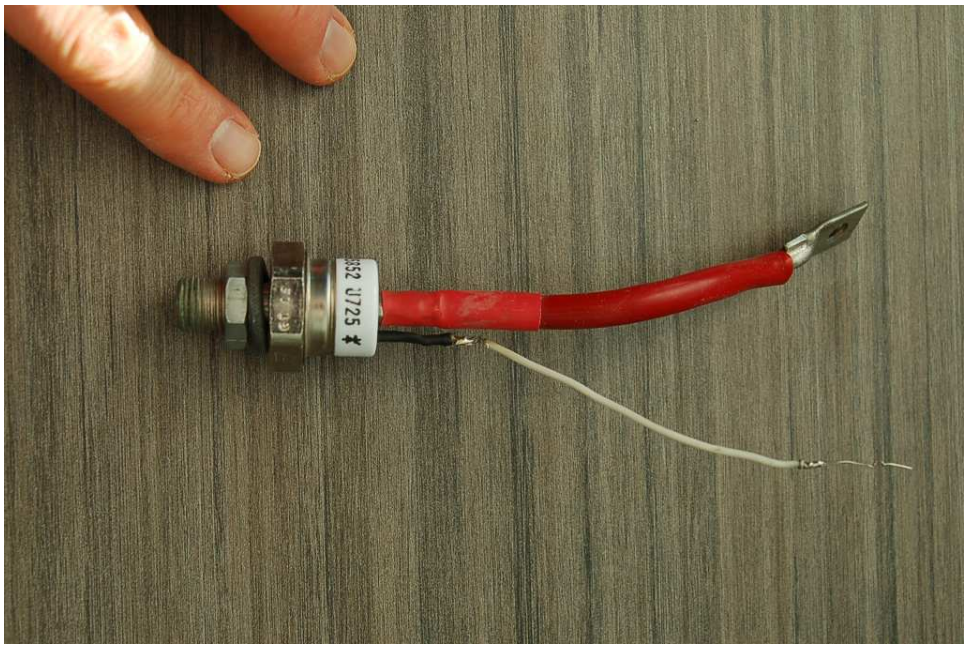

The heavy-gauge red-insulated cable connects to the cathode (K) on this SCR, while the light-gauge white wire connects to the gate (G). The Anode (A) is not a wire connection at all, but rather the threaded metal base, which means this device is intended to *bolt* to a solid metal receiving piece functioning as both a heat sink and an electrical bus.

3.5 TRIAC thyristors

Just as we created a DIAC by paralleling two reverse-facing Shockley diodes, we may create a new type of controlled thyristor by paralleling two reverse-facing SCRs. The resulting device is called a *TRIAC*:

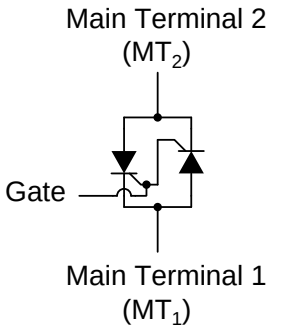

TRIAC equivalent circuit

TRIAC schematic symbol

Because individual SCRs are more flexible[5] to use in advanced control systems, they are more commonly seen in circuits like motor drives, while TRIACs are usually seen in simple, low-power applications like residential lamp dimmer switches and electric hand tool speed controls.

[5] One example of flexibility using SCRs rather than TRIACs, is that by using pairs of SCRs one for the positive half-cycle and another for the negative half-cycle, it becomes possible to independently control each half-cycle of the waveform, which can be very useful. A TRIAC cannot do the same without relying on precise timing to fire the gate during selected half-cycles.

From the equivalent circuit diagram shown earlier, one might think that main terminals 1 and 2 were interchangeable. They are not! Although it is helpful to imagine the TRIAC as being composed of two SCRs joined together, it in fact is constructed from a single piece of semiconducting material, appropriately doped and layered. The actual operating characteristics may differ slightly from that of the equivalent model. This fact is proven by contrasting two simple TRIAC test circuits, one that works and one that doesn't.

The following circuit will function to control power to the lamp using a relatively small triggering current at the TRIAC's gate. Pressing the normally-open pushbutton switch energizes the lamp, and releasing the pushbutton switch allows the TRIAC to drop out at the next zero-crossing point of the AC wave:

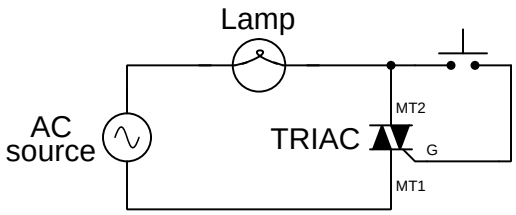

Suppose we were to swap the two main terminals of the TRIAC around. According to the equivalent circuit diagram shown earlier in this section, the swap should make no difference and the following circuit *ought* to work:

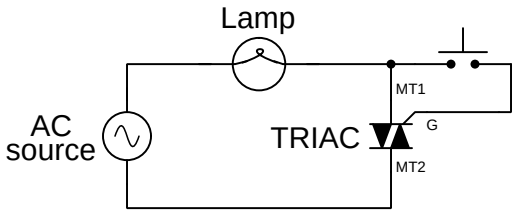

However, if this circuit is built, it will be found that it does not work! The load will receive no power, the TRIAC refusing to fire at all regardless of pushbutton status. The key to successfully triggering a TRIAC is to make sure the gate receives its triggering current from the *main terminal 2* side of the circuit (the main terminal on the opposite side of the TRIAC symbol from the gate terminal). Identification of the MT_1 and MT_2 terminals must be done via the TRIAC's part number with reference to a data sheet or book.

A more detailed characterization of TRIAC firing requirements is often represented by *operating quadrant*. This odd-sounding phrase originates from a two-dimensional plot of main terminal 2 voltage versus gate voltage (both with reference to MT1), with each quadrant of the graph representing a different combination of polarities:

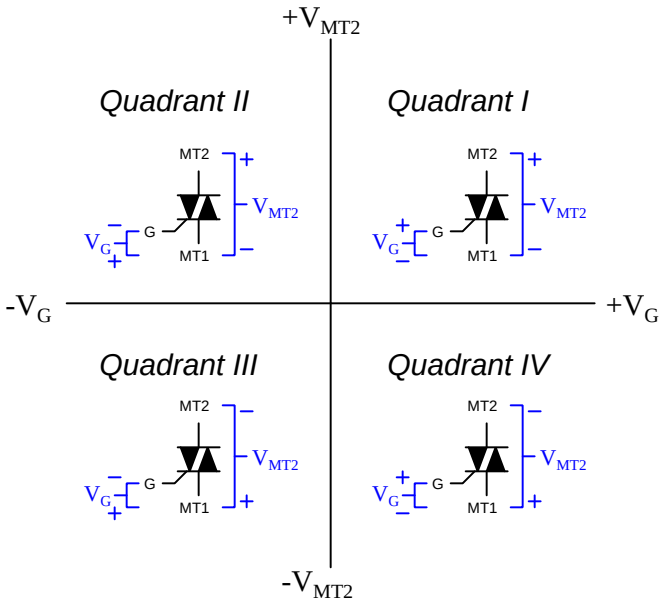

TRIAC datasheets will commonly list triggering parameters such as gate voltage and gate current separately for each quadrant of operation, with quadrants I and III typically requiring the least gate signal energy to successfully trigger.

3.6 Thyristor gate sensitivity

Thyristors differ from other active semiconductor devices primarily in their ability to "latch" into a conductive state once triggered. For gated thyristors such as SCRs and TRIACs, the amount of gate current required for triggering is very small compared to the amount of current the device is rated to conduct to a load. The positive-feedback nature of a thyristor means that the conducted load current itself acts to hold the device in its conductive state, taking over the role previously fulfilled by the gate current.

This positive-feedback behavior means the amount of gate current required to trigger an SCR or TRIAC under the right conditions may be quite trivial, less than a single microAmpere. Such sensitivity can pose a problem for these devices when used within power-control applications, since high voltages and high currents tend to make coupled "noise" inevitable. If a surge of high current and/or a transient voltage pulse at one thyristor happens to induce enough of a current through the gate terminal of an adjacent thyristor to trigger it, the result could be catastrophic!

One way to "de-sensitize" a thyristor in order to mitigate this problem is to connect an internal resistance across the gate's PN junction, like this:

De-sensitized SCR

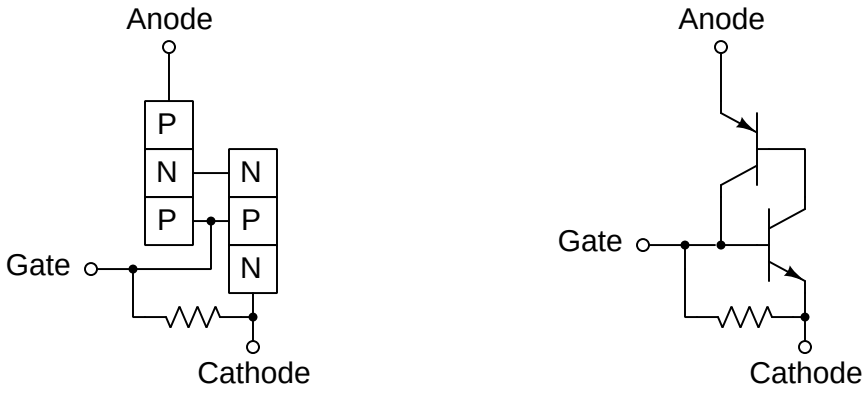

Crude diagram *Equivalent schematic*

Now there must be enough gate current flowing through this internal resistance to produce a 0.7 Volt drop before the gate-cathode PN junction becomes forward-biased to trigger the SCR into conduction. Without that resistance in place, *any* current passing through the gate-cathode PN junction is liable to trigger the device.

Gated thyristors are often labeled as being either *sensitive-gate* devices or *standard-gate* devices, to differentiate between the non-resistor and resistor versions. Again, sensitive-gate thyristor gate trigger currents are often less than 1 mA in contrast to standard-gate trigger currents which may be tens of milliAmperes.

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 Thyratron inverter patent

A predecessor to the solid-state thyristor was the gaseous-tube *thyratron*, similar in construction to a vacuum tube except with a low-pressure gas fill rather than a hard vacuum. The ionization effect of the filling gas provided a similar hysteresis effect for the thyratron tube as the interconnected-transistor structure provides hysteresis for a thyristor: once triggered, a thyratron tube would continue to conduct electricity until something *commutated* it.

The following diagram comes from a United States patent (number 2,784,365 granted 5 March 1957) submitted by inventors Ronald William Fenemore and Charles John Quartly of England, illustrating a circuit designed in invert DC into three-phase AC:

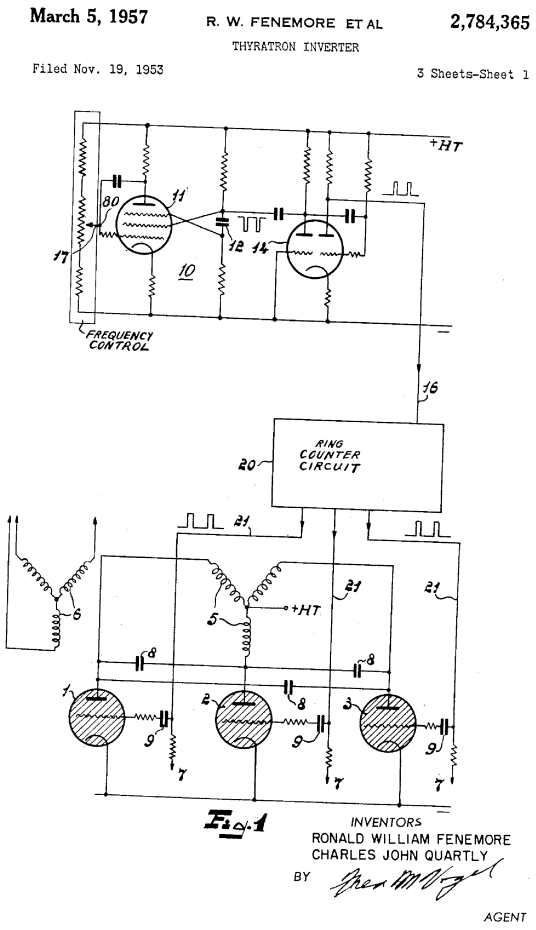

Three thyratron tubes near the bottom of the diagram handle the power switching, each one dedicated to a phase of a Wye-connected transformer primary. The positive pole of the DC supply (+HT) connects to the center of the transformer primary Wye, while the negative pole connects to

the cathodes of all three thyratron tubes. A *ring counter* circuit sends triggering pulses to the grid electrodes of the thyratron tubes, equivalent to the gate terminals of three SCRs.

Like SCRs, thyratrons are impractical to commutate via a negative pulse applied to the gate (grid) terminal, and instead must be commutated by other means. In this case, the conducting tube is commutated by firing another tube, and the key to the success of this scheme is the three capacitors (labeled 8 on the diagram). The inventors describe how this works, which could have easily described how a capacitor-commutated SCR circuit would function if SCRs were invented at the time:

> Three capacitors 8 are connected in a ring or delta between the anodes of the thyratrons 1, 2, 3 and operate in such manner as to extinguish automatically one thyratron when the next thyratron is fired. For example, if tube 2 is conducting at a particular instant and tube 3 is then fired by a pulse, the anode of the latter tube will drop in potential due to the current taken. This potential drop will be passed momentarily to the anode of tube 2 and will be sufficient to extinguish tube 2 in the absence of a firing pulse. The anode drop will also be communicated to tube 1, but ineffectively, since tube 1 will already be in the extinguished condition.

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 Lightning bolts, relaxation oscillators, and thyratron tubes

If you've ever witnessed a lightning storm, you've seen electrical hysteresis in action. The action of wind and/or rainfall accumulates tremendous static electric charges between cloud and earth, and between clouds as well. Electric charge imbalances manifest themselves as high voltages, and when the electrical resistance of air can no longer withstand the intense electric fields, the air molecules ionize and become electrically conductive, allowing huge surges of current between opposing poles of electrical charge which we call *lightning*.

The buildup of high voltage by wind and rain is a fairly continuous process, the rate of charge accumulation increasing under the proper atmospheric conditions. However, lightning bolts are anything but continuous: they exist as momentary surges rather than continuous discharges. Why is this? Why don't we see soft, glowing lightning *arcs* instead of violently brief lightning *bolts*? The answer lies in the nonlinear (and hysteretic) resistance of air.

Under ordinary conditions, gases such as air have extremely high electrical resistance. This resistance is so high, in fact, that we typically treat its resistance as infinite and electrical conduction through the air as negligible. When a sufficient electric field intensity develops across a sample of gas, though, its electrical properties change: electrons become "stripped" from their normal positions around their respective atoms and are liberated to function as mobile electric charge carriers. In this state, air is considered to be *ionized* and is referred to as a *plasma*[1] rather than a normal *gas*.

An electric current moving through the plasma dissipates energy because the plasma is not a perfect conductor. In the case of lightning, this dissipated energy translates into extremely high temperatures. High temperatures are *also* sufficient to convert gaseous air into a plasma independent of any electric fields. As the voltage between cloud and earth, or between cloud and cloud, decreases as the electric charge imbalance neutralizes due to the current of the lightning bolt, the heat dissipated by the bolt maintains the air path in a plasma state to keep its resistance low. In other words, the lightning bolt *remains* in its plasma state until the temperature drops low enough. When this occurs, the air returns to its normal (gas) state and stops conducting current.

Note the air's electrical conductivity is hysteretic: when not conducting electricity it *remains an insulator* until voltage builds up past a critical threshold point; after becoming a plasma it *remains a conductor* until voltage falls below a lower critical threshold point. Once "turned on" the air tends to stay on, and once "turned off" it tends to stay off. This hysteresis, combined with a steady buildup of voltage due to the electrostatic effects of wind and rain, explains the phenomenon of lightning as brief bursts.

[1] This usage of the word "plasma" is not to be confused with the medical term (meaning the fluid portion of blood), but is a fourth state of matter, the other three being solid, liquid, and vapor (gas). Plasma is a relatively good conductor of electricity, its specific resistance being much lower than that of the same substance in its gaseous state.

In electronic terms, lightning functions as a spectacular *relaxation oscillator*. Oscillators are electronic circuits that produce an oscillating (AC) voltage from a steady supply of DC power. A relaxation oscillator is one that works on the principle of a charging capacitance that is suddenly discharged every time its voltage reaches a critical threshold value. One of the simplest relaxation oscillators in existence is comprised of three components (not counting the DC power supply): a resistor, capacitor, and neon lamp:

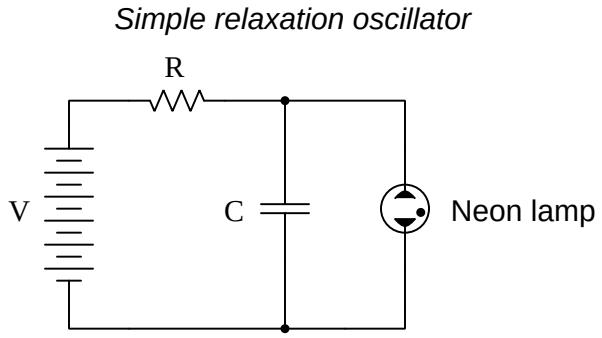

Neon lamps are nothing more than two metal electrodes inside a sealed glass bulb, separated by the neon gas inside. At room temperatures and with no applied voltage, the lamp has nearly infinite resistance. However, once a certain threshold voltage is exceeded (this voltage depends on the gas pressure and geometry of the lamp), the neon gas will become ionized (turned into a plasma) and its resistance dramatically reduced. In effect, the neon lamp exhibits the same characteristics as air in a lightning storm, complete with the emission of light as a result of the discharge, albeit on a much smaller scale.

The capacitor in the relaxation oscillator circuit shown above charges at an inverse exponential rate determined by the size of the resistor. When its voltage reaches the threshold voltage of the lamp, the lamp suddenly "turns on" and quickly discharges the capacitor to a low voltage value. Once discharged, the lamp "turns off" and allows the capacitor to build up a charge once more. The result is a series of brief flashes of light from the lamp, the rate of which dictated by battery voltage, resistor resistance, capacitor capacitance, and lamp threshold voltage.

While gas-discharge lamps are more commonly used as sources of illumination, their hysteretic properties were exploited in slightly more sophisticated variants known as *thyratron tubes*. Essentially a gas-filled triode tube (a triode being a three-element vacuum electron tube performing much a similar function to the N-channel, D-type IGFET), the thyratron tube's threshold voltage could be adjusted by a small DC control voltage applied between grid and cathode.

(Simple) Thyratron control circuit

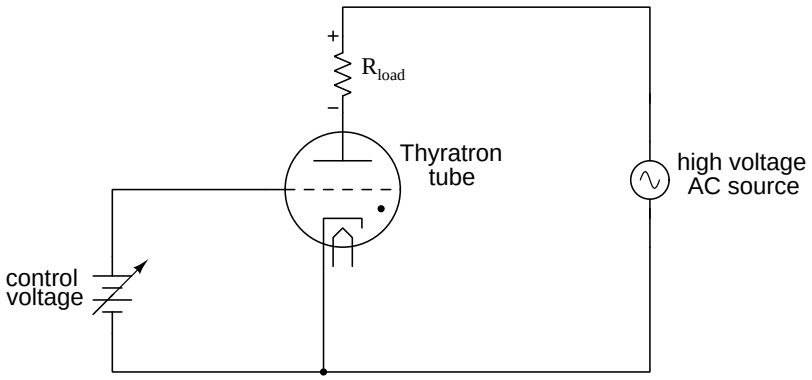

In essence, thyratron tubes were *controlled* versions of neon lamps built specifically for switching current to a load. The dot inside the circle of the schematic symbol indicates a gas fill, as opposed to the hard vacuum normally used in most other electron tube designs. In the circuit shown above, the thyratron tube allows current through the load in one direction (note the polarity across the load resistor) when triggered by the small DC control voltage connected between grid and cathode. Note that the load's power source is AC, which provides a clue as to how the thyratron turns off after it's been triggered on: since AC voltage periodically passes through a condition of 0 Volts between half-cycles, the current through an AC-powered load must also periodically halt. This brief pause of current between half-cycles gives the tube's gas time to cool, letting it return to its normal "off" state. Conduction may resume only if there is enough voltage applied by the AC power source (some other time in the wave's cycle) *and* if the DC control voltage allows it.

An oscilloscope display of load voltage in such a circuit would look something like this:

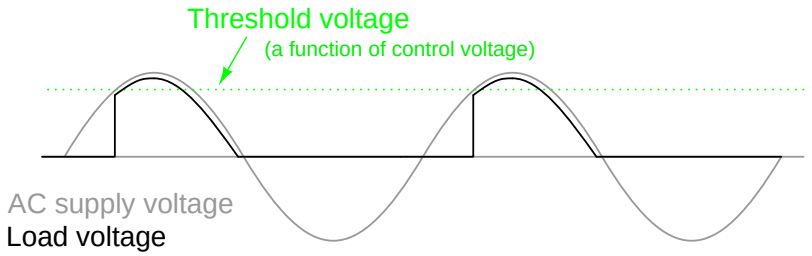

As the AC supply voltage climbs from zero Volts to its first peak, the load voltage remains at zero (no load current) until the threshold voltage is reached. At that point, the tube switches "on" and begins to conduct, the load voltage now following the AC voltage through the rest of the half cycle. Notice how there is load voltage (and thus load current) even when the AC voltage waveform

has dropped below the threshold value of the tube. This is hysteresis at work: the tube stays in its conductive mode past the point where it first turned on, continuing to conduct until there the supply voltage drops off to almost zero Volts. Because thyratron tubes are one-way (diode) devices, there is no voltage across the load through the negative half-cycle of AC. In practical thyratron circuits, multiple tubes arranged in some form of full-wave rectifier circuit to facilitate full-wave DC power to the load.

It would be possible to apply thyratron technology to a relaxation oscillator circuit and control the frequency with a small DC voltage between grid and cathode, making a crude voltage-controlled oscillator, otherwise known as a *VCO*. Relaxation oscillators tend to have poor frequency control, not to mention a very non-sinusoidal output, and so they exist mostly as demonstration circuits (as is the case here) or in applications where precise frequency control isn't important. Consequently, this use of a thyratron tube would not have been a very practical one.

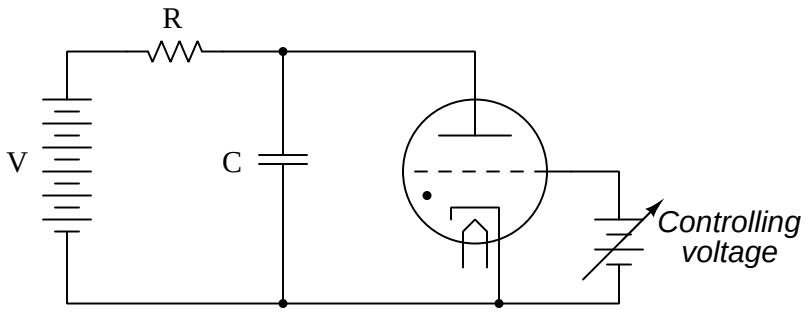

I speak of thyratron tubes in the past tense for good reason: modern semiconductor components have obsoleted thyratron tube technology for all but a few very special applications. It is no coincidence that the word *thyristor* bears so much similarity to the word *thyratron*, for this class of semiconductor components does much the same thing: use *hys*te*r*etically switch current on and off.

5.2 Latch-up in integrated circuits

An often bewildering phenomenon for students of electronics is something called *latch-up*, known to plague certain forms of integrated circuit. "Latch-up" describes a condition where the IC suddenly ceases to function properly, and instead draws an unreasonably high amount of current from the DC power source, sometimes to the point of catastrophic failure. What is more, this sudden and unexpected high-current draw refuses to stop unless and until the DC power source is interrupted.

Close examination of the afflicted IC's internal schematic diagram is unhelpful, as no apparent cause for latch-up will ever be evident upon inspection. This is because latch-up is a phenomenon rooted in *parasitic* transistor structures within the IC itself resulting from the monolithic form of the circuit (i.e. being fabricated on a single wafer of semiconducting material). In fact, if the IC were constructed from a large collection of discrete components on a printed circuit board, latch-up could never occur!

In order to understand what latch-up is and why it occurs, we must first review a class of semiconductor devices called *thyristors*. A "thyristor" by definition is any semiconductor device manifesting *hysteresis*: the tendency to remain on once triggered on and to remain off once turned off. Thyristors include devices such as DIACs, TRIACs, silicon-controlled rectifiers (SCRs), Shockley diodes, and unijunction transistors. Among these various thyristor devices, the most elementary is the Shockley diode, illustrated below:

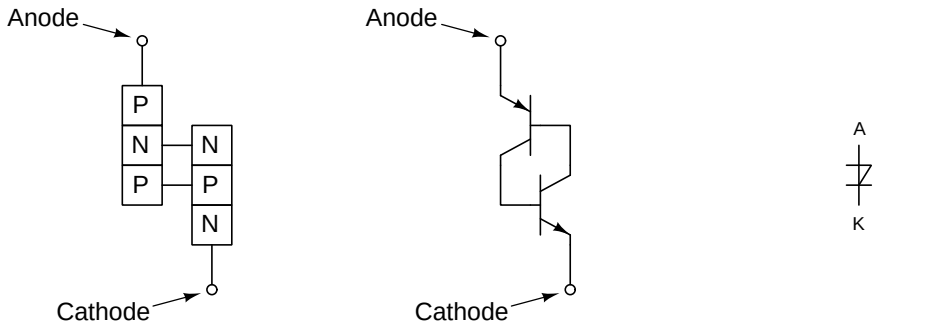

Crude diagram *Equivalent schematic* *Schematic symbol*

If we examine the equivalent Shockley diode schematic consisting of two interconnected bipolar junction transistors, one NPN and the other PNP, we find that it should not conduct current between the anode and cathode terminals. The only way for one of these transistors to ever turn on is to first initiate base current through it, but the interconnection with the other transistor makes base current for one contingent on collector current of the other, which cannot occur unless and until that other transistor's base is also activated by current. In other words, we have a "Catch-22" scenario where the NPN transistor cannot turn on until the PNP transistor turns on, and vice-versa. Therefore, Shockley diodes – and by extension all thyristors – naturally assume an "off" state and tend to maintain that state even as external voltage is applied across the two end-terminals.

If, however, we imagine *something* turning on one of these transistors, the resulting current would immediately turn on the other transistor, which in turn would hold the first transistor in its "on" state after the initiating event passed. Therefore, Shockley diodes – and by extension all

thyristors – tend to remain in their "on" states once triggered. Turning off a thyristor necessitates some external interruption of current. That "something" initially triggering the thyristor could be excessive applied voltage resulting in BJT breakdown, an extremely rapid rate-of-voltage change applied across the thyristor resulting in displacement current "through" one of the transistor's otherwise non-conducting depletion regions, excessive device temperature, a photon of light striking the depletion region and liberating an electron-hole pair, etc. Thyristors designed as controllable devices (e.g. SCRs, TRIACs, unijunction transistors) provide a third terminal to apply a triggering current.

Monolithic integrated circuit construction – where the entire circuit is built on a single wafer of semiconducting material – often results in parasitic PN junctions, and sometimes even parasitic NPN and PNP junction sets. If parasitic NPN and PNP junctions happen to coexist on the silicon die to form the equivalent of a Shockley diode thyristor structure, we will find that circuit susceptible to the unfortunate phenomenon of latch-up. These parasitic thyristors exist within many different types of integrated circuits, from logic gates to amplifiers to comparators to timers such as the popular model 555. They are guaranteed to exist within integrated circuits using complementary pairs of N-channel and P-channel MOSFETs as final "push-pull" output driver stages.

Sometimes the potential for latch-up is communicated in the form of a warning within IC datasheets, but other times not. Inconsistent and incomplete documentation, added to the fact that an examination of an IC's internal schematic does not necessarily reveal the presence of parasitic thyristors, makes latch-up risk a difficult thing to predict.

Let's begin our exploration of these parasitic structures by closely examining one of the most basic building-blocks of CMOS (Complementary Metal-Oxide Semiconductor) circuitry, the so-called *inverter*. A CMOS inverter network is nothing more than a P-channel enhancement-type MOSFET connected to an N-channel enhancement-type MOSFET, as shown below both in schematic form as well as semiconductor wafer cross-section. This complementary MOSFET pair acts as a "push-pull" transistor amplifier, capable of either sourcing current to a load or sinking current from one:

Schematic of an integrated CMOS inverter

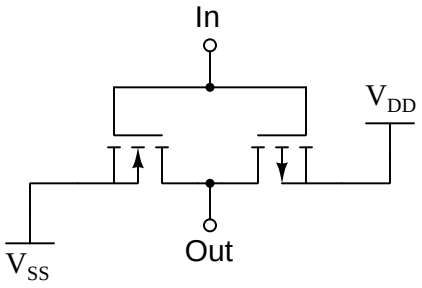

Cross-section of an integrated CMOS inverter

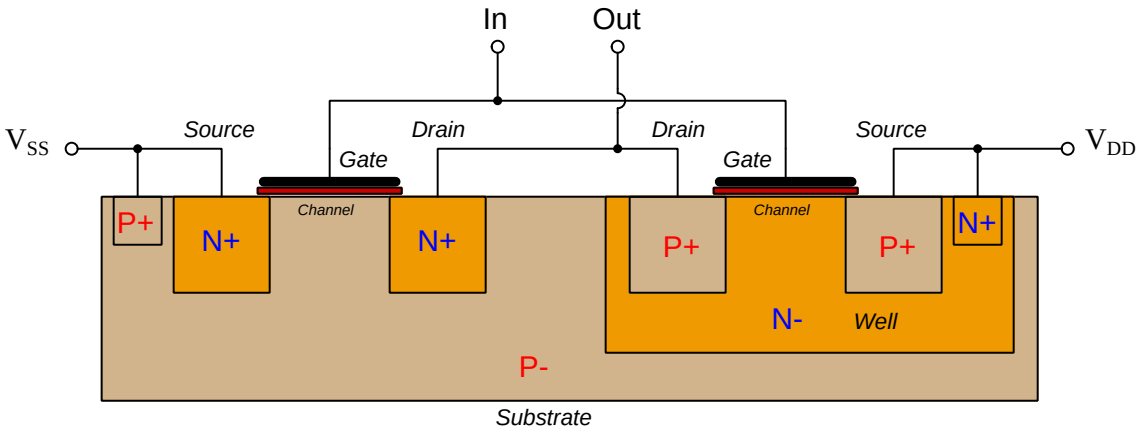

Each of the doped semiconductor layers in the cross-section are noted with their respective dopant concentration levels, a "+" meaning a high concentration of dopant and "−" meaning a low concentration of dopant. Strongly-doped P-type and N-type layers are shown as P+ and N+, respectively. Lightly-doped P-type and N-type layers are shown as P− and N−, respectively. Note how the source and drain regions are all heavily-doped in order to provide a rich reservoir of majority charge carriers to fill in each of the channels when prompted by a sufficient gate-to-substrate electric field.

Now we will include within this cross-sectional illustration some the parasitic transistors[2] and resistances formed by the P and N type doped regions:

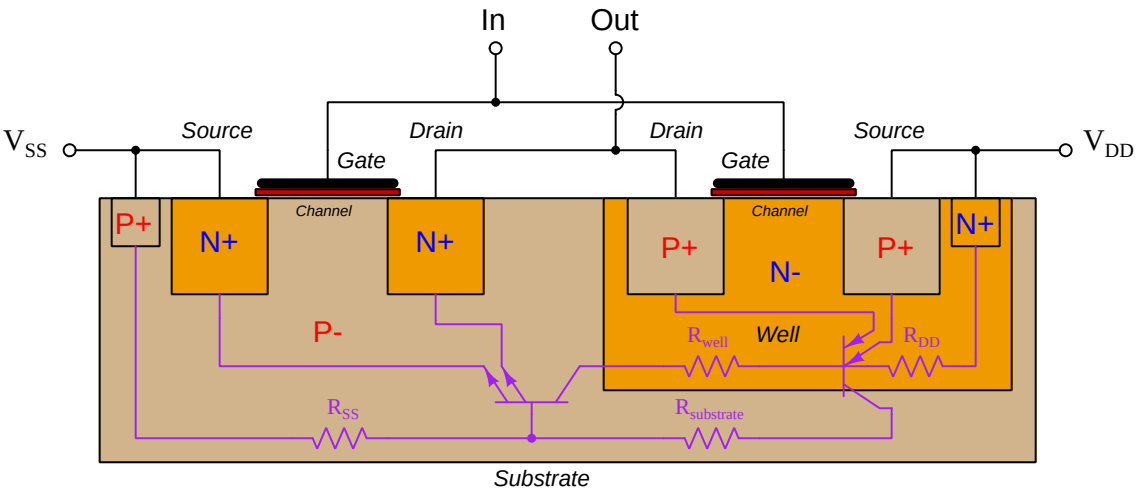

Including these parasitic components into the CMOS inverter schematic diagram:

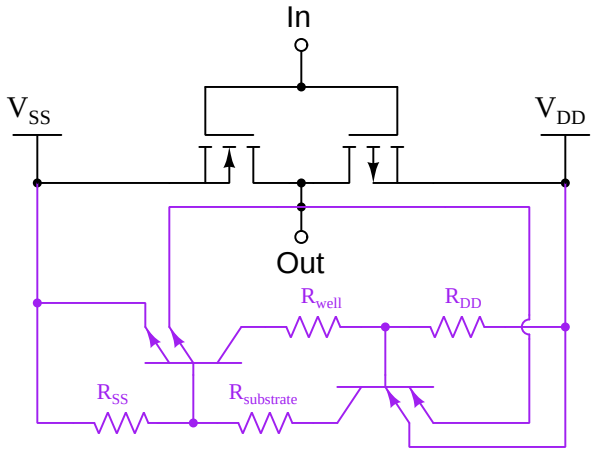

Now it should be clear that a Shockley diode structure exists within the CMOS inverter's integrated circuit. Furthermore, with each bipolar transistor having two emitters, one each connected to the inverter's output terminal, that output terminal becomes a pathway to externally trigger this thyristor structure and thereby throw the IC into a state of latch-up.

[2]It is helpful to recall that the emitter region of any bipolar junction transistor is more heavily doped than the collector, in order to provide an abundance of charge carriers to inject into the base region as minority carriers. This is why the emitters of the PNP are P+ doped and the collector P− doped, and why the emitters of the NPN are N+ doped and the collector N− doped. Multiple emitters in a BJT simply mean more than one location from which charge carriers may be injected into the base and thereby render the transistor conductive between emitter and collector.

Triggering a latch-up condition is as simple as something external to the inverter IC driving its output terminal below V_{SS} potential or above V_{DD} potential. In the next example we will see a below-V_{SS} potential initiating latch-up:

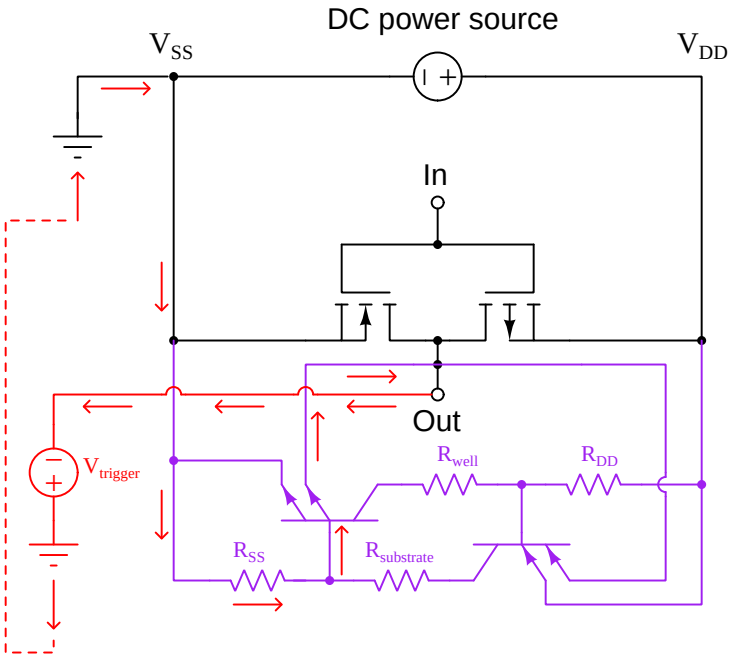

In this scenario, external source $V_{trigger}$ forced the inverter's output terminal to a potential lower than ground (V_{SS}), and this in turn prompts a current through one of the emitters of the parasitic NPN transistor, injecting electrons into its P-type base region and turning it on. Once the NPN transistor turns on, its collector current permits base current through the PNP transistor, turning it on as well, and thus the two parasitic transistors latch each other into conductive states.

Once latched on, the parasitic thyristor creates a V_{DD}-to-V_{SS} short-circuit. Excessive current now passes through the DC power source and through the parasitic resistances of R_{well} and $R_{substrate}$, possibly damaging the entire integrated circuit by excessive heating. Only an interruption of the DC power source will halt the latch-up condition.

Next, we will see an above-V_{DD} potential initiating latch-up:

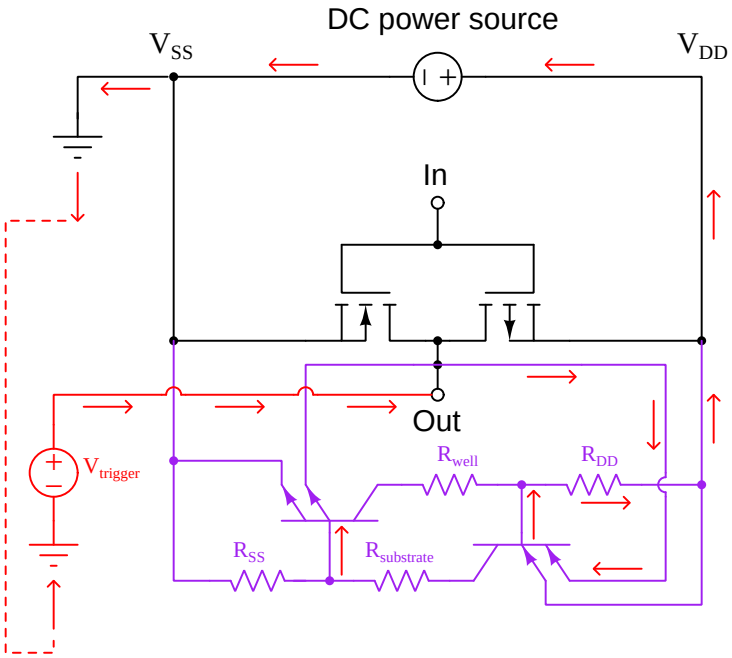

In this scenario, external source $V_{trigger}$ forced the inverter's output terminal to a potential higher than that of the DC source voltage (V_{DD}), and this in turn prompts a current through one of the emitters of the parasitic PNP transistor, injecting holes into its N-type base region and turning it on. Once the PNP transistor turns on, its collector current permits base current through the NPN transistor, turning it on as well, and thus the thyristor formed by the two parasitic transistors becomes latched in its "on" state.

Again, once the parasitic thyristor is latched in its "on" state it will create a V_{DD}-to-V_{SS} short-circuit. Excessive current will now pass through the DC power source and through the parasitic resistances of R_{well} and $R_{substrate}$, possibly damaging the entire integrated circuit by excessive heating. Only an interruption of the DC power source will abort the latch-up condition.

Clearly, one of the causes of latch-up is some external electrical source attempting to force the output terminal either above V_{DD} potential or below V_{SS} potential, and as such we should strive to avoid such possibilities in all our circuit designs.

Another potential cause of latch-up is if the load driven by an IC output is too inductive. Inductance stores energy in magnetic fields, and the amount of energy is a function of the current through the inductance. If an inductive load is energized by the output of an IC and then the IC switches state to de-energize that load, the load inductance will become a temporary current source and try to force current through the IC's output, often generating surprisingly high voltages in the process. Such voltage transients may easily trigger latch-up in the IC.

Another possible cause of latch-up is power supply rail instability. That is, if something is causing an IC's positive and/or negative rail potentials to vary, and that same IC connects to some other device with higher positive or lower negative power supply rails, this may constitute a triggering mechanism for latch-up. Consider the following example:

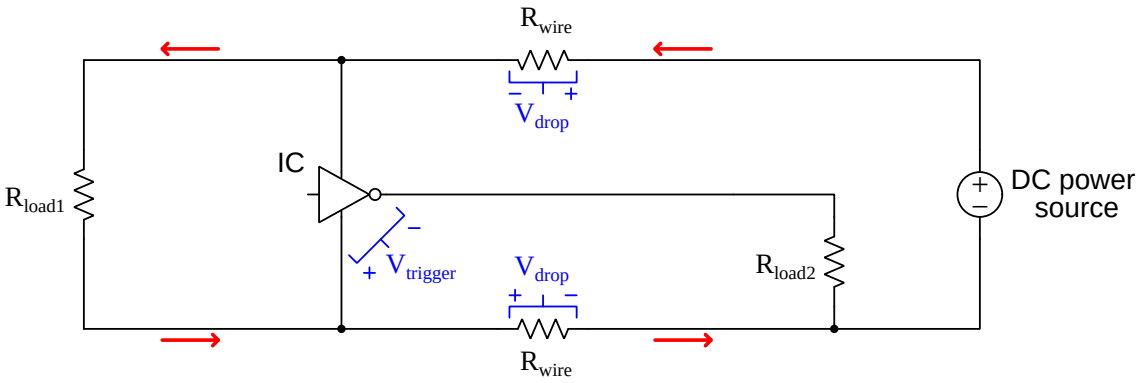

Here we see some voltage dropped along the conductors carrying current to a heavy load (R_{load1}), the effect of these voltage drops being to suppress the potential of the IC's positive rail terminal slightly below that of the power source's positive terminal, and to elevate the potential of IC's negative rail terminal slightly above that of the power source's negative terminal. In this case, with the IC's output connected to a second load (R_{load2}), we see that the voltage drop along the negative power conductor has the effect of making the IC's output terminal fall to a potential slightly less than its negative power terminal. As we have seen previously, this may be enough to trigger latch-up in that IC.

If R_{load2} happened to be connected to the positive conductor rather than the negative, we would similarly risk latch-up due to the load imposing a potential on the IC's output greater than that of its positive power terminal:

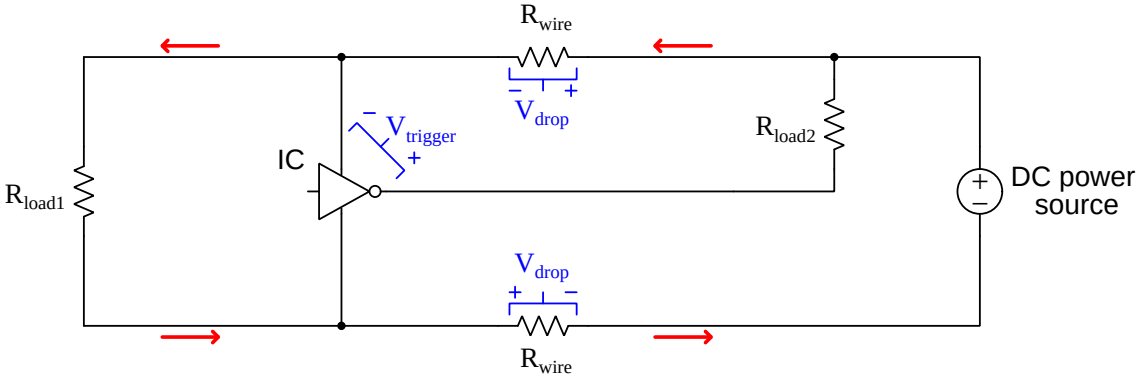

In either scenario the wiring voltage drop caused by the high current demand of R_{load1} risks latch-up in the IC by skewing its power-supply terminal potentials compared to the load it connects to. A

strategy for avoiding this problem is to provide high-current loads with their own dedicated power conductors, and not connect them to the same power conductors that sensitive ICs use for power delivery. This strategy is often referred to as *star-point grounding* or star-point power distribution:

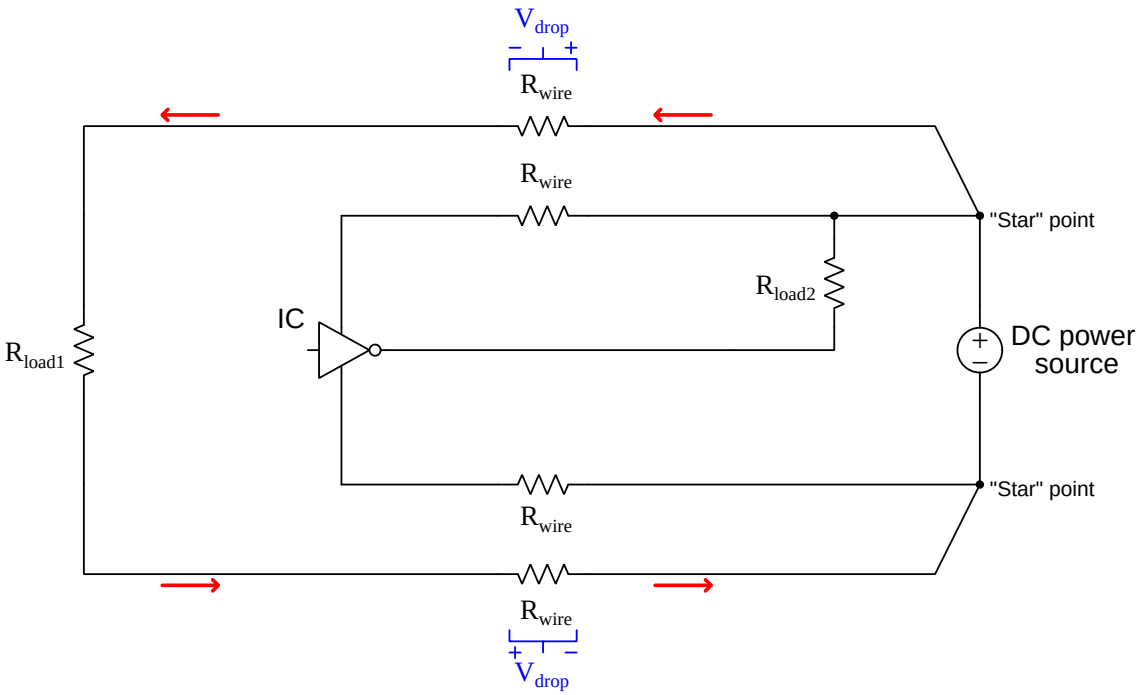

Star-point power connections do not eliminate voltage drop due to high load current, but they do isolate those voltage drops from affecting *other* loads such as ICs sensitive to latch-up.

Given the fact that Complementary MOSFET transistor pairs are so prevalent in modern integrated circuits (e.g. *all* CMOS digital logic) we must find mitigations for latch-up or else risk self-destruction of a great many ICs any time an output pin is subjected to even the briefest transient voltage outside of the DC power supply's rail potentials.

One such mitigation was invented by Alastair Stevenson of the Hughes Aircraft Company in 1978, consisting of an extra layer of implanted material in the semiconductor structure serving as a new collector of charge carriers. His patent ("CMOS Monolithic Integrated Circuit", US patent 4,173,767 filed 5 May 1978 and granted 6 November 1979) contrasts the "prior art" of CMOS inverter construction[3] (figures 1 through 3) against his innovation shown in figures 4 and 5:

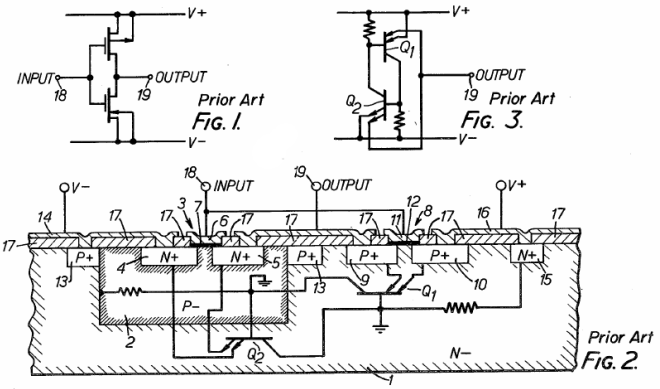

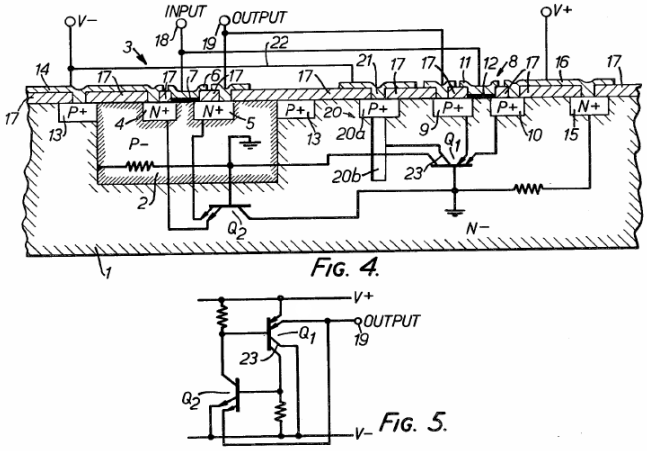

In Stevenson's design we see the addition of a heavily-doped P+ region (20a in figure 4) near the surface of the semiconductor die connecting with a deeper region of lightly-doped P− material (20b in figure 4) inset into the N− substrate. This P− region serves to collect holes injected as minority

[3] Stevenson's prior art example happens to use an N-type substrate rather than a P-type substrate as I showed earlier, thus necessitating a P-type well surrounding the N-channel MOSFET rather than the N-type well surrounding the P-channel MOSFET in my illustrations.

carriers into the N− substrate that might otherwise travel to the P− well region of the N-channel MOSFET, essentially adding a second collector to the PNP parasitic transistor. The purpose for the P+ region interfacing the metal surface conductor to the deeper P− inset is to avoid "rectifying" contact between metal and semiconductor surfaces, the heavy doping of that P+ region (20a) serving instead to form an "ohmic" contact with the metal. This additional collector region serves to divert current away from the parasitic NPN transistor that would otherwise turn it on and ensure latch-up once the PNP transistor was triggered by an offending potential at the output terminal.

Stevenson describes how such current-collecting channels inserted into the substrate need only be placed between complementary MOSFETs at device input or output terminals, not between every pair of complementary MOSFETs, due to the fact that only those connecting to device pins are susceptible to the triggering that cause latch-up:

> It will be appreciated that the circuit arrangement shown in FIG. 4 represents only one small portion of the integrated circuit which will in fact include many more complementary MOS transistors. However, channels such as 20 need not necessarily be included between all adjacent complementary transistors in the circuit, and preferably channels such as 20 are only included between complementary transistors that constitute input or output stages of the circuit since it is only these transistors which will be subject to spurious voltage fluctuations which could trigger a bipolar conductive latch up in the circuit.

A final note on the topic of latch-up is that this term is unfortunately mis-used to describe a phenomenon plaguing certain comparators and operational amplifiers. In this other meaning of the term "latch-up", the comparator or opamp output terminal assumes an incorrect potential because one or more of the input terminal potentials has approached or exceeded a power supply rail potential. This kind of "latch-up" is sometimes referred to as *pseudo latch-up* because unlike the thyristor-based latch-up we've been exploring here, pseudo latch-up immediately ceases as soon as the offending input potential returns to normal. In the "real" form of latch-up, only interruption of DC power will cause the latching action to cease.

Chapter 6

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?

6.1 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.

6.1.1 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.

6.1.2 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

Thought experiments as a problem-solving strategy

Hysteresis

Positive feedback

Breakdown

Holding current

Crowbar

False triggering

Commutation

Electrical source

Electrical load

6.1.3 SCR versus TRIAC circuits

Explain what happens in each of these circuits when the pushbutton switch is actuated and then released:

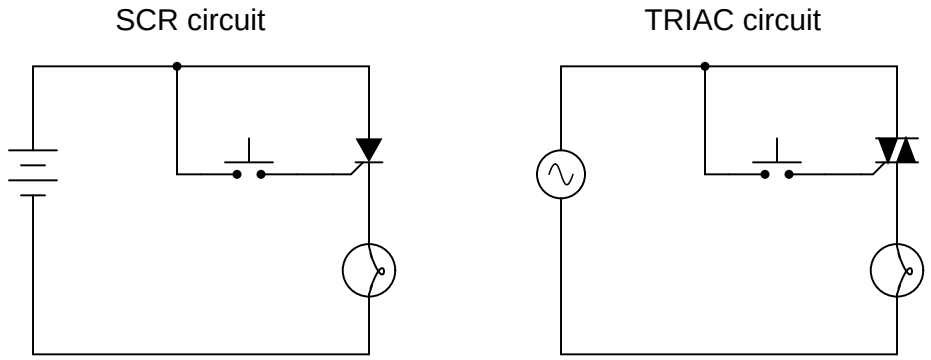

Which one of these circuits requires a second pushbutton switch to turn the lamp off? Where would this second switch need to be connected?

Challenges

- If an oscilloscope were connected across the load in each of these circuits, what would its display show?
- How would these two circuits function if both were sourced by DC?
- How would these two circuits function if both were sourced by AC?

6.1.4 Turning an SCR on and off

Explain what must be done to the SCR to make it turn on and send power to the light bulb:

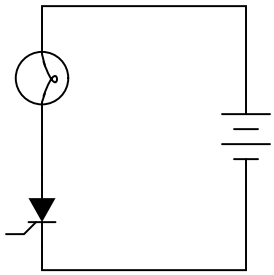

Next, explain what must be done to turn the SCR off so that the light bulb de-energizes.

Challenges

- Explain how SCR behavior differs from that of BJT behavior.

6.1.5 SCR voltage drop

When an SCR latches "on", it drops very little voltage between anode and cathode. Explain why this is, and what advantage this gives SCRs over transistors when conducting heavy load currents.

Challenges

- Explain why SCRs typically exhibit very fast turn-on times.

6.1.6 Sensitive-gate thyristors

Some SCRs and TRIACs are advertised as *sensitive-gate* devices. What does this mean? What is the difference between a "sensitive gate" SCR and an SCR with a "non-sensitive gate"?

Challenges

- Identify a practical application where a non-sensitive-gate thyristor would be preferable.

6.1.7 SCR-based amplifier

A student of electronics has just recently learned how to build audio amplifier circuits, and this inspires dreams of designing a super-powerful amplifier for a home entertainment system. One day, this student comes across a donation of electronic components from a local business, and in this donation are several industrial SCRs, rated at 20 Amperes each.

"Wow," says the student, "these components look like really big transistors, but they're rated for a lot of current. I could build a *huge* amplifier with these!"

The student approaches you for advice, because you've just recently learned how SCRs function in your electronics class. What do you tell the student, concerning the use of SCRs as audio amplification devices? How do you explain to this excited student that these devices will not work in an amplifier circuit?

Challenges

- Devise a demonstration circuit contrasting transistor behavior against SCR behavior, to make it clear to the enthusiastic student why their plan cannot work.

6.1.8 Different thyristor types

Identify at least three different types of thyristors (besides SCRs):

Challenges

- Which of these thyristors can *only* switch DC?
- Which of these thyristors can switch AC?
- Identify some unique characteristic of each thyristor listed.

6.1.9 UJT test circuit

Describe what happens to the UJT as the potentiometer is slowly adjusted upward to provide a variable voltage at point **A** in this circuit, starting from 0 Volts and ending at the trigger voltage V_P:

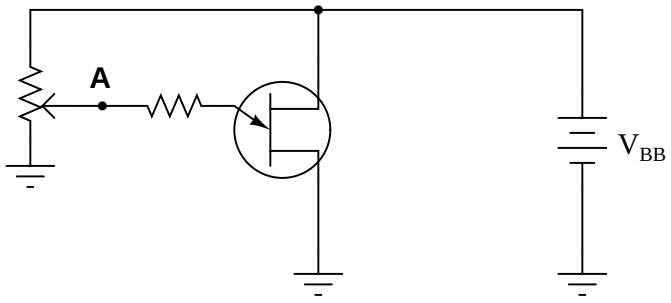

Now describe what must be done to the potentiometer to cause the UJT to turn back off again.

Challenges

- Explain how *hysteresis* plays a role in this circuit's behavior.

6.2 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.

6.2.1 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.

6.2.2 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.*

6.2.3 Voltage rate-of-change from oscillographs

First, identify at which of the numbered points is the waveform experiencing a positive rate-of-change versus a negative rate-of-change:

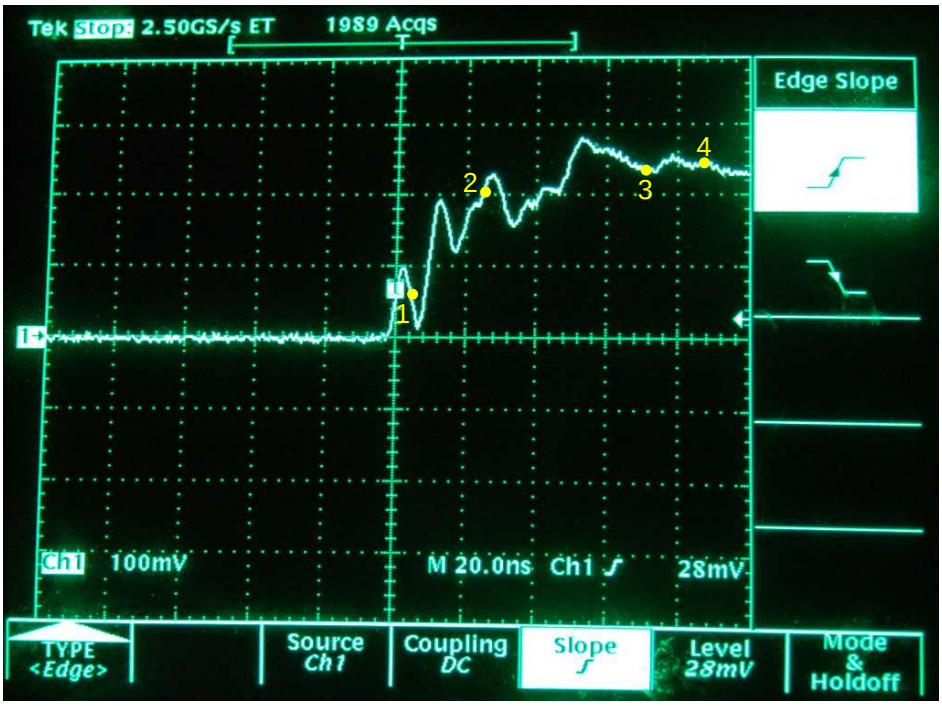

Next, estimate the voltage rate-of-change for this voltage signal at the points specified:

- Point #1 $\frac{dV}{dt} \approx$
- Point #2 $\frac{dV}{dt} \approx$
- Point #3 $\frac{dV}{dt} \approx$
- Point #4 $\frac{dV}{dt} \approx$

Challenges

- Identify points on the waveform where the rate-of-change is zero.
- Why, exactly, is a high voltage rate-of-rise troublesome for thyristors?

6.2.4 2N6403 specifications

Find a datasheet for a thyristor with the part number 2N6403, and answer the following questions:

- What type of thyristor is the 2N6403?
- What differentiates the 2N6403 from others in this series (e.g. 2N6401, 2N6402, etc.).
- Identify the holding current rating for this thyristor, and explain its significance.
- Identify the critical rate-of-rise for voltage, and explain its significance.
- Would a linear rise of V_{AK} from 0 Volts to 14 Volts over 0.2 μs be acceptable?
- Identify the typical gate trigger current for this thyristor. Would you classify this as a *sensitive-gate* or *standard* thyristor?
- How much voltage would you expect this thyristor to drop in the "on" state?

Challenges

- Suppose you needed a thyristor to build a circuit and the only one on hand had a "sensitive" gate. How could you intentionally "de-sensitize" the gate?

6.2.5 2N6344 specifications

Find a datasheet for a thyristor with the part number 2N6344, and answer the following questions:

- What type of thyristor is the 2N6344?
- Identify the holding current rating for this thyristor, and explain its significance.
- Identify the critical rate-of-rise for voltage, and explain its significance.
- Would a linear rise of V_{AK} from 0 Volts to 25 Volts over 11 μs be acceptable?
- Identify the typical gate trigger current for this thyristor. Would you classify this as a *sensitive-gate* or *standard* thyristor?
- How much voltage would you expect this thyristor to drop in the "on" state?

Challenges

- Suppose you needed a thyristor to build a circuit and the only one on hand had a "sensitive" gate. How could you intentionally "de-sensitize" the gate?
- Some of the parameters for this thyristor vary with its *operating quadrant*. What does this refer to, and how is it significant?

6.2.6 UJT trigger voltage

The *unijunction transistor*, or *UJT*, is an interesting device, exhibiting hysteresis just like SCRs and TRIACs. Its schematic symbol is as follows:

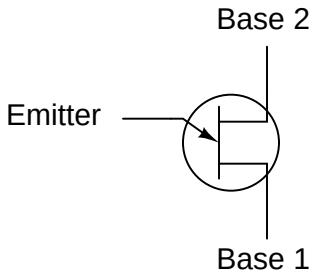

One equivalent circuit diagram for the UJT uses a pair of transistors and a pair of resistors:

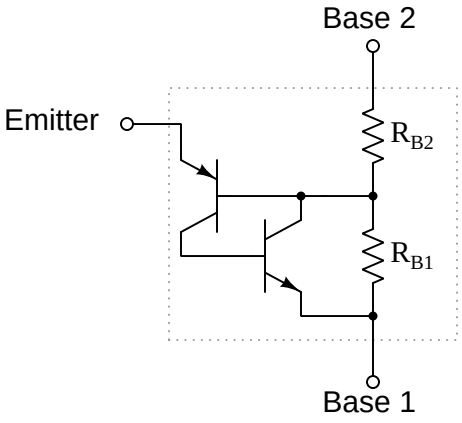

When the two base terminals of a UJT are connected across a source of DC voltage, the two base resistances (R_{B1} and R_{B2}) form a voltage divider, splitting the applied voltage into lesser portions:

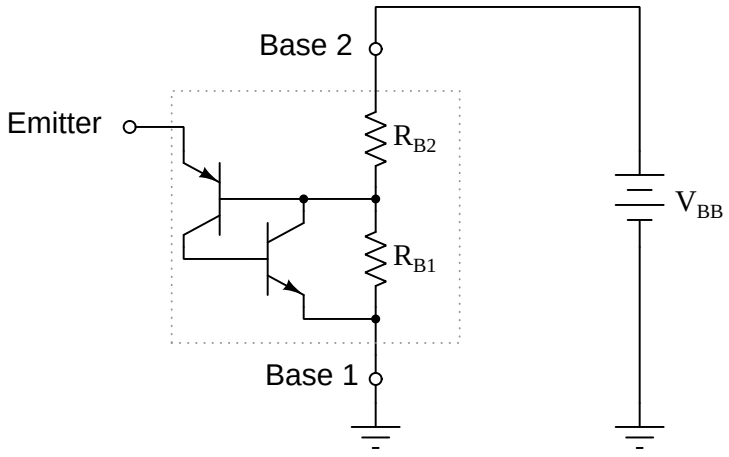

Calculate the voltage and determine the polarity which must be applied to the emitter terminal of the UJT to turn it on. Write an equation solving for the magnitude of this triggering voltage (symbolized as V_P), given R_{B1}, R_{B2}, and V_{BB}.

Challenges

- If V_{BB} increases, does V_P increase or decrease?
- How is the the *standoff ratio* defined for a UJT, and how might this equation be re-written to include it?

6.2.7 UJT standoff ratio

A unijunction transistor with an intrinsic standoff ratio (η) of 0.8 is powered by a 15 volt DC source. Calculate the emitter voltage needed to "trigger" this UJT into its conductive state.

Challenges

- If the standoff ratio were less than it is, how would this affect V_P?

6.3 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.

6.3.1 Identifying SCR terminals

Shown here is an illustration of a large "stud mount" type of SCR, where the body is threaded so as to be fastened to a metal base like a bolt threads into a nut:

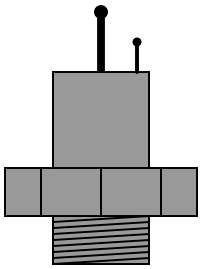

With no test instrument other than a simple continuity tester (battery and light bulb connected in series, with two test leads), how could you determine the identities of the three terminals on this SCR? Note that the threaded metal base of the SCR constitutes one of its three terminals.

Challenges

- When testing an SCR, it is best to do so under load. Describe how to set up a "loaded" test circuit and an SCR, and why load-testing is best.

6.3.2 Faults in a UJT circuit

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

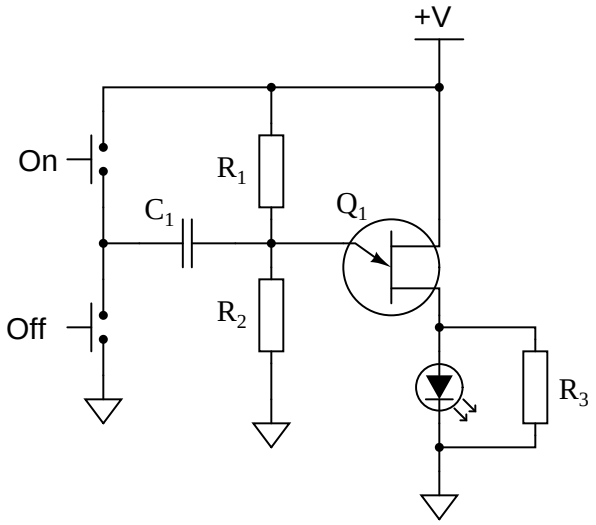

- Capacitor C_1 fails open:
- Capacitor C_1 fails shorted:
- Resistor R_1 fails open:
- Solder bridge (short) past resistor R_1:
- Resistor R_2 fails open:
- Solder bridge (short) past resistor R_2:

Challenges

- Explain how *momentary* pushbutton switches are able to make the LED *toggle* on and off.

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

Cowles, Laurence G., *Transistor Circuits and Applications*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1968.

Fenemore, Ronald William, and Quartly, Charles John *US Patent 2,784,365*, "Thyratron Inverter", application 19 November 1953, patent granted 5 March 1957.

Mazda, FF, editor, *Electronics Engineer's Reference Book*, Fifth edition, Butterworth & Company Limited, London, 1983.

Ochoa, Alfredo; Lara, Alex; Gonzalez, Gabriel; "Thyristor Theory and Design Considerations Handbook", document HBD885/D, Revision 0, ON Semiconductor, June 2005.

Pierret, Robert F., *Semiconductor Device Fundamentals*, Addison-Wesley Publishing Company, 1996.

Rose, MJ, *Electronics Engineer's Reference Book*, chapter "Thyristors and Triacs" in Mazda, 1983.

Schweizer, Anton, *GTO Databook*, section 3 "Data Sheet User's Guide", ABB Semiconductors AG.

Schweizer, Anton and Carroll, Eric, *GTO Databook*, section 5 "Application Aspects", ABB Semiconductors AG.

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.

24 July 2026 – renamed the Tutorial sections for Shockley diodes, DIACs, SCRs, GTOs, and TRIACs. Also made some minor edits to the Tutorial chapter for readability.

8 November 2025 – typograhical error corrections courtesy of Wyatt Gleason. Also made some minor edits to the Tutorial for better readability.

20 September 2025 – added a new Derivations and Technical references section exploring the phenomenon of latch-up.

13 November 2024 – minor clarifications in the Tutorial section on critical rate of voltage rise to clarify concepts for readers less familiar with oscilloscopes. Also corrected a couple of typographical errors.

11 November 2024 – added a new Case Tutorial chapter with a section on measuring rates of voltage change over time.

25 October 2024 – added a new challenging concept to the Introduction chapter.

18 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.

13 November 2023 – added some more comments on the action of positive feedback at the beginning of the Tutorial. Also removed references to "frequency" within the Tutorial chapter for those readers yet unfamiliar with AC. Also added a footnote to help clarify voltage rate-of-change.

7 November 2023 – edited image_6600 to better show rate-of-voltage rise.

21 April 2023 – added Tutorial content and Questions regarding $\frac{dV}{dt}$ rate-of-rise of voltage, and also edited Tutorial section names.

13 April 2023 – added some more detail about thyristor breakover and rate-of-rise firing.

2 December 2022 – added photograph of a large (500 Ampere rated) SCR to the Tutorial.

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.

10 November 2022 – added more questions to the Introduction chapter.

15 November 2021 – added a new section to the Tutorial to differentiate sensitive-gate versus non-sensitive-gate thyristors.

2 October 2021 – divided Tutorial into sections.

9 May 2021 – commented out or deleted empty chapters.

18 March 2021 – corrected multiple instances of "volts" that should have been capitalized "Volts".

17 December 2020 – minor edit to Tutorial.

12 November 2020 – minor edit to image_2640.eps, labeling the switch as the "on switch" rather than just "switch" because there are actually two switches in the SCR circuit. Minor edit to image_2636, adding a blocking diode to the SCR circuit (and also commenting on this in the text). Also, added to the instructor notes in some of the questions.

24 September 2020 – added more content to the Introduction chapter, and also added a footnote to the Tutorial regarding voltage rate-of-change. Instantaneous rates of change (e.g. $\frac{dV}{dt}$) are a calculus concept, and therefore present a conceptual challenge for many students lacking a background in calculus. The purpose of the footnote was to relate this concept to something immediately familiar: *speed* as a time-based function of *distance* or *position*.

6 September 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.

7 July 2020 – minor edit to wording in Tutorial.

24 January 2020 – added Foundational Concepts to the list in the Conceptual Reasoning section.

23 December 2019 – minor additions to a question.

18 December 2019 – minor edits to diagnostic questions, replacing "no multiple faults" with "no coincidental faults".

6 December 2019 – finished the Tutorial chapter, added some questions, and also added a US patent reference.

5 December 2019 – finished most of the Tutorial chapter.

4 December 2019 – began writing Tutorial.

3 December 2019 – document first created.

Index

4-layer diode, 15

Active component, 20
Adding quantities to a qualitative problem, 74
Annotating diagrams, 73

Breakdown voltage, 16
Breakover voltage, 17
Bus, 24

Checking for exceptions, 74
Checking your work, 74
Code, computer, 83
Commutation, 22
Concentration, doping, 40
Contact, ohmic, 47
Contact, rectifying, 47
Critical rate of voltage rise, 19
Critical voltage rate of rise, 18
Crowbar, 17
Current probe, 8
Current, holding, 17, 21

Dimensional analysis, 73
Diode, four-layer, 15
Diode, PNPN, 15
Diode, Schottky, 15
Diode, Shockley, 15
Displacement current, 8
Distance versus speed, 10
Dropout, 17

Edwards, Tim, 84
Electronic circuit, 20

Feedback, positive, 14
Firing, 17
Forced commutation, 22
Four-layer diode, 15

Gas, 34
Gate turn-off thyristor, 22
Graph values to solve a problem, 74
Greenleaf, Cynthia, 49
GTO, 22

Heat sink, 24
Holding current, 17, 21
How to teach with these modules, 81
Hwang, Andrew D., 85
Hysteresis, 14

Identify given data, 73
Identify relevant principles, 73
Intermediate results, 73
Inverted instruction, 81
Inverter gate, 40

Knuth, Donald, 84

Lamport, Leslie, 84
Latch, 17
Latch-up, 38
Latch-up, pseudo, 47
Light-triggered thyristor, 21
Limiting cases, 74
Line commutation, 22
Low-current dropout, 17
LTT, 21

Maxwell, James Clerk, 29
Metacognition, 54
Moolenaar, Bram, 83
Murphy, Lynn, 49

Natural commutation, 22

Ohm's Law for a capacitor, 18, 19
Ohm's Law for an inductor, 19
Ohmic contact, 47
Open-source, 83
Operating quadrant, TRIAC, 27
Overvoltage protection, 17

Parasitic effect, 8
Parasitic transistor, 38
Passive component, 20
Pitch, 8
Plasma, 34
PNPN diode, 15
Positive feedback, 14
Power dissipation, 14
Problem-solving: annotate diagrams, 73
Problem-solving: check for exceptions, 74
Problem-solving: checking work, 74
Problem-solving: dimensional analysis, 73
Problem-solving: graph values, 74
Problem-solving: identify given data, 73
Problem-solving: identify relevant principles, 73
Problem-solving: interpret intermediate results, 73
Problem-solving: limiting cases, 74
Problem-solving: qualitative to quantitative, 74
Problem-solving: quantitative to qualitative, 74
Problem-solving: reductio ad absurdum, 74
Problem-solving: simplify the system, 73
Problem-solving: thought experiment, 4, 16, 23, 73
Problem-solving: track units of measurement, 73
Problem-solving: visually represent the system, 73
Problem-solving: work in reverse, 74
Protection, overvoltage, 17
Pseudo latch-up, 47
Push-pull amplifier, 40

Quadrant, TRIAC, 27
Qualitatively approaching a quantitative problem, 74

Rate of rise of voltage, 18
Reading Apprenticeship, 49
Rectifying contact, 47
Reductio ad absurdum, 74, 80, 81
Rise over run, 8

Schoenbach, Ruth, 49
Schottky diode, 15
Scientific method, 54
SCR, 20
Sensitive-gate thyristor, 28
Shockley diode, 15
Silicon-controlled rectifier, 20
Simplifying a system, 73
Sinking, 40
Slope, 8
Socrates, 80
Socratic dialogue, 81
Sourcing, 40
Speed versus distance, 10
SPICE, 49
Stallman, Richard, 83
Standard-gate thyristor, 28
Star grounding, 45
Star-point power distribution, 45

Tangent line, 8
Thought experiment, 4, 16, 23, 73
Thyratron, 36
Thyristor, 38
Torvalds, Linus, 83
Transistor, parasitic, 38
TRIAC, 25
TRIAC operating quadrant, 27
Triggering, 17
Tube, thyratron, 36

Units of measurement, 73

Visualizing a system, 73
Voltage rate of rise, 18
Voltage rise, critical rate of, 19
Voltage, breakdown, 16
Voltage, breakover, 17
Volts per second, 18

Work in reverse to solve a problem, 74
WYSIWYG, 83, 84