CHAPTER 1

Overview

The title of this first chapter, "Foundations," describes its place pretty well: here you will learn techniques that will underlie circuitry that later produces impressive results. Chapter 1's circuits are humbler than what you will see later, and the devices you meet here are probably more familiar to you than, say, transistors, operational amplifiers—or microprocessors: Ohm's Law will surprise none of you; $I = C\, dV/dt$ probably sounds at least vaguely familiar.

But the circuit elements that this chapter treats—*passive* devices—appear over and over in later *active* circuits. So, if a student happens to tell us, 'I'm going to be away on the day you're doing Lab 2,' we tell him he will have to make up the lab somehow: that the second lab, on RC circuits, is the most important in the course. If you do not use that lab to cement your understanding of RC circuits—especially filters—then you will be haunted by muddled thinking for at least the remainder of *analog* part of the course.

Resistors will give you no trouble; diodes will seem simple enough, at least in the view that we settle for: they are one-way conductors. Capacitors and inductors behave more strangely. We will see very few circuits that use inductors, but a great many that use capacitors. You are likely to need a good deal of practice before you get comfortable with the central facts of capacitors' behavior—easy to state, hard to get an intuitive grip on: they pass AC, block DC, and *sometimes* cause large phase shifts.

We should also restate a word of reassurance offered by the Text (p. 29), but seldom believed by students: you can manage this course perfectly even if you cannot follow the mathematical arguments that begin in sec. 1.18 (use of complex quantities to represent voltage and current), and even if, after reading the spectacularly-dense *Math Review* in appendix B you feel that *you* must be spectacularly dense. This is the place in the Text and course where the squeamish usually begin to wonder if they ought to retreat to some slower-paced treatment of the subject. Do not give up at this point; hang on until you have seen transistors, at least. The mathematical arguments of 1.18 are not at all characteristic of this Text or of this course. To the contrary, one of the most striking qualities of this Text is its cheerful evasion of complexity whenever a simpler account can carry you to a good design. The treatment of transistors offers a good example, and you ought to stay with the course long enough to see that: the transistor chapter is difficult, but wonderfully simpler than most other treatments of the subject. You will begin designing useful transistor circuits on your first day with the subject.

It is also in the first three labs that you will get used to the lab instruments—and especially to the most important of these, the oscilloscope. It is a complex machine; only practice will teach you to use it well. Do not make the common mistake of thinking that the person next to you who is turning knobs so confidently, flipping switches and adjusting trigger level—all on the first day of the course—is smarter than you are. No, that person has done it before. In two weeks, you too will be making the scope do your bidding—assuming that you don't leave the work to that person next to you—who knew it all from the beginning.

The images on the scope screen make silent and invisible events visible, though strangely abstracted as well; these scope traces will become your mental images of what happens in your circuits. The scope will serve as a time microscope that will let you see events that last a handful of *nanoseconds*: the length of time light takes to get from you to the person sitting a little way down the lab bench. You may even find yourself reacting emotionally to shapes

Cambridge University Press
978-0-521-37709-6 - Student Manual for The Art of Electronics
Thomas C. Hayes and Paul Horowitz
Excerpt
More information

on the screen: feeling good when you see a smooth, handsome sine wave; disturbed when you see the peaks of the sine clipped, or its shape warped; annoyed when fuzz grows on your waveforms.

Anticipating some of these experiences, and to get you in the mood to enjoy the coming weeks in which small events will paint their self-portraits on your screen, we offer you a view of some scope traces that never quite occurred, and that nevertheless seem just about right: just what a scope *would* show if it could. This drawing has been posted on one of our doors for years, now, and students who happen by pause, peer, hesitate—evidently working a bit to put a mental frame around these not-quite-possible pictures; sometimes they ask if these *are* scope traces. They are not, of course; the leap beyond what a scope can show was the artist's: Saul Steinberg's. Graciously, he has allowed us to show his drawing here. We hope you enjoy it. Perhaps it will help you to look on your less exotic scope displays with a little of the respect and wonder with which we have to look on the traces below.

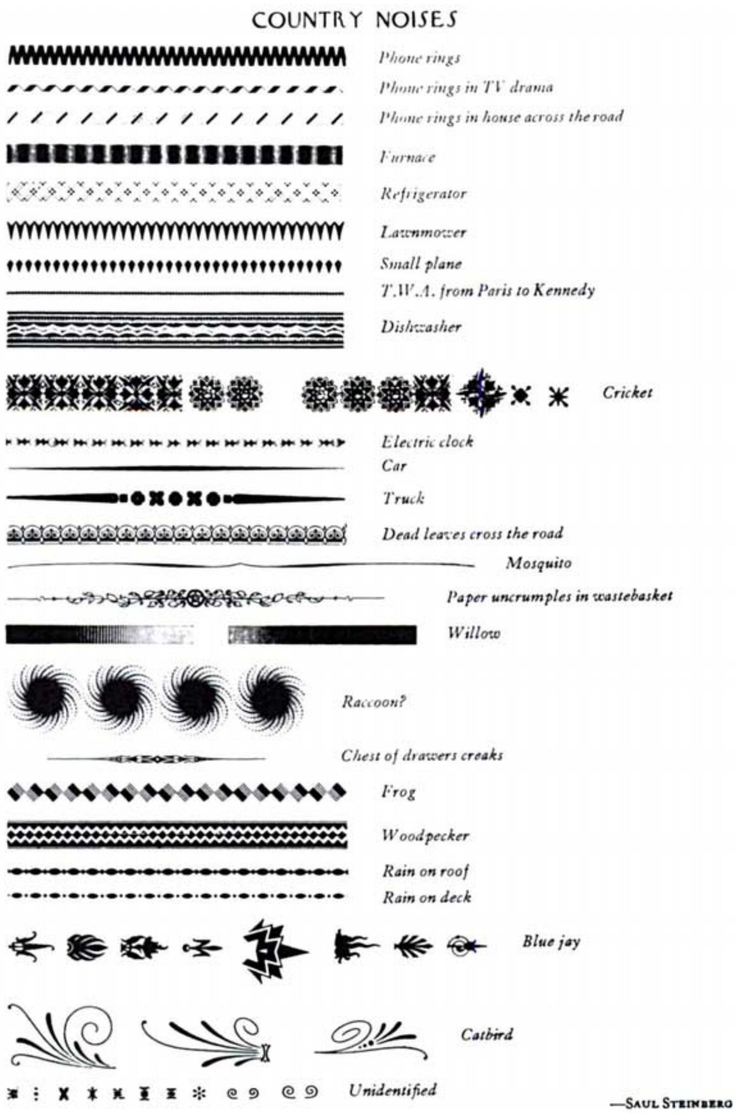

Figure IN1.1: Drawing by Saul Steinberg; copyright 1979 The New Yorker Magazine, Inc.

Class 1: DC Circuits

Topics:

- What this course treats: Art? of Electronics
  DC circuits
  Today we will look at circuits made up entirely of
  - DC voltage sources (things whose output voltage is constant over time; things like a battery, or a lab power supply);
    and
  - resistors.

Sounds simple, and it is. We will try to point out quick ways to handle these familiar circuit elements. We will concentrate on one circuit fragment, the voltage divider.

Preliminary: What is "the art of electronics?"

Not an art, perhaps, but a craft. Here's the Text's formulation of what it claims to teach:

> ...the laws, rules of thumb, and tricks that constitute the art of electronics as we see it. (P. 1)

As you may have gathered, if you have looked at the text, this course differs from an engineering electronics course in concentrating on the "rules of thumb" and the "tricks." You will learn to use rules of thumb and reliable tricks without apology. With their help you will be able to leave the calculator-bound novice engineer in the dust!

Two Laws

t sec. 1.01

First, a glance at two *laws*: Ohm's Law, and Kirchhoff's Laws (V,I).

We rely on these rules continually, in electronics. Nevertheless, we rarely will mention Kirchhoff again. We use his observations *implicitly*. We will see and use Ohm's Law a lot, in contrast (no one has gotten around to doing what's demanded by the bumper sticker one sees around MIT: *Repeal Ohm's Law!*)

Ohm's Law: $E = IR$

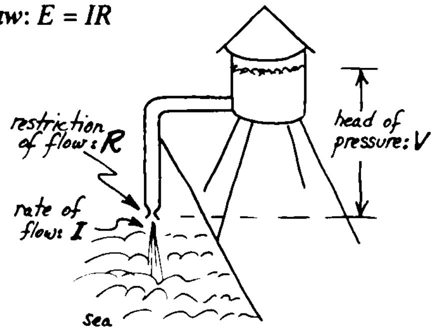

E (old-fashioned term for voltage) is analog of water pressure or 'head' of water
R describes the restriction of flow
I is the rate of flow (volume/unit time)

Figure N1.1: Hydraulic analogy: voltage as head of water, etc. Use it if it helps your intuition

The homely hydraulic analogy works pretty well, if you don't push it too far—and if you're not too proud to use such an aid to intuition.

Cambridge University Press
978-0-521-37709-6 - Student Manual for The Art of Electronics
Thomas C. Hayes and Paul Horowitz
Excerpt
More information

Ohm's is a very useful rule; but it applies only to things that behave like *resistors*. What are these? They are things that obey Ohm's Law! (Sorry folks: that's as deeply as we'll look at this question, in this course[1].)

We begin almost at once to meet devices that do *not* obey Ohm's Law (see Lab 1: a lamp; a diode). Ohm's Law describes *one* possible relation between V and I in a component; but there are others.

As the text says,

> Crudely speaking, the name of the game is to make and use gadgets that have interesting and useful I vs V characteristics. (P. 4)

Kirchhoff's Laws (V,I)

These two 'laws' probably only codify what you think you know through common sense:

Sum of voltages around loop (circuit) is zero.

Sum of currents in & out of node is zero (algebraic sum, of course).

Figure N1.2: Kirchhoff's two laws

Applications of these laws: series and parallel circuits

Series: $I_{total} = I_1 = I_2$
$V_{total} = V_1 + V_2$
(current same everywhere; voltage divides)

Parallel: $I_{total} = I_1 + I_2$
$V_{total} = V_1 = V_2$
(voltage same across all parts; current divides)

Figure N1.3: Applications of Kirchhoff's laws: Series and parallel circuits: a couple of truisms, probably familiar to you already

Query: Incidentally, where is the "loop" that Kirchhoff's law refers to?

This is *kind of boring*. So, let's hurry on to less abstract circuits: to applications—and tricks. First, some labor-saving tricks.

1. If this remark frustrates you, see an ordinary E & M book; for example, see the good discussion of the topic in E. M. Purcell, Electricity & Magnetism, cited in the Text (2d ed., 1985), or in S. Burns & P. Bond, Principles of Electronic Circuits (1987).

Cambridge University Press
978-0-521-37709-6 - Student Manual for The Art of Electronics
Thomas C. Hayes and Paul Horowitz
Excerpt
More information

Parallel Resistances: calculating equivalent R

The *conductances* add:
$\text{conductance}_{\text{total}} = \text{conductance}_1 + \text{conductance}_2 = 1/R_1 + 1/R_2$

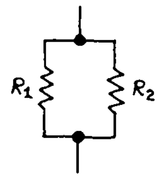

Figure N1.4: Parallel resistors: the *conductances* add; unfortunately, the *resistances* don't

This is the easy notion to remember, but not usually convenient to apply, for one rarely speaks of conductances. The notion "resistance" is so generally used that you will sometimes want to use the formula for the effective resistance of two parallel resistors:

$$R_{tot} = R_1 R_2 / (R_1 + R_2)$$

Believe it or not, even this formula is messier than what we like to ask you to work with in this course. So we proceed immediately to some tricks that let you do most work in your head.

t sec. 1.02

Consider some easy cases:

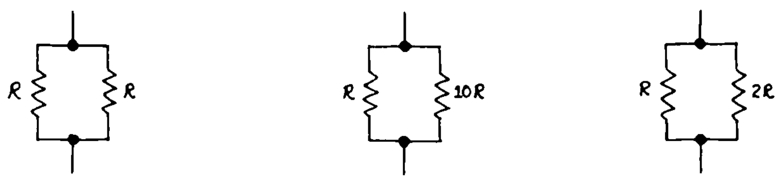

two equal R's two very unequal R's R, 2R

Figure N1.5: Parallel R's: Some easy cases

The first two cases are especially important, because they help one to *estimate* the effect of a circuit one can liken to either case. Labor-saving tricks that give you an estimate are not to be scorned: if you see an easy way to an estimate, you're likely to make the estimate. If you have to work to hard to get the answer, you may find yourself simply *not* making the estimate.

In this course we usually are content with answers good to 10%. So, if two parallel resistors differ by a factor of ten, then we can ignore the larger of the two.

Let's elevate this observation to a rule of thumb (our first). While we're at it, we can state the equivalent rule for resistors in series.

Parallel resistances: shortcuts

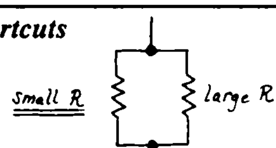

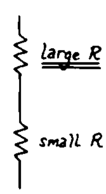

In a parallel circuit, a resistor much *smaller* than others dominates.
In a series circuit, the *large* resistor dominates.

Figure N1.6: Resistor calculation shortcut: parallel, series

Voltage Divider

Text sec. 1.03

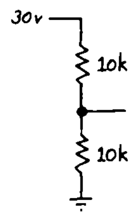

Figure N1.7: Voltage divider

At last we have reached a circuit that does something useful.

First, a *note on labeling*: we label the resistors "10k"; we omit "Ω." It goes without saying. The "k" means kilo- or 10^3, as you probably know.

One can calculate V_{out} in several ways. We will try to push you toward the way that makes it easy to get an answer in your head.

Three ways:

1. Calculate the current through the series resistance; use that to calculate the voltage in the lower leg of the divider.

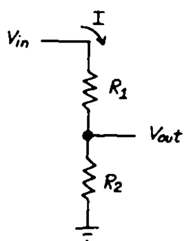

$I = V_{in} / (R_1 + R_2)$
Here, that's 30v / 20kΩ = 1.5 mA
$V_{out} = I \cdot R_2$
Here, that's 1.5 mA · 10k = 15 v

Figure N1.8: Voltage divider: first method (too hard!): calculate current explicitly

That takes too long.

2. Rely on the fact that I is constant in top and bottom, but do that implicitly. If you want an algebraic argument, you might say,

$$V_2/(V_1 + V_2) = IR_2 / (I[R_1 + R_2]) = R_2 /(R_1 + R_2)$$

or,

$$(1),\ V_{out} = V_{in}\ \{R_2 /(R_1 + R_2)\}$$

In this case, that means

$$V_{out} = V_{in}\ (10k/20k) = V_{in}/2$$

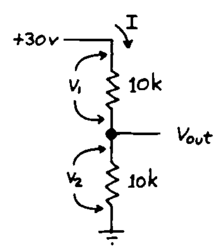

Figure N1.9: Voltage divider: second method: (a little better): current implicit

That's *much better*, and you will use formula (*1*) fairly often. But we would like to push you not to memorize that equation, but instead to—

3. Say to yourself in words how the divider works: something like,

> *Since the currents in top and bottom are equal, the voltage drops are proportional to the resistances* (later, *impedances*—a more general notion that covers devices other than resistors).

So, in this case, where the lower R makes up half the total resistance, it also will show half the total voltage.

For another example, if the lower leg is 10 times the upper leg, it will show about 90% of the input voltage (10/11, if you're fussy, but 90%, to our usual tolerances).

Loading, and "output impedance"

t sec. 1.05,

Now—after you've calculated V_{out} for the divider—suppose someone comes along and puts in a third resistor:

t exercise 1.9

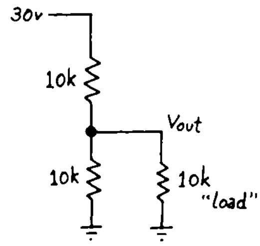

Figure N1.10: Voltage divider *loaded*

(*Query:* Are you entitled to be outraged? Is this no fair?) Again there is more than one way to make the new calculation—but one way is tidier than the other.

Two possible methods:

1. Tedious Method:

t exercise 1.19

Model the two lower R's as one R; calculate V_{out} for this new voltage divider:

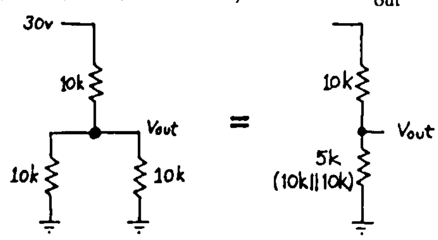

The new divider delivers 1/3 V_{in}

Figure N1.11: Voltage divider *loaded*: load and lower R combined in model

That's reasonable, but it requires you to draw a new model to describe each possible loading.

2. *Better method*: Thevenin's.

Text sec. 1.05

Thevenin Model

Thevenin's good idea:

Model the actual circuit (unloaded) with a simpler circuit—the *Thevenin model*—which is an idealized voltage source in series with a resistor. One can then see pretty readily how that simpler circuit will behave under various loads.

Figure N1.12: Thevenin Model: perfect voltage source in series with output resistance

Here's how to calculate the two elements of the Thevenin model:

V_{Thevenin}: Just $V_{\mathrm{open\ circuit}}$: the voltage out when nothing is attached ("no load")

R_{Thevenin}: *Defined* as the quotient of $V_{\mathrm{Thevenin}} / I_{\mathrm{short\text{-}circuit}}$, which is the current that flows from the circuit output to ground if you simply *short* the output to ground.

In practice, you are not likely to discover R_{Thev} by so brutal an experiment; and if you have a diagram of the circuit to look at, there is a much faster shortcut:

Shortcut calculation of $\underline{R}_{Thev}$

Given a circuit diagram, the fastest way to calculate R_{Thev} is to see it as *the parallel resistance of the several resistances viewed from the output*.

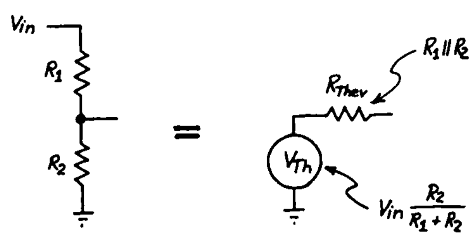

Figure N1.13: $R_{Thev} = R_1$ parallel R_2

(This formulation assumes that the voltage sources are ideal, incidentally; when they are not, we need to include their output resistance. For the moment, let's ignore this complication.)

Cambridge University Press
978-0-521-37709-6 - Student Manual for The Art of Electronics
Thomas C. Hayes and Paul Horowitz
Excerpt
More information

You saw this result above, but this still may strike you as a little odd: why should R_1, going up to the positive supply, be treated as *parallel* to R_2? Well, suppose the positive supply were set at 0 volts. Then surely the two resistances would be in parallel, right?

Or suppose a different divider (chosen to make the numbers easy): twenty volts divided by two 10k resistors. To discover the *impedance* at the output, do the usual experiment (one that we will speak of again and again):

A general definition and procedure to determine impedance at a point:
To discover the *impedance* at a point:
apply a ΔV; find ΔI.
The quotient is the impedance

This you will recognize as just a "small-signal" or "dynamic" version of Ohm's Law.

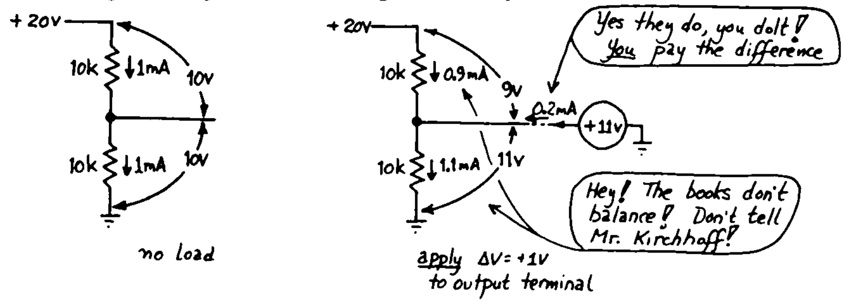

Figure N1.14: Hypothetical divider: current = 1 mA; apply a wiggle of voltage, ΔV; see what ΔI results

In this case 1 mA was flowing before the wiggle. After we force the output up by 1v, the currents in top and bottom resistors no longer match: upstairs: 0.9 mA; downstairs, 1.1 mA. The difference must come from you, the wiggler.

Result: impedance = $\Delta V / \Delta I$ = 1v / 0.2 mA = 5 k

And—happily—that is the parallel resistance of the two R's. Does that argument make the result easier to accept?

You may be wondering why this model is useful. Here is one way to put the answer, though probably you will remain skeptical until you have seen the model at work in several examples: Any non-ideal voltage source "droops" when loaded. How much it droops depends on its "output impedance". The Thevenin equivalent model, with its R_{Thevenin}, describes this property neatly in a single number.

Applying the Thevenin model

First, let's make sure Thevenin had it right: let's make sure his model behaves the way the original circuit does. We found that the 10k, 10k divider from 30 volts, which put out 15v when not loaded, drooped to 10V under a 10k load. Does the model do the same?

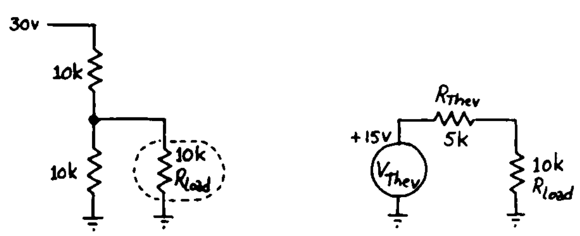

Figure N1.15: Thevenin model and load: droops as original circuit drooped

Yes, the model droops to the extent the original did: down to 10 v. What the model provides that the original circuit lacked is that single value, R_{Thev}, expressing how droopy/stiff the output is.

If someone changed the value of the load, the Thevenin model would help you to see what droop to expect; if, instead, you didn't use the model and had to put the two lower resistors in parallel again and recalculate their parallel resistance, you'd take longer to get each answer, and still you might not get a feel for the circuit's *output impedance*.

Let's try using the model on a set of voltage sources that differ *only* in R_{Thev}. At the same time we can see the effect of an instrument's *input impedance*.

Suppose we have a set of voltage dividers, dividing a 20v input by two. Let's assume that we use 1% resistors (value good to ±1%).

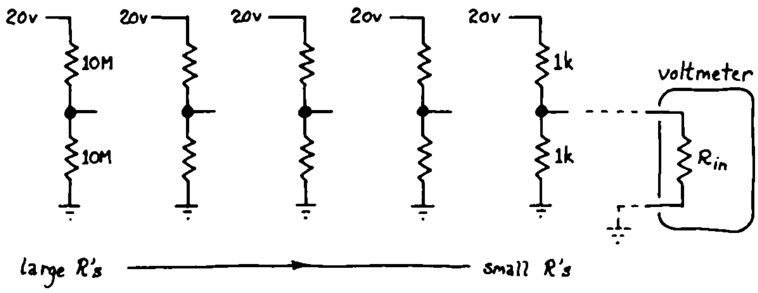

Figure N1.16: A set of similar voltage dividers: same V_{Th}, differing R_{Th}'s

V_{Thev} is obvious, and is the same in all cases. R_{Th} evidently varies from divider to divider.

Suppose now that we try to *measure* V_{out} at the output of each divider. If we measured with a *perfect* voltmeter, the answer in all cases would be 10v. (*Query*: is it 10.000v? 10.0v?)

But if we actually perform the measurement, we will encounter the R_{in} of our imperfect lab voltmeters. Let's try it with a VOM ("volt-ohm-meter," the conventional name for the old-fashioned "analog" meter, which gives its answers by deflecting its needle to a degree that forms an analog to the quantity measured), and then with a DVM ("digital voltmeter," a more recent invention, which usually can measure current and resistance as well as voltage, despite its name; both types sometimes are called simply "multimeters").

Suppose you poke the several divider outputs, beginning from the right side, where the resistors are 1kΩ. Here's a table showing what we find, at three of the dividers:

R values, divider	Measured V_{out}	Inference
1k	9.95	within R tolerance
10k	9.76	loading barely apparent
100k	8.05	loading obvious

The 8.05 v reading shows such obvious loading—and such a nice round number, if we treat it as "8 v"—that we can use this to calculate the meter's R_{in} without much effort:

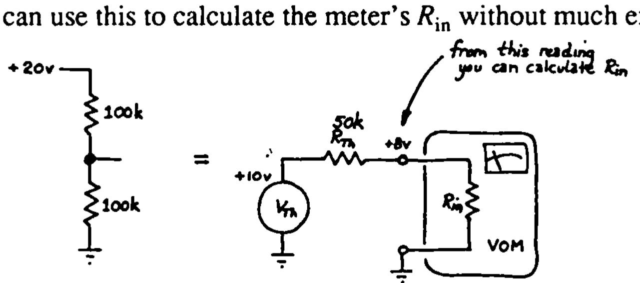

Figure N1.17: VOM reading departs from ideal; we can infer $R_{in\text{-}VOM}$