Applied Analog Electronics

a first course in electronics

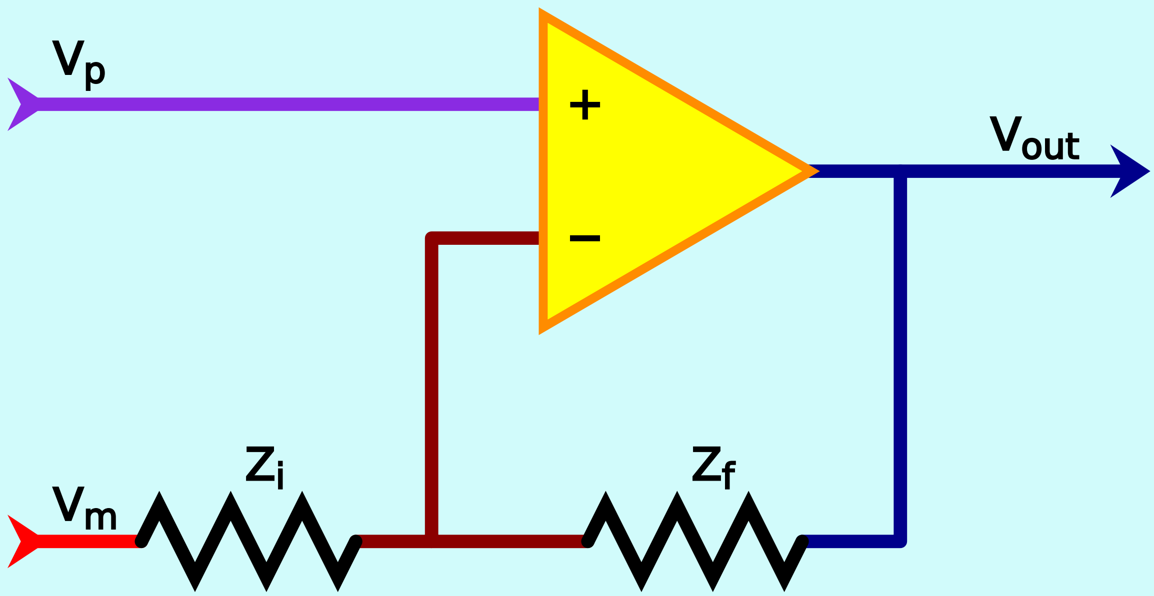

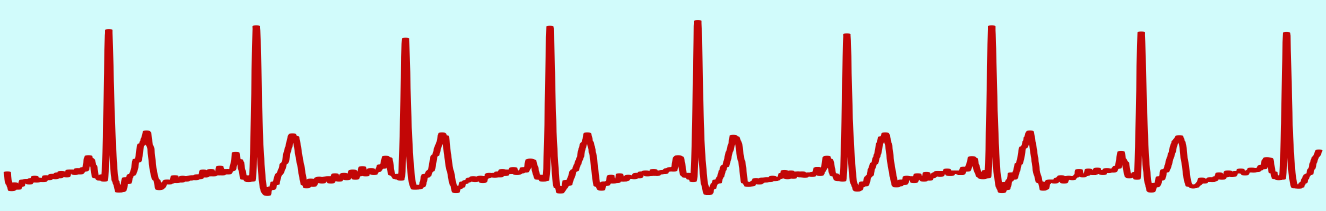

Kevin Karplus

Applied Analog Electronics:

a first course in electronics

Kevin Karplus

Edition 1.3.2—April 28, 2023

Published by the author through Leanpub

Printed copies available from World Scientific Publishing, with somewhat different pa
tion and less color.

Published through Leanpub.
Available at https://leanpub.com/applied_analog_electronics

Preface

This preface is intended primarily for people who plan to teach from this book, including autodidacts who plan to teach themselves. The introduction for students is in Chapter 1.

Why I wrote this book

I started creating the course that this book is based on in June 2012, teaching it for the first time in January 2013. There was a pressing need to have a more accessible and useful electronics course for bioengineering majors—the existing EE circuits course that they were required to take was really an applied math course, with almost no engineering design. The bioengineers (except the few who went into bioelectronics) had no idea why they were taking circuits, saw no point to all the math, and generally did poorly in the course.

Thinking about the applied circuits course started from what design labs I wanted the students to do, and only afterwards filled in what concepts were needed to allow them to do the designs. I had lots of ideas for labs, and I tried out several of them at home. Many of the labs that I initially thought would be great got rejected, because they were too hard, too easy, didn't teach anything useful, or would require working in a wet lab, rather than an electronics lab. Mixing wet-lab and electronics equipment is a bit risky, particularly for first-time lab students.

Once I'd chosen several of the labs, I started looking for a textbook. I was not able to find one that came close to teaching what I wanted in the order I wanted. Almost all wanted to do 5–10 weeks of preparation before students did any design (if they ever got to design at all), but I wanted students doing design from the first or second week.

I tried putting together an online text out of Wikipedia pages, but that turned out to be difficult—the pages either had so little information that students learned nothing from them, or so much esoterica that students couldn't extract the key ideas from them.

I ended up writing detailed lab handouts that provided background information students needed, as well as the design goals and requirements of the labs.

Unfortunately, bioengineering students have been conditioned to look at lab handouts as protocol sheets to glance at ahead of time, but only really read once they are in the lab. As a result students were coming to lab unprepared, with none of the design work done, and large amounts of lab time was being wasted on students reading and doing design work that should have been done at home.

This book is an attempt to rewrite those lab handouts, separating out the background material from the design assignment, and supplementing the background with material that had been covered in lectures, but had not worked its way into the lab handouts. It is my hope that the extra gravitas of the *textbook* format will make students more willing to do the reading and the homework ahead of time, so that lab time can be used more productively.

This book is designed for a specific course and includes some material that is normally found in the course syllabus, rather than in a textbook. As the book has matured, I've removed much of the most course-specific material, but I've chosen to retain some in the book to try to convey some of my teaching style and to make it easier for an instructor to teach the course the way I would teach it. Anyone using the book can, of course, override these sections with policies presented in their own syllabus.

Detailed notes on the development and presentation of the course can be found on my blog, where I have over 500 posts specific to the course development:
https://gasstationwithoutpumps.wordpress.com/circuits-course-table-of-contents/
It is not possible, nor desirable, to include all that information in this book, but anyone thinking of duplicating or adapting this course is likely to find some food for thought on the blog.

If instructors at other institutions or in other fields find this course design to be useful, I'll be pleased. I firmly expect that any instructor who undertakes such a course will find much that they want to do differently (the course has come out different with somewhat different labs and different order of instruction each time I've taught it), so this book should be treated more as a jumping-off place for exploring the world of applied circuits classes than as a finished product.

Who the book is for

This book and the course it is associated with are intended for anyone who wants a practical introduction to electronic circuits. We work at the op-amp level, not the transistor level, so that simple models are sufficient for most of the design work.

Although examples were chosen for bioengineers, the applications are mostly ones of interest to anyone: blood pressure, pulse, EKG, temperature, and sound. Only the electrode characterization is rather bio-specific.

The course expects students to have seen some circuits before—at the level covered in a high-school or college freshman physics class, for example. Students should have heard of Ohm's law ($V = IR$) and of capacitors ($Q = CV$), and know that current is the movement of charged particles (generally electrons in metal wires, though we do look at ionic currents in the electrode labs, and talk about holes as charge carriers in p-type semiconductors). I try to review this material in Chapter 2, so that students with a few holes in their education can fill in the gaps. Over the years of teaching the course, I've made fewer and fewer assumptions about what students retain from prior courses, so it may now be possible for students to take the course with no prior notion of circuits.

Students are expected to have been exposed to complex numbers, so that Euler's formula $e^{j\theta} = \cos(\theta) + j\sin(\theta)$ does not come as a shock to them. We also do some differentiation ($I = \frac{\mathrm{d}Q}{\mathrm{d}t}$ and derivatives of sinusoids when looking at complex impedance). They are also expected to understand and be able to manipulate logarithms and fractions. Even college engineering students seem to have difficulty understanding logarithms or adding and multiplying fractions—skills that used to be considered high-school and middle-school math.

Because of the math and physics requirements, the course is aimed primarily at sophomores in college, but it should be accessible to interested high-school students who have had calculus. Generally, students who have had single-variable differential calculus (AP Calculus AB) should have no trouble with any of the math in this book, and students might be able to get by with a good deal less math. Facility with simple algebra and fractions is essential, though.

Autodidacticism

As is the case for many of the courses I've created, I taught myself circuits. I had a little help (my Dad taught me a little when I was in high school, and I had a digital logic course and a VLSI design course as a grad student), but mostly I learned on my own from books and from experimentation.

Because so much of what I've taught is material that I've had to teach myself, I tend to take a different approach to teaching than many other faculty. I see my role as trying to provide guidance for students to learn the material faster than I did, with less time chasing down blind alleys, not to just dump some predigested knowledge into their heads for them to memorize and regurgitate. I don't teach them as I've been taught, but as how I wish I had been taught. I tend to pose them problems to guide their learning, rather than giving them information that they are expected to repeat back to me.

I've tried to write this book so that other autodidacts can teach themselves the rudiments of electronics from this book, providing pointers to other sources (mainly Wikipedia) where there is no room in the book for all the material that one should know. Citations in the book are heavy on web pages and light on traditional (book and journal sources), to make it easier for people without access to academic libraries to find the material.

Because the book is intended as a jumping-off place for exploring electronics, in places it contains more than the bare minimum needed to do the labs. In some places, students will complain that the book provides insufficient detail. My goal is *not* to provide a lab handbook that tells students precisely what they have to do, but to give them enough resources to figure things out for themselves—they will certainly have to supplement the book with data sheets, and occasionally with other outside sources.

To enable people to work through the material themselves, I've tried to design the labs in three levels:

1. using the professional-level test equipment in some of our university teaching lab (setups that cost about $10,000 a station),
2. using USB-based test equipment (requiring about $450 a station, plus a computer),
3. using home-brew equipment (requiring about $100 a station, plus a computer).

The labs with professional-level equipment have been used repeatedly in the classroom, but we switched to the USB-based approach in the 2018 offering of the course. The USB-based test equipment allowed larger lab section sizes and better data collection, because of the integration of the test equipment with the student laptops that they prepared their reports on. Unless the course is being taught in an already-equipped lab with training on that lab equipment as a major goal, I recommend using USB-based test equipment.

Don't be too helpful

This section has previously appeared in a slightly different form on my blog as Showing is better than telling, but not by much [52].

Robert Talbert, in a blog post *Examples and the light bulb* [96] wrote

> *I have a confession to make: At this point in the semester (week 11), there's a question I get that nearly drives me to despair. That question is:*
>
> **Can we see more examples in class?**
>
> *Why does this question bug me so much? It's not because examples are bad. On the contrary, the research shows (and this is surely backed up by experience) that studying worked examples can be a highly effective strategy for learning a concept. So I ought to be happy to hear it, right?*

The difficulty, of course, is that the students are asking to see examples, rather than working on the examples themselves—they are asking to be spoon-fed mush rather than chewing for themselves.

I have found in my own learning that I can get a certain amount by reading, but that really understanding material requires me to work out problems for myself. Sometimes this just means doing exercises from the textbook (a boring task which I have trouble forcing myself to do without the structure of a course), and sometimes it means struggling with making something work to solve a real problem. Real problems are both motivating and frustrating—just doing carefully drafted exercises that are designed to work out easily doesn't always help much in applying ideas to the real world.

Talbert gets the point across well:

> *Of course at the beginning of a semester, students aren't experts, and showing them examples is important. But what I also have to do is (1) teach students how to study examples and (2) set and adhere to an exit strategy for giving examples. My job is not to give more and more examples. Instead it's to say: Rather than give you more examples, let me instead give you the tools to create and verify your own examples. And then, at some point in the semester, formally withdraw from the role of chief example-giver and turn that responsibility over to the students.*

This is the same idea as in my post *Descaffolding* [48], which was prompted by a post by Grant Wiggins, *Autonomy and the need to back off by design as teachers* [104]. It also fits in with Dan Meyer's theme to "be less helpful" [70]

Given how frequently teachers and teacher leaders discuss it, I think that over-scaffolding is a common problem for many teachers. We all want to help the struggling student succeed, but too often we make them incapable of succeeding without us. If they always outsource their thinking, they'll never develop their own skills.

To use analogies from other fields: over-scaffolding is like showing the students only great literature and telling them about the writing process, but never having them struggle through 5 to 10 drafts of a piece of writing, or teaching art by showing only cast bronzes and mosaics, but never having them do a sketch or sculpt in clay. Showing or telling students how to do something is often necessary (students can't be expected to guess non-obvious methods), but it needs to be followed by students doing things for themselves.

A lot of us put a lot of time into polishing our presentations so that the students see the cleanest, most elegant way of doing a proof or solving a problem, but never see the debugging and refinement process that creates such elegant results. I've never been guilty of the over-polished lecture: I give my lectures as extemporaneous performances that are never the same twice. For one course (not the electronics course that this book is for), I did not even prepare any lectures, but had the students give me problems from the homework that they wanted to see how to do, a process I called *live-action math*. That approach required a thorough understanding of the material and a confidence that I could do any of the problems in front of an audience without prior prep.

Not all my classes are so extreme, but when I give examples I always try to make them examples of problem solving (as opposed to examples of solved problems). In the first, prototype run of the electronics course I probably did about the right number of examples in class and got the students involved in solving them, but I did not give the students enough simple problems to practice on. I was withdrawing the supports too quickly and trying to have them jump from the material in the reading (which they were not doing) directly to design problems. In subsequent offerings I have gradually increased the number of routine exercises (though I've always hated the drill work) to help them build their skills, but more are probably still needed.

So too many examples is not a big problem in my teaching style. The bigger teaching difficulty I have is keeping myself from doing debugging for the students. In labs and programming courses I can find student errors much more quickly than they can, and I have to restrain myself from just pointing out the (to me) obvious problem. I can think of several times in the electronics lab when I glanced at a breadboard that students had asked for help with and just asked them "where's the connection to ground for this component?" or "why are all these nodes shorted together?" That was not quite the right approach—it got them unstuck and left them some of the debugging still to do (that is, it was better than just moving the wires around for them), but did not help them develop the skills needed to see the problem at a glance themselves.

Some other approaches, like "Show me your schematic—I can't debug without a clear schematic of what you are trying to build," were probably more effective—walking away from the students and telling them to call me back when they had a schematic to debug from was very useful in the second run of the class. By the third week, everyone had a schematic drawn before asking for help. (The students requested that 2017 T-shirts have "Show me your schematic!" on them, because I said it so much in lab.)

It might be better for me to go through a checklist with the students—for example, having them check that each component has the right number of connections and check the breadboard against the schematic to see if the wiring is the same. I tried to do this in the second run of the class, particularly in consistently checking that the breadboard had the same circuit as the schematic, one wire at a time. A few of the students picked up that habit, but many still called me over for help before they had done consistency checks themselves. I've now incorporated many of the suggestions I give students into Chapter 13.

Occasionally I still have to step in to correct a misunderstanding (particularly at the beginning when some students don't understand how the holes of the breadboard are connected together underneath and put components in sideways), but by stepping them through a process I think I could eventually get more of them debugging on their own.

After all, the point of programming assignments and labs is to teach students how to debug, not just to get them to produce working programs or circuits. It is much harder to teach a student how to debug than to demonstrate debugging—I'm still working on better ways to do that. I think that what I've done in the applied electronics course worked for some students (they were debugging pretty independently by the end of the course), but others were still relying too much on help even at the end.

A big chunk of learning how to teach is figuring out how to withdraw the initial support without students failing. Suddenly yanking it out from under them will make many collapse, but being too slow to remove support will leave them still leaning on the crutch when they should be running on their own.

Setting up a course based on the book

To clone this course, an instructor should first go through and do all the labs him- or herself. Even people who know the theory very well and can do the designs in their sleep will learn a lot about the little problems that will plague the students. Doing the designs and making the measurements will give a better understanding of where the students might stumble, where they will need to be very efficient to finish on time, and what results to expect.

Video lectures for much (but not all) of the material in the book is available on YouTube on two playlists that correspond to the two halves of the course as it was taught in 2020–21. The playlists are at https://tinyurl.com/electronics-A (about 27 hours) and https://tinyurl.com/electronics-B (about 12 hours).

This course was originally intended to take about 220±10 hours over one 10-week quarter, but has been redesigned to take 300±15 hours over two 10-week quarters. Squeezing it into less than 200 hours is not advisable—the labs and the write-ups take time, especially if the students are expected to make mistakes and correct them.

The course is centered around the lab, not the lectures. The one-quarter version of the course scheduled 6 hours of lab a week, 3.5 hours of lecture, and 11.5 hours a week for reading, doing pre-lab exercises, and writing up design reports for 10 weeks. The revised two-quarter version schedule 190 minutes (3:10) of lab a week, 195 minutes (3:15) of lecture, and 515 minutes (8:35) of reading, homework, and write-ups for two 10-week quarters.

The primary instruction occurs in the lab, with the lectures as support sessions to learn tools and concepts needed for doing the labs. Trying to do this course as a huge lecture with lots of barely trained TAs running the labs is almost a guarantee of an unsuccessful course.

The lab equipment needed for the course is described in the next section and the parts needed in Section 3.1.

In addition to standard, off-the-shelf parts, I custom made several parts:

stainless-steel-electrode pairs: I cut up some 1/8” 316L stainless-steel welding rods (the type for inert-gas welding, not coated for arc welding) into 13 cm pieces and ground the ends a bit to eliminate any sharp burrs. I also cut up an old plastic cutting board into approximately 35 mm squares, and drilled two 1/8” holes in them 2 cm apart. I hammered the welding rod into the holes with about 3 cm sticking out on one side. The short end can be immersed in salt water, and the long end used to attach a wire with an alligator clip (see Figure 39.2).

holders for silver electrodes: I used a laser cutter to cut several holders for 24-gauge silver wire out of acrylic (see Figure 39.1a). The SVG and DXF files for the holder design can be found at https://users.soe.ucsc.edu/~karplus/bme51/pc-boards/electrode-holder/ (the DXF files are more likely to be useful).

pressure sensors: Because the pressure sensors were too expensive at about $12 each to require each student to buy one, I bought one per station for the lab and soldered them to custom breakout boards for easier handling. The design can be found at https://users.soe.ucsc.edu/~karplus/bme51/pc-boards/MPX2050DP-breakout-v1_3/

hysteresis-oscillator boards: To give students practice with soldering on a board with few parts, I designed a custom PC board (2.5 cm by 5 cm) that they could use as soldering practice. The design can be found at https://users.soe.ucsc.edu/~karplus/bme51/pc-boards/hysteresis-oscillator-rev0.7/. Because the Teensy boards come without headers, the hysteresis-oscillator board is now their second soldering project, but it is still good practice at through-hole soldering.

capacitance touch sensors: Students make their own capacitance sensors out of aluminum foil and packing tape in Lab 4.

op-amp prototyping boards: I designed 5 cm-by-5 cm prototyping boards with places for an INA126P instrumentation amp and an MCP6004 op-amp chip, as well as several resistors, 2 4-pin screw terminals, and a potentiometer. I replaced these with a revised design that eliminated the INA126P chip and provided more room for resistors and a less confusing wiring grid (see Figures 31.2–31.4). The new design can be found at https://users.soe.ucsc.edu/~karplus/bme51/pc-boards/op-amp-proto-rev0/.

unknown impedance boards: My son, Abraham Karplus, designed an "impedance token" board for me to make simple RC circuits that students can use for practice at impedance spectroscopy and fitting impedance models. The board designs are available at https://users.soe.ucsc.edu/~karplus/bme51/pc-boards/ImpedanceToken/. The directory there also has a Python program for selecting resistor and capacitor values from a set of values, to make each impedance token different.

I've considered ordering a large number of the PC boards used in the course and reselling them, to get lower costs per board through volume production, but even buying the boards in 10s of units, the cost is only about $2.50 per student. Eagle design files and Gerber files for all the PC boards can be found at https://users.soe.ucsc.edu/~karplus/bme51/pc-boards/

This course is also a writing course, with ten substantial design reports required, each going through two drafts. Detailed feedback on the design reports is essential for students to learn engineering writing. Insistence that serious errors be corrected before work is accepted is also important, even though this usually means more time spent on reading the student work that is most painful to read. The instructor should budget at least 30 minutes per student per week for providing feedback to the students on their writing.

I generally have design reports due shortly after labs are over, with a grading schedule that returns them well before the next labs are due. The short writing time between the end of lab and the due date for the report encourages students to write stuff up before lab, rather than leaving the documentation to the end. Writing as you go is an extremely important habit to develop in engineers, but difficult to instill. Furthermore, this course is a design course, so most of the thinking and writing should be happening *before* students enter the lab, not after the lab is over.

With ten design reports in the course, the grading schedule for a one-quarter course was insane, with reports due the Friday morning after the Thursday lab session and graded reports returned on Monday. The same ten reports spread over two quarters is a much more

equipment	which labs												
laptop	1	2	3	4	5	6	7	8	9	10	11	12	13
soldering station	1			4		6			9	10			13
fume extractor	1			4		6			9	10			13
board holder(optional)				4					9	10			13
microcontroller (for PteroDAQ)	1	2	3	4	5	6	7						13
function generator			3			6	7	8		10	11	12	
ohmmeter		2	3	4	5	6	7	8	9		11	12	13
voltmeter		2					7	8				12	
oscilloscope			3	4	5	6	7		9	10	11		13
bench power supply							7			10	11	12	
thermometer		2										12	
hot water & ice		2											
beakers		2										12	
secondary containment		2										12	
micrometer				4								12	
calipers												12	
drill press (or drill)					5	6							
pressure sensor					5								
blood pressure cuff					5								
electrodes/holders												12	
NaCl solutions												12	
EKG electrodes												12	13

Table P.1: The equipment needed is listed here, along with which labs it is needed for. The numbers are the numbers of the labs that use that piece of equipment. The voltmeter, oscilloscope, and power supply can all be replaced by an Analog Discovery 2, but an ohmmeter is needed for Lab 2. For other labs, an ohmmeter is useful for checking that resistor values are correct, as students often misread the color codes.

reasonable schedule, with a report due about every two weeks and a week to do the grading. Students can be given a weekend to finish a report, with still a weekend for grading before the next report is due. Even this schedule became insane with 80 students in the course, as 40 reports took over 40 hours to grade.

What bench equipment is needed for the course

The book, as currently written, assumes that every student has a microcontroller and computer with the PteroDAQ software installed.

The instructions for the labs assume that students will be using an Analog Discovery 2, which combines an oscilloscope, function generator, and power supply in one compact package. Using separate function generators, bench power supplies, and oscilloscopes is certainly possible, but would increase the time needed for data collection and recording, especially in Lab 7, Lab 8, Lab 10, and Lab 12.

Table P.1 lists the equipment needed for each lab for the versions of the labs using professional equipment. The rest of this section will describe both what equipment is available in the lab that we use for the course, and what minimal features are needed for setting up a lab elsewhere, as most of the equipment in the lab is overkill for what we need.

soldering equipment: We need a soldering station, preferably one with a small tip, like the Weller ETU 0.015" or ETV 0.024" sloped tip and a solder vacuum (also called a *solder sucker*) for unsoldering. Some engineers prefer to use solder wick to a solder sucker, but I've had less success with solder wick. Useful, but not essential, is a board holder for holding PC boards while soldering—I like the Pana-vise Jr. PV-201 holders, but much cheaper ones are usable. Soldering equipment is needed for all levels of the labs, but can be obtained for about \$25–\$30 a station, though a temperature-controlled station like the Weller WESD51 Digital Soldering Station is a good investment at about \$130. We've had pretty good luck with the Aoyue 9378 soldering stations at \$90 a station. For at-home labs, we have used extremely cheap soldering irons (\$5 from Harbor Freight), as shipping the soldering stations was too expensive.

host computer for PteroDAQ and Analog Discovery 2: Students should have their own laptops in lab—PteroDAQ works with Windows, linux, and Mac OS X operating systems. Getting PteroDAQ to work on Windows 8 or 10 systems is possible, but difficult, because the serial drivers are unsigned, and so require rather awkward workarounds to circumvent Windows 8's insistence on signed drivers. The installation instructions on the PteroDAQ site provide the necessary instructions (see also Section 3.5).

function generator: The circuits lab that we used earlier had Agilent 33120a arbitrary waveform function generators, which is definitely overkill for the needs of the course. A low-frequency sine-wave generator is sufficient for Lab 3 and a medium-frequency (around 50 kHz–100 kHz) triangle wave generator suffices for Lab 11. The impedance measurements lab benefit from having an oscillator that goes up to 1 MHz, but that isn't strictly necessary.

The function generators in the Analog Discovery 2 are better for this course than a standalone function generator, as they can be coupled to the oscilloscope using the network-analyzer and impedance-meter functions to automate measurements in Lab 8, Lab 11, and Lab 12.

The function generator built into the Analog Discovery 2 USB oscilloscope provides much cleaner signals than most of the low-cost function generators I've used, it has good offset voltage control, and it has powerful modulation and sweep capabilities. The Analog Discovery 2 is certainly more function generator than needed for the course, and the academic price makes it very cheap.

A low-cost function generator kit (there are several on the market) probably suffices for this course. I've tested JYE Tech's FG085 function generator, and found it usable, but the waveforms are not high quality and the buttons on the device are unreliable.

I've also tried using the function generator built into the BitScope BS10 USB oscilloscope, but it does not produce outputs centered at 0 V—its outputs are between 0 V and 3.3 V, which is problematic for measuring the impedance of electrodes in Lab 12. Adding external circuitry to recenter the voltage is doable, but a nuisance for most uses of a function generator.

oscilloscope: In the instructional lab used for the applied electronics course at UCSC, we had Tektronix TDS3054 4-channel and TDS3052 2-channel digital scopes, and Kikusui C0S5041 analog scopes. All of these have far more bandwidth and functionality than are needed for this course. Initially, I favored teaching with the analog scopes, which are easier to set up, but I now prefer to have students use digital scopes, because many of the waveforms we look at are based on heartbeats, and so are rather too low frequency for an analog scope. The digital scopes are also capable of measuring (rather crudely) a number of properties of the waveforms. The university has since replaced the instructional scopes with Keysight EDUX1002A, which are much cheaper and easier to use than the old Tektronix and Kikusui scopes.

Standalone oscilloscopes are a very traditional engineering tool, and students probably should have some exposure to them, but they are being replaced more and more with USB-connected and wireless instruments that use a laptop or phone for the user interface. We have switched to using the Analog Discovery 2 USB oscilloscope, which combines a dual power supply, a voltmeter, a pair of function generators, and a 2-channel digital oscilloscope in a single low-cost unit. This allows us to equip a lab for about $300 a bench ($450 a bench including a soldering station), rather than over $10,000 a bench. Overall, the Analog Discovery 2 looks like the best investment that a student could make in a "bench" instrument.

For a hobbyist, a USB oscilloscope with a 10–20 Msample/s sampling rate may be enough. I have a BitScope BS10, because it was one of the first USB oscilloscopes to provide software that runs on Mac OS X, Linux, and Windows. It is a more featureful instrument than their newer BitScope Micro, which should still be adequate for most of the labs in the book.

Digilent also makes a low-cost open-source instrument (the OpenScope MZ), which is not as good as the Analog Discovery 2, but which should be adequate if budgets are very tight. It looks like a better deal than the BitScope instruments, but not nearly as useful as the Analog Discovery 2.

For a student who cannot afford a $100–$400 USB oscilloscope, the PteroDAQ data-acquisition system with a $12 Teensy LC board should be adequate for most of the labs, though a function generator will also be needed. The sampling rate is limited to about 20,000 samples per second, depending on the speed of the host computer, which is adequate for looking at EKGs and other heart-rate-based signals, but is a bit slow for looking at audio signals for the audio-amplifier labs. The analog-to-digital converters on the microcontroller boards also have a rather limited voltage range (0 V–3.3 V on the Teensy boards), which can make debugging signals outside that range difficult. The rather limited sampling rate limits the frequency at which impedance measurements can be made—this limitation may make Lab 8 and Lab 12 much less doable.

multimeter: In the instructional lab each bench used to have two Agilent 34401A multimeters, which was very convenient, but we could get by with just one. The multimeters got used for two functions: measuring resistor values and measuring AC voltage for the impedance characterization in Lab 8 and Lab 12. For the impedance characterization it was useful to have a multimeter like this one that can measure RMS voltage from 3 Hz to 300 kHz (and we can push the 34401A meters to 1 MHz if we sacrifice a little accuracy), but the network analyzer function of the Analog Discovery 2 is better for this purpose, and the impedance analyzer function better still.

When measuring currents in this course, we always provide our own current-sensing resistor, and measure the voltage across it, rather than using an ammeter. That's a safer way to measure current anyway, as under-supervised students in EE courses are always blowing the fuses in the ammeters, rendering them useless for other lab courses. (The joys of shared teaching labs!) The voltmeter function of the Analog Discovery 2 can replace the bench multimeters for both voltage and current measurement.

The only function we used of the bench multimeters that the Analog Discovery 2 lacks is an ohmmeter for measuring resistance. For this purpose, a handheld digital multimeter suffices, though students put too much faith in the numbers reported by the cheap multimeters, which are often highly inaccurate.

I had an old hand-held Fluke 8060A meter that worked from 20 Hz to 100 kHz for AC voltage measurement, which is sufficient, especially as it can be pushed to 5 Hz–300 kHz with some loss of accuracy. More recent Fluke meters seem to have sacrificed bandwidth for price—for the Fluke 175, 177, and 179 meters, the AC response is only 45 Hz to 1 kHz, which is too narrow to be useful for the impedance labs. The Fluke 287 and 289 meters still have 100 kHz bandwidth, but these cost $440 or more, and the Agilent (now Keysight) U1252B has 100 kHz bandwidth, but costs over $500 (the slightly cheaper U1251B has only 30 kHz bandwidth for AC voltage measurement).

Cheap hand-held multimeters like those sold at hardware stores often don't measure reliably above 2 kHz, which is too low for the impedance labs (some only go to 400 Hz!). The cheap meters often do not have a specification for the AC bandwidth, presuming that the user will only be measuring line voltage at 60 Hz (or 50 Hz, depending on what part of the world you live in). They are generally fine for measuring DC voltage and DC resistance, though, at least with the low precision needed for this course.

The $10 aiyun DT-9205A multimeter can be used, as it can measure voltages and get the ratios right up to about 40 kHz (though for accurate single-voltage measurements, the range is limited to about 1 kHz). The voltages we are interested in measuring are often fairly small, so having a low-voltage range on the meter is essential. A cheap digital multimeter like this one, combined with a good USB oscilloscope like Digilent's Analog Discovery 2, is probably the best investment for a hobbyist.

A cheap multimeter and a good USB oscilloscope is probably a better investment than a good multimeter and low-quality oscilloscope.

In a pinch, the PteroDAQ data-acquisition system can be used for measuring small DC voltages fairly reliably (between 0 V and 3.3 V), and the range can be extended by adding a voltage divider, so a DC multimeter is a low-cost convenience rather than an absolute necessity for the lab.

thermometer: We used to use a $2 28-cm-long glass thermometer with markings every 1°C. These worked fairly well, but were easily broken when students carried them in their backpacks. We switched from including them in the parts kit to having them as lab equipment, then in 2014 we were loaned some digital thermometers used in another lab class. These thermometers turned out to be a very poor choice, despite the easy-to-read 0.1°C resolution—several of the digital thermometers read 1°C to 2°C off, making the students' calibrations of the thermistors that far off from the specs. Accuracy is more important than precision for Lab 2, and that is the only lab we use the thermometers for.

There are low-cost digital thermometers available that can be easily calibrated with ice water. We are now using "CDN DTQ450X Digital ProAccurate Instant-Read Thermometers" which seem to be accurate enough for this course and cost only about $13 each.

beakers, secondary containment tub, hot water, and ice bucket:

For beakers, I bought ceramic coffee cups for about $0.50 each at the thrift store (as opposed to 150 mL beakers at around $2.50 each). One could use disposable coffee cups for the thermistor lab, Lab 2, but I prefer having as little waste produced as possible. The ceramic cups also have a fairly high thermal mass, which provided opportunity for reminding students of that concept, as they could not get the highest and lowest temperatures unless they preheated or prechilled the cups before adding the hot or cold water.

Clear plastic cups are more useful for the electrode lab, Lab 12, so that the electrolyte solution can be added to a calibrated depth on the electrodes.

Secondary containment tubs to prevent spills from spreading to surrounding electronics equipment were provided by the University, but dish-washing tubs would work as well.

For Lab 2, hot water was provided by a coffee urn, which only goes up to about 70°C–80°C and by a tea kettle that provided 100°C water. Ice water was provided from a large Thermos (another thrift store purchase), but the ice needed to be replenished from an ice machine every couple of hours.

The hobbyist or student at home probably has better access to hot and cold water sources than the electronics lab, so the thermistor lab may be easier to do at home than in the professionally equipped labs.

bench power supply: Each station in our lab had an Agilent E3631A bench power supply. These power supplies provide three independently adjustable power sources (one up to 6 V and 5 A, one up to 25 V and 1 A, and one up to −25 V and 1 A).

The settable current limitation on the power supplies is a very useful feature, as it prevents blowing fuses on ammeters and can reduce the chance of damaging chips if things are miswired, but settable voltage is sufficient.

Formerly the power-amp lab (Lab 11) used all three supplies, but that lab has been redesigned to work with a single power supply. If the power amp is powered directly from a USB cable, the current is usually limited to about 500 mA, but a 5 V or 6 V AC/DC adapter with a barrel plug could be used for the power-amp lab to get more power. The Meanwell brand sold by Jameco and Mouser seem pretty reliable.

The power-amp lab now recommends using the power supply of the Analog Discovery 2, and discusses the limitations of that power supply.

electrodes and electrode holders: There are two sets of electrodes to be tested in Lab 12: a pair made from stainless steel (polarizable) and a pair made from Ag/AgCl using fine silver wire (nonpolarizable).

The stainless-steel ones are made from 316L stainless-steel 1/8" welding rods (for inert gas welding, so no coating). I cut two pieces about 5" (12.5 cm) long with a pair of big bolt cutters, then ground the ends to round them. I then drilled two 1/8" holes about 2 cm apart in a scrap of plastic from an old cutting board (probably made of HDPE), and hammered the rods through until they stuck out about 1" (2.5 cm) on one side. The short end gets immersed in the solution being measured, and the long end provides a place to clip alligator clips on. (See Figure 39.2.)

The silver/silver-chloride electrodes are made by wrapping 24-gauge fine silver wire around an electrode holder that provides markings for immersing the electrodes to a measured depth (Figure 39.1a). The holders I made were cut from clear acrylic with a laser cutter—the design files are at https://users.soe.ucsc.edu/~karplus/bme51/pc-boards/electrode-holder/ as files holder5.svg and holder5.dxf.

bottles of salt water: I cleaned out some plastic mineral water bottles and had a colleague make up stock solutions of 1 M, 0.1 M, 20 mM, and 5 mM NaCl. I needed about 100 mL per student of each solution, with the students working in pairs. Individual measurements take much less solution, but the students needed to measure two different sets of electrodes, and many needed to remeasure at least one of the sets.

For Lab 12 the bottles are kept in a secondary containment tub well away from electronics equipment and students bring their cups (in secondary containment tubs) to the bottles to pour what they need.

If you have to make the stock solutions yourself, a centigram scale for measuring the NaCl and a graduated cylinder for measuring distilled or deionized water would be needed.

drill press: I brought in a drill press from home, so that students could drill 2 mm holes in PVC elbows for the breath-pressure apparatus (Lab 5) and 3 mm holes in LEGO® bricks for holding optoelectronics (Lab 6).

micrometer: We bought one micrometer for every 16 students in the lab, for measuring the thickness of the insulator in Lab 4 and the wire diameter in Lab 12.

calipers: We bought one set of stainless-steel vernier calipers for every 16 students in the lab, for measuring the dimensions of the electrodes.

Time expectations for the course

Hours needed

This book is now designed around two 5-unit courses—that is, the total time expected of students is about 300–330 hours. For two 10–11-week classes, that's 15 hours a week: about 3.2 in lecture, 3.2 in lab, 2 reading in preparation for class, 2 doing homework exercises, and 4.6 doing design work and write-ups (either pre-lab or post-lab).

If time constraints are tight for a course, it is better to reduce the lecture/discussion hours than the lab hours. Most of the learning happens in design work, lab time, and report writing. The reading time and class time are to enable the learning in the remaining time.

The time spent reading will not be entirely for this text book, but also reading supplemental material (particularly from Wikipedia).

The design reports are due shortly after the lab is completed, so the writing of the report should not be left until the lab is done, but should be written as much as possible while doing the design and analysis work. I am now requiring a complete draft of the report (excluding measurements and conclusions) *before* labs. At the very least, schematics should be drawn neatly before coming to lab. That way the short time after the lab is used only for making corrections and incorporating final results, not for trying to describe the entire process. The time spent each week for writing should be spread out, not concentrated into a single evening.

Some students, pressed for time, try to skimp on the reading and the pre-lab design work. This usually results in their wasting much of the lab time trying to do pencil-and-paper or calculator work, rather than building, testing, and debugging their designs. The overall result is that they run out of lab time, rush through the lab work keeping poor lab notes, and write up poor design reports. They often have to redo the design reports, costing them more time than if they had invested up-front in reading the material and doing the designs before lab.

This is a *design* class, not a science demo-lab course. The lab is where the designs are built and tested, not where the design is first thought about. Most of the thinking and writing in this class should happen *before* each lab, not after.

Possible time allocation for labs

The course that this book was written for was originally designed around a 10-week schedule, with 3 days a week of lecture/discussion (3.5 hours/week) and 2 days a week of lab (6 hours/week). That schedule turned out to be rather intense, both for the students and for the instructor, so the course was redesigned to take two 10-week quarters, with 3:15 hours a week of lecture and 3:10 hours a week of lab. For ease of scheduling, the lectures and labs alternate, with MWF lectures and TTh labs. Longer lab times may be more effective in general, but it is useful in several of the labs to have a lecture between parts of the lab, to help students learn how to analyze the data they have collected while they still have time to collect more data. Live demos/tutorials on using gnuplot in the lectures between lab sessions have been particularly useful.

In any schedule, the time spent on reading, homework, and design exercises is heavily loaded towards the beginning of the course. I have experimented with having students do the homework before lectures on the material (to encourage learning from reading and to encourage asking questions rather than passive absorption), but many students have had great difficulty with this approach.

Every time I teach the course, I end up with a slightly different schedule for the labs and lectures, based on what the class gets quickly and what they need more time on. Table P.2 gives a representative assignment of lab times, assuming that there are 20 lab sessions of 95 minutes each in each quarter. Schedules of labs and lectures for each run of the course can be found at https://users.soe.ucsc.edu/~karplus/bme51/

Lab	sessions	time
1	2*	3:10
2	4	6:20
3	2	3:10
4	4	6:20
5	4	6:20
6	4*	6:20
7	2	3:10
8	2	3:10
9	5*	7:55
10	2	3:10
11	3*	4:45
12	3	4:45
13	3*	4:45
total	40	63:20

Table P.2: Time allocation used in 2019. There was no lab report for Lab 1, and there were combined reports for Lab 7 and Lab 8 and for Lab 10 and Lab 11. The labs with asterisks required an extra weekend lab session for some students to complete the work. It is not clear that scheduling extra lab sessions for the labs would reduce this need—some students don't ask for help until too late to finish the work, no matter how much time they have.

I have not attempted to make a 15-week semester schedule for the course. The lab time per week would need to be increased to at least four hours to fit in 60 lab hours. Alternating days of lecture and lab would still be valuable.

Contents

List of Figures

List of Tables

1: Why an electronics class?

1.1 First (and sometimes last) course on electronics

This course is intended to be a first course in electronics for students who have not had anything more than a high-school physics course that covered circuits. Students should have heard of resistors and capacitors and know something about how electrons flow in wires to make circuits. Important concepts (like voltage, current, and resistance) will be reviewed, but more time will be spent on concepts that are likely to be new to students.

> The course is structured around three big concepts that are used repeatedly in different ways: *voltage dividers*, introduced in Section 5.1; *complex impedance*, introduced in Section 10.2, and *negative-feedback amplifiers* using op amps, introduced in Section 19.2. We'll use these three ideas over and over again to design a variety of different circuits.

I'll try to keep the theory and the math to a minimum, introducing just enough to make the designs in the labs possible. Students wishing to go on to more advanced electronics courses will need to follow this course with a conventional electrical engineering (EE) circuits course, which provides the math and theory in abundance.

EE circuits courses are usually taught as applied math courses, preparatory to later using the math to do design. That can work well if you later take design courses that use it, but is pretty useless if you stop with just the math and never do design. If you only take one electronics course, it should be one that does a lot of design, not one that prepares you for something you then don't do.

> This course turns around the conventional EE pedagogy, emphasizing the design elements first. If you go on in electronics, the math in the EE circuits course will make a lot more sense after this course, as you'll know what the math is useful for. But even if you don't go further in electronics professionally, this course will teach you how to design and build some simple amplifier circuits that can be useful in a lab and allow you to explore electronics at a hobbyist level.

1.2 Why teach electronics to non-EE majors?

This book was originally created for a course for bioengineers, not electrical engineers. The justification for an electronics course for bioengineers can be summed up in one word: *sensors*.

A sensor is a device that converts some physical or chemical property of interest into a more easily measured or recorded property, generally an electrical parameter: a voltage, a current, a resistance, a capacitance, an inductance, and so forth. That electrical property can be amplified, filtered, and manipulated by electronic circuits, after which it is usually converted to a numeric value that can be recorded on a computer or in a lab notebook.

This class focuses on circuits needed to connect common sensors to computers, where the information can be processed, recorded, or acted on. Because this is a first course in electronics, not a computer engineering course, we will look only at sensors that produce one-dimensional analog outputs, not more sophisticated sensors like digital cameras nor digital interfaces like I^2C and SPI. That is not to say that such topics are unimportant, just that they are beyond the scope of this course.

This course emphasizes *analog* electronics that converts the signal from whatever the sensor produces to a voltage that can be read by a low-cost computer.

We'll cover several different sensors: thermistors for measuring temperature, microphones for sound, electrodes for converting ionic current to electronic current (including EKG measurements), pressure sensors for breath pressure or blood pressure, and phototransistors for light measurements (for optical pulse monitoring).

There are, of course, other applications of electronics, and we'll look at one of them: audio amplification of sound waves. In addition to sound being a useful signal type for bioengineering work, it is pedagogically convenient, as students are already familiar with sound and electronic devices for dealing with sound. Furthermore, the pulse-width modulation and power amplifier design (Lab 11) are similar to motor-control applications in powered prostheses, wheelchairs, and mobility devices, though we'll use a loudspeaker as our output device, rather than a motor, and only use a few watts of power.

At UCSC, many students in the biomolecular concentration work in the nanopore and nanopipette labs, where electronic sensing of ionic current through small holes is a primary lab technique.

1.3 Teaching design

One reason for teaching electronics is that it provides a medium for teaching *engineering design*, which is the process of converting a specification or goal into a detailed design that achieves that goal.

To do engineering design, we need to have mental models of the real-world phenomena that we wish to measure and of the electronic components we can use for these measurements.

Models in electronics are chosen to represent just enough of the real world to allow us to do design. If we choose too simple a model, the design decisions we make will be incorrect, due to missing some important effect. If we choose too complex a model, we'll have a hard time making any design decisions, because we won't be able to do optimization or predict the effect of changes easily.

The same physical device may be modeled with many different models for different design purposes. For example, we'll look at several different models of a loudspeaker in Lab 8, ranging from a simple $8\,\Omega$ resistor to a complicated model that captures both the mechanical resonance of the loudspeaker and its electrical behavior at high frequencies. None of these models is *correct* in any absolute sense—they are just more or less useful for various design tasks.

Of course, when we work with models, we never know for sure whether the model is really capturing all that it needs to—so we must check our thinking by querying the real world and not just the model. "Try it and see" is my standard answer to any question of the form "Is this right?" or "Will this work?" (unless there is a safety concern that needs to be addressed). I reply that way so often that it becomes a mantra for the lab—we've even put "Try it and see!" on the class T-shirts.

In any given design problem we may need to use several different models of the same phenomenon. These are often at different levels of abstraction—for example, in most circuits classes, a wire is treated as a *node*, that is as if all parts of the wire have the identical voltage at all times. For many purposes, this is an excellent level of abstraction, allowing us to ignore many irrelevant properties of the wire. But when we are passing a large current through the wire, we may need to worry about the resistance of the wire, which produces a voltage drop $V = IR$ and dissipates power in the wire that may heat it up. (I have some wires whose insulation melted together when I tried using them to supply current to a motor—the wires had a much higher resistance than I expected.)

Even the model of a wire as a resistor is not always enough—if we are dealing with high frequency signals, then the inductance of the wire may become important. In still other applications, we may need to model the wire as a distributed transmission line (which is beyond the scope of this course).

Picking the right level of abstraction—the right model to use—depends partly on experience, but also on checking whether the model chosen works well enough.

> Besides checking that models are reasonable, I also want students to do what engineers refer to as *sanity checks*—checks that part values, amplifier gain, voltages, currents, or other numbers in their design "make sense".

Often these sanity checks are just arithmetic checks: for example, if I have an amplifier with a gain of 1000 and an input signal of 200 mV, I expect an output signal of 200 V, but my power supply can't deliver that, so something is wrong. Is my input signal only 200 μV not 200 mV? Is the amplifier gain only supposed to be 10? Is the circuit one that is supposed to saturate the amplifier at one of the power rails, like a comparator, so that the model of a linear amplification is the wrong model?

Sometimes the sanity check is just a polarity check: increasing current causes a larger voltage drop across a resistor—does that cause the voltage at one end of the resistor to go up or down relative to ground? (The answer depends on which way the current is flowing!)

Some sanity checks are just completeness checks: does every 2-port component have both ports hooked up in a schematic? Are both power and ground wired to every amplifier? Is every block of the block diagram expanded into a schematic?

Some are simple consistency checks—beginning students often draw wires that create short circuits to ground or power in their schematics. Is everything connected by wires in the schematic supposed to be at the same voltage?—that is what putting the wire in the schematic means.

There are lots of other sanity checks that can be done, and students need to get into the habit of looking for sanity checks that they can do on their designs—not relying on an instructor to check their work for them (or even tell them explicitly what sanity checks to do).

> I want students to learn *skills*, not facts, in this course. (All that students need to memorize is summarized on the study sheet Appendix B.) The skills I want students to have by the end of the course are to be able to design and build simple amplifier circuits and to write design reports. I don't care much whether they can work textbook problems—what I want is that they acquire the mental attitudes of engineers: that they can design and build things, that data sheets are worth consulting, that precise and accurate recording of what was designed and measured is essential, that often you have to check things for yourself (not blindly trusting the data sheets or simple models), that consistency and sanity checks are an important part of any problem solving, that breaking a problem into subproblems is an essential element of design in any engineering field, that one can improve one's ability with practice, and so forth.

Of course, developing these attitudes takes more than one textbook or one course, but for many students this course will be their first exposure to engineering ways of thinking, so I hope to make the most of it.

1.4 Working in pairs

Because of space limitations in the lab and to improve learning in the lab, all labs will be done by pairs of students (unless there are an odd number of people in class, in which case we will have a singleton, not a triple). For each lab the pairing will be different, so that no one has an unfair advantage or disadvantage from consistently being paired with a more or less competent partner.

Rotating partners for labs has the further pedagogic advantage of learning to work with people who have different styles of work—and realizing what work behaviors are particularly annoying, so that you can try to avoid those behaviors yourself.

For each lab, the partners have to choose whether to turn in a joint report with both names on it as co-authors, or separate reports with one author each, but explicitly acknowledging in writing the work done by the other partner. Both partners should keep their own lab notebooks, as they may not have access to their partner's lab notebook later in the quarter.

On joint reports, both partners are fully responsible for everything in the report and get identical grades. It is very important that you check your partners work at least as carefully as you check your own. As the Russian proverb goes Доверяй, но проверяй *"doveryai, no proverayai"* (trust, but verify). President Ronald Reagan became very fond of this Russian proverb after he learned it.

1.5 Learning outcomes

UCSC requires faculty to come up with a list of learning outcomes for any course that they create. This section lists the outcomes I used for the 2018 offering of the course.

Students will be able to

- draw useful block diagrams for amplifier design.
- use simple hand tools (screwdriver, flush cutters, wire strippers, multimeter, micrometer, calipers, …).
- hand solder through-hole parts and SOT-23 surface-mount parts.
- use USB-controlled oscilloscope, function generator, and power supply.
- use Python, gnuplot, PteroDAQ data-acquisition system, and WaveForms 3 on own computer.
- do computations involving impedance using complex numbers.
- design single-stage high-pass, band-pass, and low-pass RC filters.
- measure impedance as function of frequency.
- design, build, and debug simple op-amp-based amplifiers.
- draw schematics using computer-aided design tools.
- write design reports using LaTeX and biblatex.
- plot data and theoretical models using gnuplot.
- fit models to data using gnuplot.

1.6 Videos for the course

Video lectures for much (but not all) of the material in the book is available on YouTube on two playlists that correspond to the two halves of the course as it was taught in 2020–21. The playlists are at https://tinyurl.com/electronics-A (about 27 hours) and https://tinyurl.com/electronics-B (about 12 hours).

These videos are intended to supplement the textbook, not replace it. It is particularly important for students to read the instructions for each lab and not rely just on the partial demos in the videos.

2: Background material

Although this book was originally intended for students who have already successfully completed a calculus-based physics course on electricity and magnetism, students with less background have successfully completed courses based on the book, and the prerequisite for the course was reduced to just a first course in differential calculus.

In this chapter I'll try to review briefly the material we'll use from algebra, calculus, and physics courses, so that students can fill in gaps in their preparation by looking up and learning the background material.

2.1 Metric units

This book will use, as much as possible, standard metric prefixes and units.

Table 2.1 lists the standard metric prefixes, but not all of them are in common use for electronics. Most often, we use the powers of 1000 (kilo-, mega-, giga- going up and milli-, micro-, nano-, pico- going down). Occasionally we need slightly larger or smaller prefixes (tera-, femto-, atto-). Know the prefixes and use them correctly, paying particular attention to whether they are upper-case or lower-case. I don't want to see anyone calling for 100 MV signals, when they mean 100 mV—an error that is a factor of 1,000,000,000 is not negligible.

Some people use "u" instead of "μ" when typing documents, because "μ" was difficult to type before Unicode was defined and became widely used. In some applications, such as computer files that are encoded in ASCII rather than Unicode, this usage is still acceptable, but for the design reports in this class "μ" should be used whenever a multiplier of 10^{-6} is needed.

> The most common usage of metric prefixes is to make all numbers be in the range $[1.0, 1000.)$—whenever a number would be 1000 or larger, the next larger metric prefix should be used, and whenever a number would less than 1, then the next smaller metric prefix should be used.
>
> If you *must* use a number smaller than 1, then the leading zero has to be included, as the decimal point is lost too easily. It is OK to report "0.5 V" instead of "500 mV", but it is never OK to write ".5 V".

For some unknown reason, many electrical engineers avoid the use of "nF", preferring to report a capacitor size of 0.047 μF, rather than 47 nF. Personally, I prefer to consistently keep the numbers in the range 1 to 999.999..., and I see no reason to avoid "nF".

Avoid writing units like mm^2 or cm^2, because it is not immediately clear whether you mean $\text{m}(\text{m}^2)$ or $(\text{mm})^2$, though most engineers will read it as $(\text{mm})^2$. Either include the parentheses explicitly, or use floating-point notation for the numbers and use the raw metric unit without prefixes: $3.3\,(\text{mm})^2 = 3.3\text{E-}6\,\text{m}^2$.

Multiple	Prefix	Abbreviation
10^{30}	quetta-	Q
10^{27}	ronna-	R
10^{24}	yotta-	Y
10^{21}	zetta-	Z
10^{18}	exa-	E
10^{15}	peta-	P
10^{12}	tera-	T
10^{9}	giga-	G
10^{6}	mega-	M
10^{3}	kilo-	k
10^{2}	hecto-	h
10^{1}	deka-	da
10^{-1}	deci-	d
10^{-2}	centi-	c
10^{-3}	milli-	m
10^{-6}	micro-	μ
10^{-9}	nano-	n
10^{-12}	pico-	p
10^{-15}	femto-	f
10^{-18}	atto-	a
10^{-21}	zepto-	z
10^{-24}	yocto-	y
10^{-27}	ronto-	r
10^{-30}	quecto-	q

Table 2.1: These are the standard prefixes in the metric system. Case is very important: M and m have a ratio of 10^9 *and Y and y have a ratio of* 10^{48}*. Warning: k for kilo- is always lower-case—"K" is used for kelvin. In electronics we mostly use the prefixes that are powers of 1000, from* pico- *to* giga-*. The prefixes* hecto- *and* deka- *are not used, and* deci- *is used only for* decibel *(dB).*

Table 2.2 has a table of the standard metric units that we'll be using repeatedly in this course. It isn't necessary to remember all the translations to fundamental units, but which unit is associated with which concept is essential knowledge.

If you need to use a unit as a noun, you must spell it out, not use the unit abbreviation. This restriction is particularly important for plurals: Ω already means "ohms", and Ωs means "ohm-seconds" (a unit that has no use that I'm aware of). Be careful about capitalization—we use both seconds (s) and siemens (S) in this course.

property	unit	symbol	fundamental units
time	second	s	s
distance	meter	m	m
volume	liter	L	$1000\,\mathrm{L} = \mathrm{m}^3$
mass	gram	g	g
temperature	kelvin	K	K
frequency	hertz	Hz	s^{-1}
force	newton	N	$\mathrm{kg\ m\ s}^{-2}$
pressure	pascal	Pa	$\mathrm{N/m}^2 = \mathrm{kg\ m}^{-1}\ \mathrm{s}^{-2}$
energy	joule	J	$\mathrm{N\ m} = \mathrm{kg\ m}^2\ \mathrm{s}^{-2}$
power	watt	W	$\mathrm{J/s} = \mathrm{kg\ m}^2\ \mathrm{s}^{-3}$
current	ampere	A	A
charge	coulomb	C	A s
potential difference	volt	V	$\mathrm{W/A} = \mathrm{kg\ m}^2\ \mathrm{s}^{-3}\ \mathrm{A}^{-1}$
resistance	ohm	Ω	$\mathrm{V/A} = \mathrm{kg\ m}^2\ \mathrm{s}^{-3}\ \mathrm{A}^{-2}$
conductance	siemens	S	$1/\Omega = \mathrm{s}^3\ \mathrm{A}^2\ (\mathrm{kg})^{-1}\ \mathrm{m}^{-2}$
capacitance	farad	F	$\mathrm{C/V} = (\mathrm{kg})^{-1}\ \mathrm{m}^{-2}\ \mathrm{A}^2\ \mathrm{s}^4$
inductance	henry	H	$\mathrm{Vs/A} = \mathrm{kg\ m}^2\ \mathrm{A}^{-2}\ \mathrm{s}^{-2}$

Table 2.2: Metric units used in this book. Be careful about capitalization—unit names are not capitalized, even when the unit is named after a person, but many of the symbols are capitalized. ***Capitalization matters****—we'll use both picoamps (pA) and pascals (Pa) in this book, and it is not always easy to tell which units are meant from context. Although international usage allows either lower-case or upper-case L for liter, US standards call for upper-case only (to avoid confusion between "l" and "1").*

2.2 Dimensional analysis

One important method for making sure your computations are meaningful is to keep the units with the numbers for all intermediate results.

When you multiply two numbers together, the units multiply also, and you can only add or subtract numbers when they have identical units.

Table 2.2 is useful for translating units to the underlying fundamental units, when combining units that look like they are different.

Worked Example:
For example, one formula we often deal with is product of a resistance and a capacitance (as in Section 11.2).

If we have $470\,\Omega$ and $10\,\mu$F, their product is $4.7\,\mathrm{mF}\Omega$, but we can use the definitions of ohms and farads to convert to standard units:

$$470\,\Omega\ 10\,\mu\mathrm{F} = 4.7\,\mathrm{mF}\Omega = 4.7\,\mathrm{mC/V\ V/A} = 4.7\,\mathrm{mC/A} = 4.7\,\mathrm{ms}\ .$$

We use this particular product often enough that is worth remembering that ohms times farads is seconds ($\mathrm{F}\ \Omega = \mathrm{s}$).

We can often use dimensional analysis to help solve problems and remember formulas. For example, if we are trying to determine a capacitance (in F) and are given a voltage (in V) across the capacitor, then we can look for a charge (in C), which may come from integrating a current (in $\mathrm{A} = \mathrm{C/s}$).

Measurements in the US are often given in awkward, non-metric units (inches, feet, miles, pounds force, pounds mass, ounces, fluid ounces, pints, gallons, cubic feet, cubic yards, square inches, acres, …). One of the first steps in using such measurements is to convert to standard units:

$$\begin{aligned} 17\,\mathrm{inch} &\rightarrow 17\,\mathrm{inch}\ 25.4\,\mathrm{mm/inch} = 431.8\,\mathrm{mm}\ , \\ 4\,\mathrm{lb}_{\mathrm{force}} &\rightarrow 4\,\mathrm{lb}_{\mathrm{force}}\ 4.44822\,\mathrm{N/lb} = 17.19\,\mathrm{N}\ , \\ 20\,\mathrm{psi} &\rightarrow 20\,\mathrm{psi}\ 6894.76\,\mathrm{Pa/psi} = 137895\,\mathrm{Pa}\ . \end{aligned}$$

Exercise 2.1

If we have a current of 22 mA running for 200 ms, how much charge has been transferred? (Write out the dimensional analysis, not just the final value.)

Exercise 2.2

If we have a voltage drop of 2.7 V across a $47\,\Omega$ resistor, what is the current through the resistor? (Write out the dimensional analysis, not just the final value.)

2.3 Logarithms

2.3.1 Definition of logarithms

Logarithms are used extensively in electronics, in three forms: *base-10 logarithms* (sometimes called *common logarithms*), *base-2 logarithms*, and *natural logarithms.* All are inverses of exponentiation: $x = \log_{10}(y)$ means $10^x = y$, $x = \log_2(y)$ means $2^x = y$, and $x = \ln(y)$ means $e^x = y$.

The interesting properties of logarithms come from the interesting properties of exponentiation:

- $e^{(x+y)} = e^x e^y$, so $\ln(a) + \ln(b) = \ln(ab)$.
- $e^{(x-y)} = e^x / e^y$, so $\ln(a) - \ln(b) = \ln(a/b)$.
- $e^0 = 1$, so $\ln(1) = 0$.
- $e^{-y} = 1/e^y$, so $-\ln(b) = \ln(1/b)$.
- $\frac{\mathrm{d}e^x}{\mathrm{d}x} = e^x$, so $\frac{\mathrm{d}\ln(a)}{\mathrm{d}a} = 1/a$.
- For small ϵ, $e^\epsilon \approx 1 + \epsilon$, so $\ln(1 + \epsilon) \approx \epsilon$.

2.3.2 Expressing ratios as logarithms

Electrical engineers often express ratios in logarithmic terms, using *decibels (dB)*:

$$D = 20 \log_{10} \left(\frac{A}{A_{ref}} \right) \; ,$$

where D is in decibels, A is the amplitude of the signal, A_{ref} is the amplitude of the reference being compared to. You will often see A_{ref} given in the form "0 dB is A_{ref}".

For example, an amplifier whose output is 10,000 times its input can be described as having a gain of $20 \log_{10}(10000) = 80\,\mathrm{dB}$. Decibels are usually used only when the ratio is unitless—that is, both the numerator and denominator of the ratio have the same units, so that they cancel.

Sometimes we have a gain of less than 1 (from voltage dividers or passive filters, for example), which results in a negative value in decibels. For example, a gain of 0.1 could be expressed as $-20\,\mathrm{dB}$. Because some people are uncomfortable with negative numbers, such gains are sometimes expressed as *attenuations*: a gain of $-20\,\mathrm{dB}$ may be called an attenuation of 20 dB. Whichever terminology is used, the ratio of the output to the input is 0.1.

Because amplifiers and filters are usually set up in way that causes their gains to multiply, the decibel gains can be added. For example, if we have an amplifier that multiplies by 10, followed by a filter that cuts our signal in half, followed by another amplifier with a gain of 100, the overall gain of the system is $20\log_{10}(10\frac{1}{2}100) = 20 - 6.02 + 40 = 53.98\,\mathrm{dB}$. It is common to approximate a factor of 2 as 6 dB.

The definition of decibels given here is for amplitude (of voltage, current, or other signals)—when we are talking about *power*, the definition changes to $10\log_{10}(P/P_{ref})$. That was the original definition of decibel (hence, the *deci-* prefix), and the more commonly used amplitude definition comes from the relationship between power and voltage or current with a resistor:

$$P = VI = V^2/R = I^2R\ .$$

Occasionally one will see base-10 logarithms expressed in *decades* rather than scaled to decibels. Frequency ranges are more often given in decades than in decibels, which are used more for voltage, current, or power ratios. Each decade is a factor of 10, so a frequency range from 1 Hz to 1 MHz might be expressed as six decades. Each decade is twenty decibels, so the range could also be expressed as 120 dB.

Base-2 logarithms are mainly used in talking about numbers represented in a computer, and for the analog-to-digital and digital-to-analog converters used to convert between voltages or currents and digital numeric representations. A 16-bit analog-to-digital converter can recognize $2^{16} = 65536$ different values. You will sometimes see fractional numbers of bits, when the number of distinguishable values is not a power of two. For example, a 16-bit analog-to-digital converter may be described as having 13.7 *effective bits*, if the noise level is ± 2.5 counts, so that only $65536/5 \approx 2^{13.7}$ levels are really meaningful.

Base-2 logarithms also come up when talking about frequencies, with the unit being the *octave*: a ratio of two in frequency. This usage comes originally from music applications, where frequency ratios of factors of two are particularly important.

Musical instruments are usually tuned so that the A above middle C is 440 Hz, a standard established in 1834 in Stuttgart, though still not universally used [117]. Three octaves lower would be a frequency of $2^{-3}440 = 55\,\mathrm{Hz}$.

Natural logarithms occur ubiquitously in physics and electronics as the solutions of first-order differential equations. Perhaps the most powerful use of natural logarithms comes from their use with complex numbers (see Section 2.4). Sometimes authors use the unit *nat* for steps of 1 in natural logarithm, in analogy to *bit* for steps of 1 in logarithms base 2. We will not use nats in this book.

Exercise 2.3

Simplify $-\ln\left(\frac{a-b}{c-d}\right)$.

Exercise 2.4

How many decibels is a gain of 2000? How many decades?

2.3.3 Logarithmic graphs

Logarithms are often used to scale the axes of graphs. This rescaling is usually visible as a nonuniform set of tick marks on the axis, with ticks at 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 20, 30, … . In gnuplot, specifying a log scale on the y-axis is done with the commands

```
set logscale y      # specifies the logarithmic scaling
set mytics 10       # specifies the standard tick marks for log scale
```

Using a logarithmic scale does not change the label of the axis nor the labeling of the tick marks—if I have frequency on a log scale, I still call it “frequency [Hz]”, *not* “log frequency”, and the values on the major ticks are 1, 10, 100, … , not 0, 1, 2, … .

> The main virtue of a logarithmic scale is that distances correspond to *ratios* of the corresponding numbers, while on a linear scale distances correspond to *differences* of the corresponding numbers.

When a ratio is expressed in decibels, a linear scale is used rather than a logarithmic one, because the logarithmic conversion has already been done in the definition of decibels (see Section 2.3.2)—differences of decibels already correspond to ratios of the underlying quantities, so we want distances to correspond to differences of decibels, not ratios of decibels. (Ratios of decibels are almost never meaningful.)

Judicious use of log-scaling for axes can make functions much easier for people to understand visually. As a first approximation, people only understand straight-line graphs visually—they almost always project graphs to continue in a straight line past the ends of what is plotted, and interpolate between points with straight lines.

So what functions are easily understood? Refer to Figure 2.1, which shows the same functions on each of the four major plot types.

lin-lin Graphs with both axes expressed linearly have straight lines that follow formulas like $y = ax + b$, which are known as *affine functions*. Affine functions are more general than *linear functions*, which have the form $y = ax$ and must provide a 0 output for a 0 input. Affine functions include the linear functions, but they also include the *constant functions*, $y = b$, which are most definitely not linear.

The parameters of the straight line are the slope a, which is given in y-axis units per x-axis unit, and the offset (or intercept) b, which is given in y-axis units.

log-lin Graphs with a logarithmic y-axis and linear x-axis (known as log-lin or semilogy graphs) have straight lines that follow formulas of the form $\log(y) = ax + b$, $y = e^{ax+b}$, or $y = BC^x$ (where $B = e^b$ and $C = e^a$), called *exponential functions*.

Growth for bacteria is often well-modeled as an exponential function of time in the early stages, and so growth curves are often appropriately plotted on semilogy plots. One often sees exponential functions in electronics as a result of RC discharge curves (see Section 16.2, for example)—semilogy plots of voltage vs. time are appropriate for such curves, but only when the destination voltage is 0 V.

The slope of the line on a semilogy plot corresponds to the growth rate of the function and can be expressed in units like dB/s or doublings per day for exponential functions

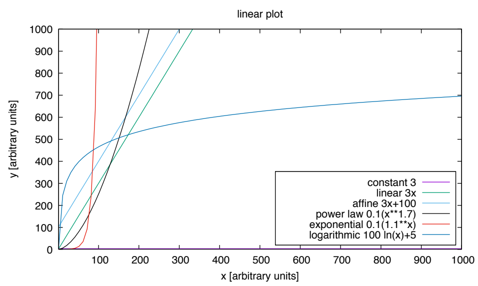

(a) The linear plot makes constant, linear, and affine functions straight lines. The constant function is hard to see, because it is nearly buried in the x-axis.

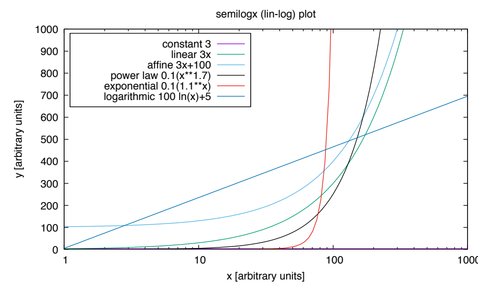

(b) The semilogx plot makes constant and logarithmic functions straight lines. The constant function is hard to see, because it is nearly buried in the x-axis.

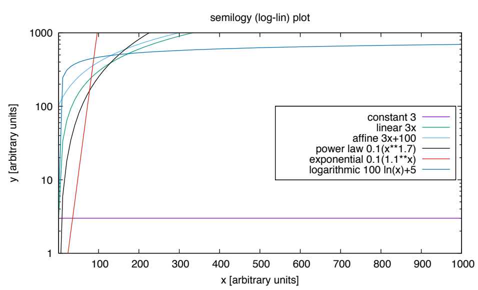

(c) The semilogy plot makes constant and exponential functions straight lines.

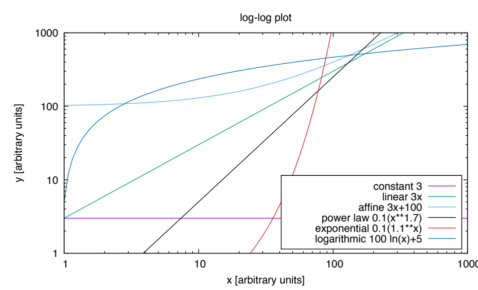

(d) The loglog plot makes constant, linear, and other power-law functions straight lines.

Figure 2.1: These four plot types are the main ones we'll use in this course. The constant function is a flat, straight line in all plots types, but other functions are straight lines in only one or two of the plot types. Choose your plot type to make the function you are interested in showing as near to a straight line as you can.

Logarithmic scales are also useful for showing functions or data that have a very wide range, when the ratios of values are more interesting than the differences between values.

Figure drawn with gnuplot [33].

of time. An exponential decay will have a negative slope, while exponential growth has a positive slope.

The offset of the line corresponds to the initial size of the function.

Semilogy plots are also useful for plotting probability density functions, when what we are interested in are the tails of the function. Probability distributions are often easier to identify and more robustly extrapolated on a log y-axis than on a linear y-axis. For example, a Gaussian distribution changes shape from a "bell-shaped" plot on linear y-axis to a simple parabola on a log y-axis. All the interesting variation in the tails disappears with the linear scaling.

lin-log Graphs with linear y-axis and logarithmic x-axis (know as lin-log or semilogx graphs) have straight lines that follow formulas of the form $y = a\log(x) + b$. Such logarithmic functions are occasionally found in electronics (as the inverses of exponential functions, swapping the roles of x and y). We will encounter them in the logarithmic transimpedance amplifier (see Section 23.2).

The proper units for slope are y-axis units per ratio unit: y-axis units/dB, y-axis units/octave, or y-axis units/decade. For example, if we want to convert a musical pitch as a frequency to the corresponding key number for a Musical Instrument Digital Interface (MIDI), the function is

$$k(f) = 69 + 12 \log_2 \frac{f}{440\,\mathrm{Hz}} ,$$

and the corresponding slope is 12 keys per octave [137].

The offset b moves the curve vertically and corresponds to the function value with an input of 1. It is often easier, however, to give the function value for some other input, as 1 may not be in the domain where the function is a good model for the phenomenon. In the example of MIDI tuning, key 69 is the key for A440, the A above middle C that has a frequency of 440 Hz. Note how the expression for the function given above expresses this offset.

log-log Graphs in which both axes as logarithmic have as straight lines functions of the form $\log(y) = a\log(x) + b$ or $y = Bx^a$, where $b = \log(B)$. These are known as *power-law functions*, and include the linear functions $y = Bx^1$. This means that linear functions form straight lines on both lin-lin (as special cases of affine functions) and log-log plots (as special cases of power-law functions).

The slope of a line on a log-log graph corresponds to the exponent on the x term in the power law—the slope is technically unitless, but can be expressed in terms like dB/decade or dB/octave. A linear function ($b = 1$) would be 20 dB/decade or approximately 6 dB/octave. An inverse linear relationship ($b = -1$) would be −20 dB/decade or about −6 dB/octave.

The offset b sets the height of the curve, by changing the scaling factor B.

Log-log plots are used extensively in electronics, because many of the phenomena we are modeling are well approximated by power laws. A particularly important class of log-log plots are the Bode plots that plot gain or impedance vs. frequency (see, for example, Section 11.2).

2.4 Complex numbers

Electrical and electronics engineers use complex numbers extensively, as they provide the most convenient way to talk about and manipulate the amplitude and phase of sinusoidal signals.

The first bit of unfamiliar notation may be the definition of j:

$$j = \sqrt{-1} .$$

Mathematicians shudder at the use of the symbol j rather than i for the square-root of minus one, but electrical engineers reserve i for current, and so have adopted j for imaginary numbers. This tradition has gone on for over 100 years, and neither the mathematicians nor the electrical engineers are likely to change their notation. For this class, we'll use the electrical engineering notation.

Euler's Formula

$$e^{j\theta} = \cos(\theta) + j\sin(\theta)$$

has been called the most beautiful formula in mathematics, connecting exponentials and trigonometric functions in a profound way.

Many of the trigonometric identities so painfully learned in trigonometry classes (and forgotten by a year later) can be replaced by simple algebra on Euler's Formula.

For example, adding two angles:

$$\begin{aligned}
\cos(\phi+\theta) + j\sin(\phi+\theta) &= e^{j(\phi+\theta)} \\
&= e^{j\phi}e^{j\theta} \\
&= (\cos(\phi) + j\sin(\phi))\,(\cos(\theta) + j\sin(\theta)) \\
&= (\cos(\phi)\cos(\theta) - \sin(\phi)\sin(\theta)) \\
&\quad + j\,(\sin(\phi)\cos(\theta) + \cos(\phi)\sin(\theta)) \ .
\end{aligned}$$

Two complex numbers are equal if, and only if, their real parts are equal and their imaginary parts are equal, so the derivation above gives the addition formulas for both cosines and sines, using only simple properties of exponentials and the distributive law.

It is often useful to view the complex numbers as being on a plane, with $z = x + jy$ being the point at coordinates (x, y) (Figure 2.2). In this representation, $z = Ae^{j\theta}$ has a simple geometric interpretation: the point z is distance A away from the origin, along a ray that has angle θ (in radians) counterclockwise from the x-axis.

In this view, Euler's formula expresses the relationship between Cartesian and polar coordinates.

A further useful application of Euler's formula (or of polar coordinates) is the interpretation of multiplication of complex numbers. If we have $z_1 = A_1e^{j\theta_1}$ and $z_2 = A_2e^{j\theta_2}$, then

$$z_1z_2 = A_1e^{j\theta_1}A_2e^{j\theta_2} = A_1A_2e^{j(\theta_1+\theta_2)} \ .$$

That formula means that multiplying by z_1 corresponds to scaling by its magnitude (A_1) and rotating about the origin by its phase (θ_1).

We'll be using complex numbers extensively for discussing impedance (Section 10.2), so thorough understanding is important.

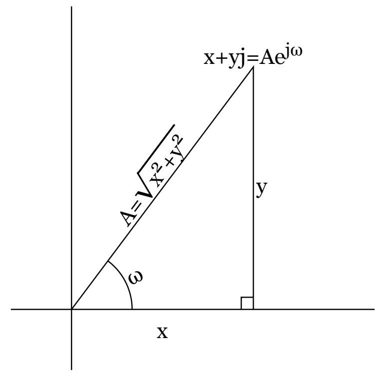

Figure 2.2: Complex numbers may be viewed in Cartesian coordinates, $x + jy$ at (x, y) or polar coordinates (radius $A = \sqrt{x^2 + y^2}$, angle ω radians).

Exercise 2.5

If $A = 1 + j$ and $B = 1 - j$, what are

- $A + B$,
- $A - B$,
- A^2,
- B^2,
- AB,
- Aj,
- Bj,
- A/B, and
- B/A,

expressed in the standard $x + yj$ form?

Exercise 2.6

Express the following in polar $re^{\theta j}$ form:

- 1,
- $1 + j$,
- $3 + 4j$,
- $5 + 12j$, and
- $-5 - 12j$.

Exercise 2.7

Express the following in standard $x + yj$ form:

- $3e^{\pi j}$,
- $e^{(\pi/4)j}$,
- $2e^{2\pi j}$, and
- $13e^{13j}$.

Exercise 2.8

Plot (by hand or with a program) all the complex numbers in Exercise 2.6 and Exercise 2.7.

2.5 Derivatives

The book assumes that students have had a semester of differential calculus, but we don't need all the material from such a course, as only a few simple rules for taking derivatives are needed.

There are three fundamental derivatives you should know:

$$\frac{\mathrm{d}}{\mathrm{d}t}\sum_n a_n t^n = \sum_n a_n n t^{n-1}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} e^{\lambda t} = \lambda e^{\lambda t}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\ln(f(t)) = \frac{\frac{\mathrm{d}f(t)}{\mathrm{d}t}}{f(t)}$$

To reduce writing, engineers often use the shorthand notation $f'(x)$ for $\frac{\mathrm{d}f(x)}{\mathrm{d}x}$ and $\dot{f}(t)$ for $\frac{\mathrm{d}f(t)}{\mathrm{d}t}$.

You should also know the following combining forms:

$$\begin{aligned}(u+v)' &= u' + v' \\ (uv)' &= u'v + uv' \\ (u/v)' &= \frac{u'v - v'u}{v^2} \\ (f(g(x))' &= f'(g(x))g'(x)\end{aligned}$$

Because this course is an electronics course, not a calculus course, it will be acceptable on homework and design reports to use computer tools such as Mathematica and WolframAlpha to take derivatives and solve equations. Electronics knowledge is needed to set up the appropriate equations—solving the equations does not have to be done by hand.

One common use for derivatives in this course is talking about *gain* or *sensitivity*, both of which are defined as the derivative of the output of a system with respect to its input. When the inputs and outputs are both electrical signals (as in an amplifier), we usually use the term *gain*, but when the input is a physical measurement (like temperature or pressure) and the output is electrical, we usually use the term *sensitivity*. Some authors refer to the sensitivity of a sensor as its *responsivity*, but that term has still not gotten widespread acceptance, so we will continue to use *sensitivity*.

The *gain* of a voltage amplifier is

$$G = \frac{\mathrm{d}V_{\mathrm{out}}}{\mathrm{d}V_{\mathrm{in}}} \ .$$

The *sensitivity* of a pressure sensor whose input is pressure P and whose output is voltage V is

$$S = \frac{\mathrm{d}V}{\mathrm{d}P} \ .$$

Some people mistakenly think of sensitivity or gain as “output over input”, but this leads to many mistakes due to DC offsets in voltage. It is much better to think of them as “change in output over change in input”.

Pay attention to units when taking derivatives. The units for $\frac{\mathrm{d}f(x)}{\mathrm{d}x}$ are the units for f divided by the units for x. For example, if both the input and the output of the function are in volts, then the gain is unitless (but is sometimes written as V/V to remind the reader that it is the ratio of voltages). If the input is pressure (in pascals) and the output is voltage (in mV), then the units for the sensitivity would be mV/Pa.

Exercise 2.9

If an amplifier circuit has the function $V_{\text{out}} = 27V_{\text{in}} + 20\,\text{V}$, what is the gain of the amplifier?

Exercise 2.10

If a temperature sensor with a resistance output has the function

$$R = 10\,\text{k}\Omega e^{3977\,\text{K}/\text{T} - 3977\,\text{K}/298.15\,\text{K}},$$

where T is the temperature in kelvins, and R is the resistance in ohms, then what is the sensitivity of the sensor at an arbitrary temperature T? What is it at 35°C?

2.6 Optimization

Another use for derivatives is in *optimization*, which is any method that finds the values of the inputs to a function that maximizes the output. We usually apply optimization to mathematical functions that describe the behavior of a physical or electronic system that we are interested in.

There are many optimization techniques, and which one to use depends on the nature of the function that we are optimizing, any constraints on the inputs to the functions, and whether we need exact or approximate solutions.

For this course, we will look only at one of the simplest optimization techniques, which is useful for continuous, differentiable functions of one variable with no domain limitations. For such functions, we can use the derivative to find the maxima. If the function $f(x)$ has a maximum at $x = x_m$, then the derivative of the function must be 0 there: $f'(x_m) = 0$. To find the maximum of f, we look at every value of x for which $f'(x) = 0$ and determine whether those points are maxima or not.

If we have simple domain constraints on a function (for example, that the input needs to be nonnegative), we can add the boundaries of the domain to the set of points to check for maxima—this is easy when we have only a single variable to optimize, but more sophisticated approaches are needed when we have multiple variables to optimize. We will not need the more sophisticated optimization techniques in this book.

Exercise 2.11

For what value(s) of x is $f(x) = -14x^2 + 7x + 32$ maximized?

Exercise 2.12

If v is constrained to the interval $v \in [0, 5]$, for what value(s) is $f(v) = v^3 - 9v^2 + 24v - 20$ maximized?

2.7 Inequalities

Many constraints and design goals are expressed as *inequalities*, not equations, and engineers are expected to be able to manipulate the inequalities to convert constraints on inputs or outputs to constraints on component values. Converting inequalities to equations is a bad idea, because doing so causes confusion about which side of the "solution" has the legal values, and it encourages engineers to design right at the limits of the constraints, rather than staying safely far from the limits.

> When we get constraints on component values, voltages, or currents, we almost always want to stay far away from the constraints in real designs. Otherwise small variations in any of the values that go into the inequality can cause the constraint to be violated.

If we have an inequality of the form

$$A < B\ ,$$

we can do the following operations to change the inequality:

$$\begin{aligned} A + x < B + x, &\qquad \text{for any real } x \\ A - x < B - x, &\qquad \text{for any real } x \\ Ax < Bx, &\qquad \text{for any positive } x \\ Bx < Ax, &\qquad \text{for any negative } x \\ A/x < B/x, &\qquad \text{for any positive } x \\ B/x < A/x, &\qquad \text{for any negative } x \\ e^A < e^B & \\ A^k < B^k, &\qquad \text{for any positive odd } k\ . \end{aligned}$$

If we have an inequality of the form

$$0 < A < B$$

(equivalent to the pair of inequalities $0 < A$ and $A < B$), we can do the following additional operations to change the inequality:

$$\begin{aligned} 0 < 1/B < 1/A & \\ \log(A) < \log(B) & \\ 0 < \sqrt{A} < \sqrt{B} & \\ 0 < A^x < B^x, &\qquad \text{for any positive } x \\ 0 < B^x < A^x, &\qquad \text{for any negative } x\ . \end{aligned}$$

If we have two inequalities

$$\begin{aligned} A &< B \\ C &< D\ , \end{aligned}$$

we can combine them to get

$$A + C < B + D \ ,$$
$$A - D < B - C \ .$$

If all the numbers are positive,

$$0 < A < B$$
$$0 < C < D \ ,$$

then we can do more combining:

$$AC < BD \ ,$$
$$A/D < B/C \ .$$

Inequalities can only be applied to real numbers, not complex numbers, because complex numbers are not an ordered field.

Let's look at a few examples.

Worked Example:
If we have a 5 V power supply with a 2 A current limit, what resistances can we use as a load?
We have to combine an equation (Ohm's law—see Section 4.3)

$$V = IR$$

with an inequality

$$I < 2\,\mathrm{A} \ ,$$

to get

$$5\,V/R = I < 2\,\mathrm{A} \ .$$

We can multiply both sides by R and divide both sides by 2 A to get

$$2.5\,\Omega < R \ ,$$

which is the desired constraint on resistance. Note that this solution does *not* mean that 2.5 Ω is a desirable value for R—in fact, it would almost always be a bad choice, as slight variations in any of the parameters (voltage, current limit, or resistance) would cause the constraint to be violated.

Worked Example:
If we are making an RC filter with corner frequency $f_c < 2$ Hz using a capacitor with capacitance $C \leq 10\,\mu F$, what values can we use for the resistor R?
The concepts and formulas for RC filters are given in Chapter 11, and the formula for the corner frequency is

$$f_c = \frac{1}{2\pi RC} \ .$$

We can apply that formula and the first inequality in the question to get

$$\frac{1}{2\pi RC} < 2\,\mathrm{Hz}\ ,$$

which we can invert (because all values are positive) to get

$$0.5\,\mathrm{s} < 2\pi RC\ .$$

Because the constraint on C sets a maximum value for C, we can do division to get

$$\frac{0.5\,\mathrm{s}}{10\,\mu\mathrm{F}} < 2\pi R\ ,$$

which we can simplify to

$$7957.747\,\Omega < R\ .$$

We can pick any larger value of R—staying close to the minimum is rarely desirable. For example, we could pick $C = 100\,\mathrm{nF}$ and $R = 820\,\mathrm{k}\Omega$ to get a corner frequency of

$$f_c = \frac{1}{2\pi RC} = \frac{1}{2\pi 820\,\mathrm{k}\Omega\ 100\,\mathrm{nF}} \approx 0.97\,\mathrm{Hz}\ .$$

Worked Example:
If we have a differential amplifier with a gain of 10, and an output range of 1 V to 2 V with a 0 V input centered at 1.5 V, what is the legal range of the input signal?
Again we need to find an equation for the behavior of the amplifier, to combine with the inequalities of the constraint

$$1\,\mathrm{V} < V_{\mathrm{out}} < 2\,\mathrm{V}\ .$$

In this case, the equation is the gain equation of a differential amplifier, from Section 18.2.1,

$$G = \frac{V_{\mathrm{out}} - V_{\mathrm{ref}}}{V_{\mathrm{p}} - V_{\mathrm{m}}}\ ,$$

where G is the gain, $V_{\mathrm{p}} - V_{\mathrm{m}}$ is the differential input signal, and $V_{\mathrm{ref}} = 1.5\,\mathrm{V}$ is the reference voltage that the amplifier uses with a $0\,\mathrm{V}$ input.

We can rearrange the gain equation to get

$$V_{\mathrm{out}} = V_{\mathrm{ref}} + G(V_{\mathrm{p}} - V_{\mathrm{m}})\ ,$$

and combine with the inequalities to get

$$1\,\mathrm{V} < V_{\mathrm{ref}} + G(V_{\mathrm{p}} - V_{\mathrm{m}}) < 2\,\mathrm{V}\ .$$

Plugging in the given values for V_{ref} and G gives us

$$1\,\mathrm{V} < 1.5\,\mathrm{V} + 10(V_{\mathrm{p}} - V_{\mathrm{m}}) < 2\,\mathrm{V}\ ,$$

which we can simplify in two steps:

$$-0.5\,\mathrm{V} < 10(V_\mathrm{p} - V_\mathrm{m}) < +0.5\,\mathrm{V}$$
$$-50\,\mathrm{mV} < V_\mathrm{p} - V_\mathrm{m} < +50\,\mathrm{mV}\ ,$$

which is the desired constraint on the input signal.

Exercise 2.13

We have an amplifier whose gain equation is $V_\mathrm{out} - 1.65\,\mathrm{V} = G(V_\mathrm{in} - 1.65\,\mathrm{V})$ and with the constraint on the output $0.05\,\mathrm{V} < V_\mathrm{out} < 3.25\,\mathrm{V}$. If the input signal is $1.65\,\mathrm{V} \pm 0.1\,\mathrm{V}$, what are the constraints on gain G?

Exercise 2.14

If we want a corner frequency $f_c = 1/(2\pi RC)$ between $1\,\mathrm{Hz}$ and $2\,\mathrm{Hz}$, and the capacitance C is $100\,\mathrm{nF}$, what are the constraints on R?

22: Optoelectronics

22.1 Semiconductor diode

A *diode* is a two-terminal device that allows current to flow in one direction (the *forward* direction), but not in the other (the *reverse* direction). Figure 22.1 shows the schematic symbol used for diodes, as well as the symbols for optoelectronic parts based on diodes.

A *semiconductor diode* is a semiconductor with two different regions: each with slightly different materials. One region has a carefully controlled impurity resulting in free electrons in the crystal structure (referred to as *n-doped*, because the charge carriers have negative charge). For example, a silicon semiconductor may be doped with phosphorus. When the crystal structure is formed, four of the five electrons in the outer shell of phosphorus form the crystal bonds and the fifth electron is free to move around in the crystal.

The other region has a different impurity that results in electrons being missing from the crystal structure (referred to as *p-doped*, because the charge carriers, called *holes*, have positive charge). For example, a silicon semiconductor may be doped with boron, which has only three valence electrons. A fourth electron for the crystal structure is borrowed from a nearby atom, resulting in a positively charged hole that can move around, though it is not quite as mobile as the free electrons of n-doped silicon.

The junction between the two types of semiconductor is where everything interesting happens, as shown in Figure 22.2.

When the n-doped semiconductor is connected to a *higher* voltage than the p-doped semiconductor, then the majority charge carriers on each side (electrons on the n-doped side, positively charged holes on the p-doped side) are pulled away from the junction, and a non-conducting *depletion layer* or *depletion region* is formed. This connection *reverse biases* the diode, so that no current flows. When there is no applied voltage from the outside, there is still a small depletion layer formed (see Figure 22.2a). A reverse-biased diode acts like a capacitor, with two conducting plates separated by an insulator, but the capacitance varies with the bias voltage.

When the n-doped semiconductor is connected to a *lower* voltage than the p-doped semiconductor, then the majority carriers in each type of semiconductor are pushed towards the junction and cross over, allowing current to flow (*forward biasing*, see Figure 22.2b). The voltage has to be large enough to eliminate the normal depletion layer before conduction starts—this minimum forward voltage is often called the *diode voltage* or *threshold voltage*.

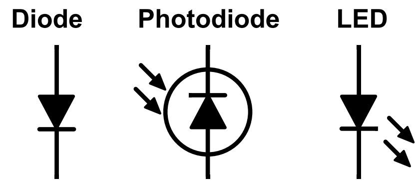

Figure 22.1: The schematic symbols for diodes, light-emitting diodes (LEDs), and photodiodes are essentially all the same, with the addition of arrows to indicate light coming out of LEDs and going into photodiodes. LEDs sometimes have circles around them, and photodiodes sometimes do not—the arrows are the essential difference, not the circles. Sometimes the arrows are drawn with wavy lines instead of straight ones. The triangles in the diode symbols indicate the direction of conventional current flow (forward biasing). The positive end (the anode*) points to the negative end (the* cathode*)—the -ode names can be kept straight if you remember that the cathode rays of cathode-ray tubes are electron beams, so the cathode is where the electrons leave the wire.*

The photodiode is normally used reverse biased (with the cathode more positive than the anode), and so I've drawn it with the opposite orientation from the other diodes.

Figure drawn with Digi-Key's Scheme-it [18].

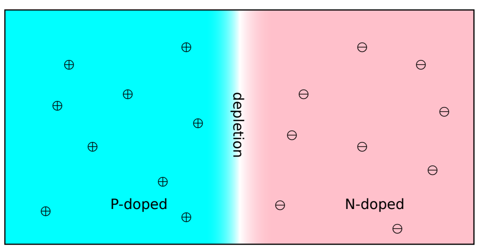

(a) A diode junction with no applied voltage has a small, non-conducting depletion region, and so acts as a capacitor.

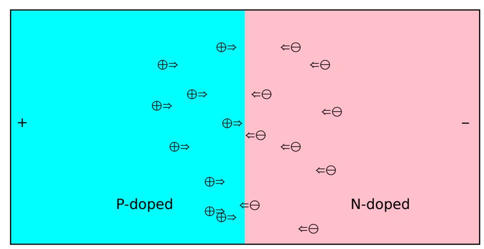

(b) When the p-doped region is set to a higher voltage than the n-doped region, the charge carriers are pushed towards the junction. When the voltage is high enough, the depletion region is eliminated, and current flows through the diode.

Figure 22.2: A semiconductor diode does not conduct until a sufficiently high voltage is applied to the p-doped region (the anode) relative to the n-doped region (the cathode).

22.2 Light-emitting diodes (LEDs)

Light-emitting diodes are semiconducting diodes that emit light when there is sufficient forward current through them. LEDs often have a fairly large diode voltage compared to other diodes, and you need to exceed the diode voltage before you get enough current to turn on the LED.

Light is emitted when a free electron from the n-doped side combines with a hole from the p-doped side. The drop in the energy of the electron is released as a photon. Because the electrons vary in energy, the photons do not all have the same wavelength, so LED light is not spectrally pure, but the variation is generally not very large, at least when compared to thermally produced light, like black-body radiation and incandescent light bulbs.

For a no-math view of how LEDs (and semiconductor diodes in general) work, try the *How Stuff Works* website [41].

Data sheets for visible-wavelength LEDs often report two wavelengths: the *peak* wavelength, where the energy from the LED is largest, and the *dominant* wavelength, which takes into account the varying sensitivity of the human eye. The dominant wavelength is closer to the maximum of human sensitivity, around 555 nm [158], than the peak wavelength is. Infrared and ultraviolet LEDs generally report only peak wavelength, because human eye sensitivity is irrelevant for light that is not visible.

When you are trying to look at the LEDs with photodetectors, rather than human eyes, the adjustments for human eye sensitivity are just a nuisance, so you should use only the *peak* wavelength, not the *dominant* wavelength.

The light output of the LED may be reported in different ways. For example, the output luminous intensity of the LED may be reported in candela (lumen per steradian) or watts per steradian, or the total luminous flux of the LED may be reported in lumens or watts. A *steradian* is a measure of solid angle that would cut an area of r^2 out of a sphere of radius r, so a full sphere is an area of $4\pi\,\mathrm{sr}$.

Some of the units are purely physical measures (watts, milliwatts per steradian, watts per square cm), while others have been adjusted for human eye sensitivity (lumens, candelas, lux, foot-candles). The Wikipedia article on photometry [147] gives tables of the SI units for both photometric (adjusted for human eye sensitivity) and radiometric (physical) measurements.

For infrared emitters, only the purely physical units make any sense, as the human-adjusted ones would all report the output as zero. Sometimes a manufacturer stupidly reports the light output of IR emitters in human-adjusted units anyway, relying on the conversion that there are 683.002 lumen/watt at 540 THz (a wavelength of about 555 nm), ignoring the correction for wavelength that is part of the definition of lumens.

The data sheet will also give the forward voltage of the LED, either for a typically used current or as a plot of current vs. forward voltage. The current rises rapidly with voltage (it is well fit by an exponential), and so engineers often use an approximation that the forward voltage is constant over the range of currents they are interested in. The variation from one LED to another from the same batch may be larger than errors introduced by this approximation.

For example, the 151034BS03000 blue LED from Würth Electronics has a peak wavelength of 465 nm and a dominant wavelength of 470 nm [169]. The output at 20 mA is typically 2500 mcd, but could go as low as 850 mcd. (The abbreviation “mcd” stands for “millicandela”, a luminous intensity of one millilumen per steradian.) Also at 20 mA, the forward voltage is typically 2.8 V, but could go as high as 3.6 V. We would not want to specify this part for a 3.3 V system, as it might work fine in prototyping, but then fail to light up in production units, because of changes in the characteristics.

The data sheet for an LED often provides a forward-current vs. forward-voltage plot—a typical cheap red LED might have a 7 mA current at 2 V but 35 mA at 2.2 V. Because an LED is a semiconductor diode, the current goes up exponentially with voltage at low voltages, growing by a factor of ten for every 100 mV. The growth rate slows down for higher voltages, to about a factor of two change in current per 100 mV at the normal operating range of the LED. This means that small changes in the forward voltage can produce large changes in the current, but the forward voltage specification for an LED is often fairly wide. The amount of light produced by an LED is proportional to the current through the LED, not to the voltage.

> From a design standpoint, having light proportional to current and only a roughly specified diode voltage means that we usually try to control the current through an LED, rather than trying to control the voltage.
>
> If you connect an LED directly across the power supply, there is a very high probability of burning it out, as the current goes up exponentially with voltage, and the part has an absolute maximum current.

For the 151034BS03000 blue LED, the maximum continuous current is 30 mA. This maximum current is mainly due to heating the device, so if you only turn the LED on a tenth of the time for 100 μs at a time, you can push the peak current for the 151034BS03000 up to 100 mA through.

To control the current through a small LED, the simplest approach is to use a current-limiting resistor, as shown in Figure 22.3. This approach takes advantage of the forward voltage of the LED being roughly constant over a wide range of currents, so that the voltage across the resistor is just the power-supply voltage minus the forward voltage of the LED. Using Ohm’s law, we can pick a resistance to give us the desired current.

For example, for 151034BS03000 with a 5 V power supply, we would choose $R = (5\,\mathrm{V} - 2.8\,\mathrm{V})/I$, so that to get 6 mA, we would choose a resistor around 366 Ω. The closest in the E12 series is either 330 Ω or 390 Ω, giving currents of 6.7 mA or 5.6 mA, respectively.

If the 151034BS03000 LED does not have its typical value for forward voltage, but the maximum (3.6 V), then 330 Ω would provide only 4.2 mA and 390 Ω only 3.6 mA. The data sheet does not give the minimum forward voltage, which is an unfortunate omission, as we cannot compute how high the current could get in the worst case.

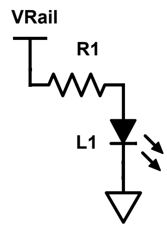

Figure 22.3: One common circuit for setting the current through an LED is a current-limiting resistor. The resistor value R is chosen so that the voltage drop across the resistor produces the desired current via Ohm's law: $V_{\text{rail}} - V_{\text{f}} = IR$. The forward voltage V_{f} of the LED increases with increasing current, but slowly enough that one can often get away with a constant approximation for V_{f}.

Figure drawn with Digi-Key's Scheme-it [18].

Exercise 22.1

Look up the data sheet for the LTL-4234 green LED. Find the peak wavelength, the dominant wavelength, and the forward voltage. Figure out what size current-limiting resistor to use to get approximately a 5 mA current with a 3.3 V power supply—pick a resistor from the E12 series. How much current would there be with the typical forward voltage? What if the forward voltage is the maximum, rather than the typical value?

22.3 Photodiode

A *photodiode* is a semiconductor diode whose junction is exposed to light [146]. To use it as a photodiode, the diode is *reverse-biased*—that is, the positive voltage is connected to the n-doped semiconductor, and the negative voltage to the p-doped semiconductor.

The reverse biasing attracts the majority charge carriers away from the junction between the two types of semiconductor, creating a *depletion* region that does not normally conduct, but when a photon is absorbed in the semiconductor in or near the depletion layer, it can knock an electron loose, resulting in a pair of charges (a hole and an electron), that can move under the influence of the electric field to the conducting regions. This movement of charges creates a current which is proportional to the number of photons absorbed (as long as the hole and electron do not recombine).

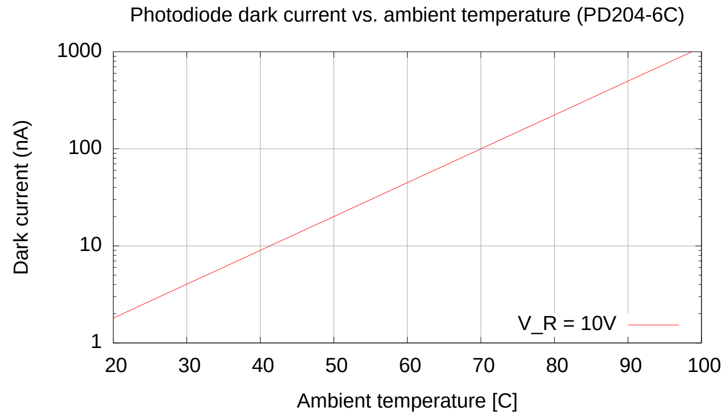

Figure 22.4: The dark current for a photodiode grows exponentially with temperature. This plot gives approximate typical dark currents for a PD204-6C photodiode reverse-biased by 10 V—drawn by fitting a simple exponential to two points from the photodiode data sheet [24].

The photodiode has a photocurrent of only 3.5 μA at $1\,\mathrm{mW/(cm)^2}$ [24], and so a hot photodiode is difficult to tell from one that is detecting light.

Figure drawn with gnuplot [33].

There is a very small current due to thermal effects even with no photons knocking loose electrons (referred to as the *dark current*). To avoid noise problems, the dark current should be much smaller than the smallest current you plan to measure, but the dark current typically grows exponentially with temperature, as shown in Figure 22.4.

Photodiodes tend to be used in one of two ways:

photoconductive mode A fixed reverse-bias voltage is used and the photocurrent is measured. The larger the reverse-bias voltage, the thicker the depletion region, and the lower the capacitance of the photodiode. The reduced capacitance makes the photodiode faster at responding to changes in light level (which is particularly important if light is being used to transmit information—the speed of information transfer may be limited by the response speed of the photodiode).

Increasing the reverse-bias voltage, however, also increases the dark current without changing the photocurrent much, reducing somewhat the dynamic range of the sensor.

photovoltaic mode If we don't allow much current to flow, then the photocurrent charges up the capacitance until the voltage is high enough that the diode starts to conduct in the forward direction. This voltage is almost independent of the amount of light, but the current we can take from the photodiode at that voltage is linear with the amount

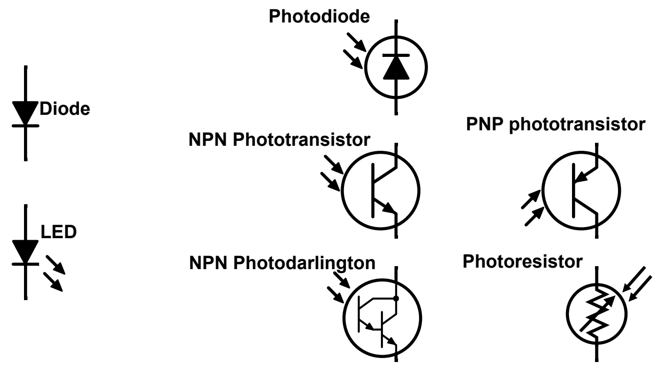

Figure 22.5: The triangles in the diode symbols and arrows in transistor symbols indicate the direction of conventional current flow (forward biasing). All the symbols in this diagram have the end connected to the more positive voltage on top (remember that photodiodes are reverse-biased).

Photodarlingtons are basically phototransistors with an extra transistor for more current gain. They provide larger currents than phototransistors (32 mA seems to be a common number) and are even slower and less linear.

Photoresistors provide a spectral response close to human eyes, but have very slow response to changes in light. Like resistors, they have no preferred direction of current flow.

Figure drawn with Digi-Key's Scheme-it [18].

of light. (This is how a solar cell works—a solar cell is just a very large photodiode used in photovoltaic mode.)

22.4 Phototransistor

Figure 22.5 shows the schematic symbols for photodiodes, phototransistors, photodarlingtons, and photoresistors. We will not use photodarlingtons nor photoresistors in this course.

A phototransistor is a two-terminal device whose current is proportional to the illuminance of the transistor. Unlike photodiodes, phototransistors do not *generate* current—instead they *control* the current that passes through them from an external voltage source.

A phototransistor can be thought of as a photodiode combined with an amplifying bipolar transistor. The two terminals are the collector and the emitter (in the schematic symbols in Figure 22.5, the emitter is the terminal with an arrow, and the collector is the other terminal).

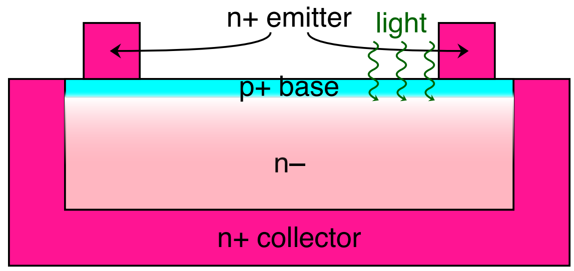

Figure 22.6: Simplified cross section of an NPN phototransistor. The emitter covers only a small portion of the base, so that photons can easily reach the depletion region at the p+/n- junction that is light sensitive.

There are two types of phototransistor: *NPN* and *PNP*, referring to the sandwiching of the types of semiconductor making up the transistor. The phototransistor we have in our kits is an NPN silicon transistor, which means that the emitter and collector are n-doped silicon and the base is p-doped silicon.

The symbols for bipolar transistors (of which phototransistors are a special case) have the collector and emitter as diagonal lines connecting to a vertical line representing the base. The emitter always has an arrow on it, and the arrow points in the direction of conventional current flow when the transistor is on. For NPN transistors, the arrow points away from the base, while for PNP transistors it points towards the base. In Figure 22.5, all the symbols are drawn so that the lower end would be connected to the more negative voltage in normal usage.

For an NPN phototransistor, the collector and emitter are n-doped semiconductor regions, and the base is a thin layer of p-doped semiconductor between the conductor and emitter. The base in a regular bipolar transistor controls the flow of current between the collector and emitter:

- If the base-emitter junction is reverse-biased, then the transistor is turned off, and essentially no current flows.
- If the base-emitter junction is forward-biased, then current flows from the collector to the emitter proportional to the base-emitter current (the current gain of a bipolar transistor is usually quite large—more than 100).

In a phototransistor (see Figure 22.6), the base is not wired to anything—a configuration known as a *floating base*. The floating base normally results in the base having a voltage intermediate between the collector and the emitter. In order to have the forward and reverse biases of the floating base correct, we need to have the base be more positive than the emitter (forward bias) and less positive than the collector (reverse bias). That means that we want the collector wire to be connected to a more positive voltage than the emitter wire: $V_\mathrm{C} > V_\mathrm{E}$.

In fact, the difference $V_\mathrm{C} - V_\mathrm{E}$ (usually called V_CE) needs to be bigger than a constant, $V_\mathrm{CE(sat)}$, called the *saturation voltage* for the phototransistor, in order for the transistor to conduct. V_CE generally needs to be much higher than $V_\mathrm{CE(sat)}$, up in the active region for the bipolar transistor, for the current specifications from the data sheet to be accurate.

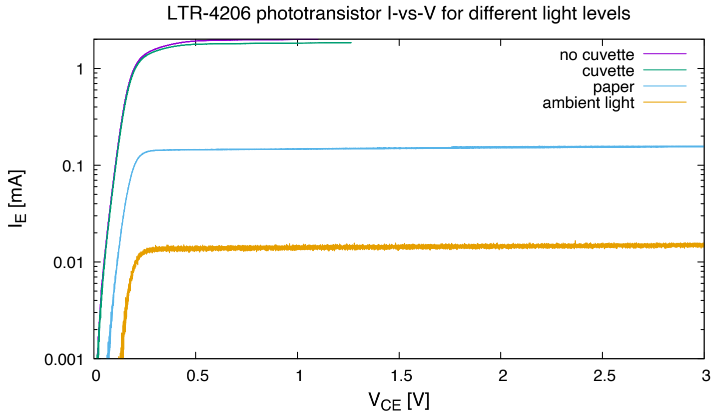

Figure 22.7: This plot shows the current through a phototransistor (an LTR-4206, which may be a different NPN phototransistor than the one in your parts kit) for different levels of illumination. This phototransistor was being used as part of a homemade colorimeter for measuring OD600, using a 333-2UYC/H3/S400-A6 LED as the light source. The current is roughly constant once the collector-to-emitter voltage is high enough (this is called the active *region in bipolar transistors).*

Data collected with PteroDAQ using a Teensy LC board [46]. Figure drawn with gnuplot [33].

> The V_{CE} bias voltage can be arranged either with a bias resistor (like the one in Lab 2) or with a transimpedance amplifier (to be explained in Chapter 23).

The base-emitter current that controls the transistor comes from the photocurrent generated at the reverse-biased base-collector junction. This current then controls the collector-emitter current that we can measure from outside the device.

A bipolar transistor acts as a current amplifier (with the collector current proportional to the base current) when the collector-to-emitter voltage V_{CE} is sufficiently high. The gain of the amplifier drops rapidly when the V_{CE} gets too small. The *saturation voltage* $V_{\mathrm{CE(sat)}}$ of a bipolar transistor is the V_{CE} value at which the current gain makes this sudden change.

We want to use phototransistors substantially above $V_{\mathrm{CE(sat)}}$, so that the current gain is roughly independent of the voltage, and the collector current that we can measure is proportional to the base current, which is the photocurrent proportional to the light input.

We can easily measure I-vs.-V plots at different light levels using PteroDAQ or the Analog Discovery 2, as shown in Figure 22.8 measured using the circuit in Figure 22.9. These plots can help us decide what voltage to use even more than $V_{\mathrm{CE(sat)}}$ does. For example, the LTR-4206 data sheet's value for $V_{\mathrm{CE(sat)}}$ at 0.5 mA is only 0.4 V, but at the high currents used in the colorimeter design, we probably want to keep $V_{\mathrm{CE}} > 0.6$ V.

Because of the current gain of the bipolar transistor, the current through a phototransistor is much larger than through a photodiode, often 100–1000 times larger. However, the

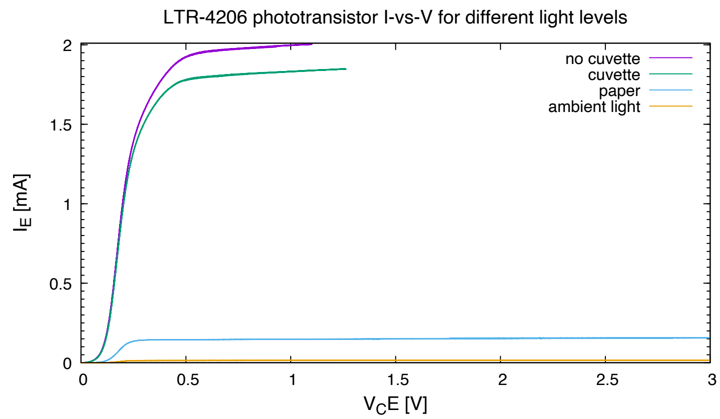

Figure 22.8: This plot shows the same data as Figure 22.7, but with a linear scale for current. It is much harder on this plot to see what is happening at low currents, but the slope of the current at high currents is now apparent.

Changing between log and linear scales for plots is commonly used to accentuate different features of the data—here the not-quite-constant nature of current at high current is apparent, while Figure 22.7 makes the low-current behavior clearer.

Data collected with PteroDAQ using a Teensy LC board [46]. Figure drawn with gnuplot [33].

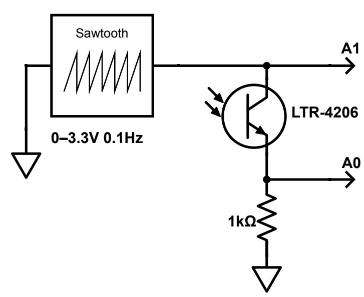

Figure 22.9: The plots in Figures 22.7–22.8 were created with this simple test fixture, except that a triangle wave was used instead of a saw-tooth wave. PteroDAQ was used to record the voltages at A0 and A1, V_{CE} *was computed as* $V_{\mathrm{A1}} - V_{\mathrm{A0}}$*, and the emitter current was* $I_{\mathrm{E}} = V_{\mathrm{A0}}/1\,k\Omega$*.*

Figure drawn with Digi-Key's Scheme-it [18].

high gain comes at a price—the capacitance of the base-collector junction is also multiplied by the gain (the *Miller Effect*), so phototransistors are about 1000 times slower also. For looking at low-frequency signals like in pulse monitoring, the slowness of the phototransistor is unimportant, and phototransistors are still faster than photoresistors.

A phototransistor also has a narrower dynamic range than a photodiode. The phototransistor has a fairly linear response over about 3 decades, while a photodiode is fairly linear over 7 or more decades. (That is, there is a range of about 1000:1 in irradiance over which the phototransistor can be well modeled as current proportional to irradiance, but there is about a 10,000,000:1 range for a photodiode.) The increased linearity makes photodiodes more popular for measuring instruments, even though the photocurrents are so much smaller.

Both photodiodes and phototransistors provide a current output that is proportional to the light power input, so their sensitivity can be expressed in A/W. The data sheet usually gives the current for bright light—often $1\,\mathrm{mW/(cm)^2}$, calling it the *on current*. If a photodetector collects all the light from a circle 3 mm in diameter ($7.07\,\mathrm{(mm)^2}$), then a light intensity of $1\,\mathrm{mW/(cm)^2}$ corresponds to a power input of about $71\,\mu\mathrm{W}$. The data sheet also gives the dark current (which is very small at room temperature).

A phototransistor may have a sensitivity of 2.8 A/W, while a photodiode in a similar package might have only 50 mA/W. The larger current is due to gain of the bipolar transistor in the phototransistor, ameliorated somewhat by different areas for the light-sensitive part of the devices that the lens of the package focuses on.

The current-vs.-voltage curves of Figure 22.8 show that the active current is not constant, particularly for larger current values. In order to get more linear measurements from a phototransistor, we need to hold V_{CE} constant and measure the current at that constant voltage. Chapter 23 explains a circuit that does just that: the transimpedance amplifier.

Exercise 22.2

Look up the data sheet for phototransistor SFH325FA.
What is the minimum photocurrent for group 4A with an illumination of $0.1\,\mathrm{mW/(cm)^2}$ at bias voltage of $V_{\mathrm{CE}} = 5\,\mathrm{V}$?
What bias voltage V_{CE} do you need across the phototransistor to get at least 0.3 times the photocurrent at 5 V? (Hint: look through the values in the specification table for things in units of voltage.)
What current would you expect with a 5 V bias, if the irradiance (brightness of light falling on the photodetector) is $2\,\mathrm{mW/(cm)^2}$?

22.5 Optical properties of blood

For Lab 6, you will design an optical pulse monitor to detect pulse by shining light through a finger and measuring how much light comes through. The amount of light changes as the amount of blood in the blood vessels changes—when the blood vessels have more blood, they block more of the light and the photocurrent drops.

Each time the heart beats, the surge in the flow causes red blood cells to accumulate where the blood vessels narrow, but in the time between pulses the blood cells redistribute

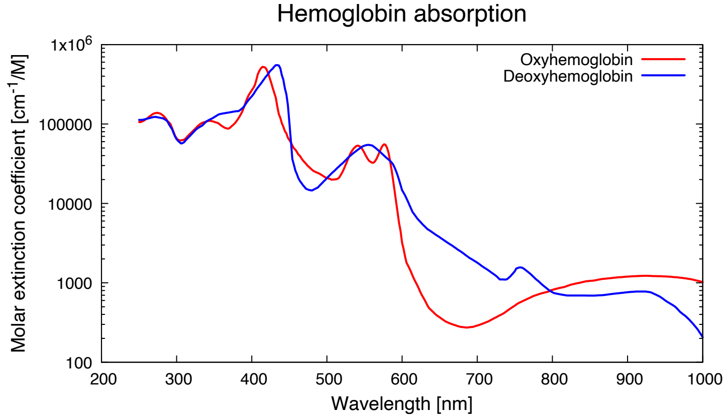

Figure 22.10: Oxyhemoglobin is most transparent around 686 nm (272.8 (cm)$^{-1}$/M*) and does not get much more opaque for* 700 nm ± 25 nm *(284–368* (cm)$^{-1}$/M*). At a shorter wavelength, which looks brighter to the eye,* 627 nm ± 22.5 nm*, absorption is higher (370–2130* (cm)$^{-1}$/M*). At the peak of our photodetector sensitivity,* 950 nm ± 27.5 nm*, oxyhemoglobin absorbs more (1136–1225* (cm)$^{-1}$/M*).*

Deoxyhemoglobin is not very important here, since oxyhemoglobin in the blood is usually well over 95% in healthy individuals.

Data for figure from https://omlc.org/spectra/hemoglobin/summary.html Figure drawn with gnuplot [33].

more uniformly. The fingers get rapidly more opaque as the pulse reaches them, then gradually get clearer again.

In this section, we'll look a little at the opacity of blood, to help choose what wavelength(s) to use for a pulse monitor.

Our phototransistor has its maximum sensitivity in the infrared portion of the spectrum (around a wavelength of 940 nm), which is typical for silicon photodiodes and phototransistors. But that doesn't mean that the best LED to use for illumination is an infrared one—we need to balance photodetector sensitivity with the opacity of blood.

> We want to monitor how much blood is in the finger based on how much light is absorbed, so we need to look at the light absorption of hemoglobin. For detection to be easy, a lot of the light needs to make it through the finger, so we need a wavelength at which flesh is not too opaque. As a good first approximation, we could look at the absorption spectrum of hemoglobin, the main coloring agent of blood—see Figure 22.10.

Because we don't want enormously bright lights, we'll want a wavelength at which there is only moderate absorption. A red LED at 627 nm (a common peak wavelength for red

LEDs) seems reasonable, but a different LED with a peak around 700 nm (like MT1403-RG-A) might be even better, as there would be less attenuation of the signal.

We can look up typical silicon photodetector sensitivity [151] and see that silicon photodetectors are only about 70% as sensitive at 627 nm as at their peak around 950 nm, and 80% at 700 nm. The three-fold greater transparency of blood at 700 nm more than compensates for the 20% reduction in sensitivity compared to using 950 nm.

If we were making a pulse oximeter, to measure the ratio of oxyhemoglobin and deoxyhemoglobin, then we'd need to use two wavelengths, choosing ones at which the molar extinction ratios were quite different. The minima of oxy/deoxy absorbance are around 654 nm, 438 nm, and 370 nm, while the maxima are somewhere greater than 1000 nm, 464 nm, 412 nm, 578 nm, and 538 nm. A good pair of wavelengths might be 650 nm (or 660 nm) and 940 nm, which are readily available LED wavelengths that have very different oxy/deoxy ratios.

Of course, hemoglobin is not the only substance in the body, so we should also look at the absorbances of other substances that might interfere with our signal: water, lipids, and melanin. The Wikipedia article *Near-infrared Window in Biological Tissue* [139] has absorption spectra for many of the relevant substances, though the spectra there have different ranges of wavelengths, with melanin's mainly in the visible range, so it may be necessary to look at other sources if you are interested in using infrared illumination.

A lot of wearable pulse monitors rely on light being scattered back to a photodetector next to the LED, not on opposite sides of a finger. The scattering of light by biological tissues decreases with wavelength, favoring short wavelengths for this application, but the opacity of melanin in the skin generally decreases with wavelength, so we need to make trade-offs. Reflection pulse monitors often use green LEDs (peak around 565 nm), where blood is much more opaque than in the red or infrared, and biological tissue scatters light back to the surface, but the melanin in skin does not yet block too much light.

When the Apple Watch first came out with a pulse monitor, it did not work well on people with dark skin or with wrist tattoos—a long-standing problem with reflection-based heart monitors [11]. Picking wavelengths at which melanin absorbs relatively little light compared to hemoglobin could reduce this problem, but one does have to check that the total amount of light scattered back to the sensor is large enough to provide a usable signal.

Tattoo pigments are a harder problem to design around, as they may be opaque at arbitrary wavelengths—either broad-spectrum illumination should be used or the user advised to put the sensor somewhere away from tattoos.

Researchers have developed fairly sophisticated optical models of skin and the flesh underneath it to model optical biosensors more accurately. One of these models has been released by Maxim Integrated, to support the use of their pulse oximeter and heart-rate sensors [67]. This model includes parameters for five different optical layers: stratum corneum, epidermis, papillary dermis, vascularized dermis, and subcutaneous adipose tissue. The model takes into account all the major sources of absorption and scattering, and traces the ray path of the light, so that either transmission or reflection modeling can be done.

An example is given of reflection modeling for someone with moderately light skin (1%–10% melanosomes, when the normal human range is given as 1%–40%). In that model, there is a minimum penetration depth around the wavelength of green LEDs, with the light penetrating about 1 mm, while for the wavelengths we've been considering (650 nm or more), the light penetrates about 3 mm. The choice of green LEDs for cis-illumination may be a deliberate choice to get light mainly from the vascularized dermis layer, where it will be modulated by the pulse, rather than from the underlying adipose layer.

25: Lab 6: Optical pulse monitor

Bench equipment: 1/8" (or 3 mm) drill for LEGO® bricks, jig or vise for holding bricks while drilling, function generator, oscilloscope (optional)

Student parts: 2–3 two-stud black LEGO® bricks, 3 2×4 black LEGO® bricks, phototransistor, rubber bands, black electrical tape, MCP6004 op amp, resistors, diode, capacitors, breadboard, PteroDAQ, LED (optional), heat-shrink tubing (optional)

25.1 Design goal

For this lab we'll be designing and building an optical pulse monitor to detect pulse by shining light through a finger and seeing the change in the opacity of the finger as the amount of blood in the blood vessels changes.

The first half of the lab consists of determining characteristics for a phototransistor and building a transimpedance amplifier to convert its current output to a voltage signal for the levels of light that pass through fingers.

The second half of the lab adds a second-stage amplifier and filter to make the signal have an appropriate voltage range to observe the pulse using the PteroDAQ system.

25.2 Design choices

In all the designs we'll be looking at, all the light arriving at the phototransistor has to pass through a finger (or other part of the body with blood pulses).

The first choice to make is to decide where the light to be detected by the pulse monitor is going to come from. There are several possibilities:

trans-illumination By mounting an LED on the opposite side of the finger from the phototransistor, you can shine light through the finger, resulting in a moderately uniform path length for the light. This approach is commonly used in fingertip pulse oximeters and used to be common in ear-clip optical pulse monitors.

ambient light You can use room light or sunlight to illuminate the finger. This has the advantage of not requiring an LED and requiring less mechanical design. It has the disadvantage of losing control over the wavelength of illumination and the intensity of the light. Furthermore, indoor illumination is often modulated (at 60 Hz, 120 Hz, or several kHz for fluorescent or LED bulbs) resulting in large AC signals in the light intensity that we might have to filter out.

cis-illumination Putting an LED next to the phototransistor is mechanically almost as simple as using ambient light, but relies on back-scattering of light from the finger, rather than on light passing all the way through the finger. The LED and the sensor can be placed on almost any exposed patch of skin where there are blood vessels near the surface. Thick layers of fat can cause problems, as fatty tissue has rather little blood flow, and the light can be scattered back to the sensor with little absorption from hemoglobin in the blood.

Ambient light is added to the LED light, so unless the LED is fairly bright or the finger is covered with opaque material, there can still be unwanted modulation of the light from ambient light sources. Making the LED very bright can cause problems with excessive heating of the skin, as all the power delivered to the LED gets dissipated as heat.

The cis-illumination design is probably the most popular now for optical pulse monitors, as the LED and phototransistor can be surface mounted on the same printed-circuit board, resulting in very low manufacturing cost, and bright LEDs can provide large signals for the phototransistor, reducing interference problems from picking up small stray 60 Hz currents.

pulsed light Both the trans-illumination and the cis-illumination designs do not really require continuous light. You can turn on the LED shortly before making the measurement, then turn it off again. Because the LED only needs to be on for 10 μs–100 μs (depending on the response speed of the phototransistor and amplifier) and the measurements only need to be made 30–60 times a second, power consumption can be greatly reduced by pulsing the LED.

For example, my Verily Life Sciences study watch has two green LEDs that it illuminates for 600 μs every 33.33 ms (a 30 Hz sampling rate). This short illumination reduces the average power needed by about a factor of 56, allowing fairly bright illumination without excess battery drain.

Furthermore, ambient light interference can be greatly reduced by measuring the light just before turning the LED off and just before turning it on, taking the difference between the readings as the measurement. This cancels out any ambient light modulation whose frequency is very small compared to the time of the LED pulse. Synchronizing the measurement with the light pulses is easy in a microcontroller program, but is a little trickier if one is using PteroDAQ or an oscilloscope.

Pulsed illumination provides a large fluctuation in light intensity that is not due to the blood pulse, so limits the amplification possible in the analog circuitry.

Look up the data sheets for the phototransistor (WP3DP3BT or LTR4206, for example) and LED (MT1403-RG-A, WP710A10F3C, WP710A10ID, or WP3A8HD, for example) in your parts kit.

Pre-lab 6.1

What current-limiting resistor would you need to place in series with your LED to limit the current to about 80% of the LED's maximum continuous current rating with a 3.3 V power supply? (Be sure to report which LED you are designing around!) Round the resistance to the E24 series and report both the forward voltage across the LED and what current you expect with the selected resistor.

25.3 Procedures

25.3.1 Try it and see: LEDs

Hook up your LED and series resistor, then power them from a 3.3 V supply. Does the LED light up? (If not, debug!) If you have an infrared LED, you may not be able to see it lighting up, but looking at it with a digital camera (including a cell phone camera) will probably show it quite visibly, as the blue sensors in the cameras are generally sensitive to near IR light. (Warning: the iPhone cameras supposedly have IR-blocking filters, so may not see the IR LEDs.)

Measure the voltage and current—remember that we always measure current by measuring the voltage across a known resistor, and you can use the current-limiting resistor as your known resistor. Does the LED have the forward voltage and current you computed it should have?

Bonus: you can use the Analog Discovery 2 to measure the I-vs.-V characteristics of the LED. Use a current-sense/current-limiting resistor that is large enough that the current through the LED will not exceed the absolute maximum, even with a 5 V supply. Hook up the function generator as the power supply for the LED and series resistor, and use the two differential oscilloscope channels to measure the voltages across the LED and across the resistor. Set the function generator to have a slow triangle wave (say 1 Hz) with an amplitude of 5 V and record the voltage waveforms. You can plot the current (the voltage across the resistor divided by the resistance) vs. the voltage to get a typical I-vs.-V curve for the LED.

25.3.2 Set up log amplifier

Wire your phototransistor and logarithmic current-to-voltage converter (see Figure 23.5) on a breadboard and measure the average DC voltage for your complete sensor in the ambient light of the lab. Shadow the phototransistor by pinching it between your fingers and measure the average DC voltage there. Try illuminating the phototransistor with the LED.

These measurements should give you an idea of the range of DC voltages (and, hence, photocurrents) to expect with the light levels you will be using.

Bonus: you can do the same sort of I-vs.-V plot for the phototransistor as you did for the LED, though you may have to use the offset on the waveform generator to avoid too large a negative voltage. Because the phototransistor current depends on the amount of light (that's the whole point of it!), doing an I-vs.-V plot only makes sense if you can be sure to keep the light level consistent throughout the measurement.

25.3.3 Extending leads

You will be putting the phototransistor (and LED, if you use an illuminated design) into a shroud that blocks all light except that which passes through the finger.

> Because the shroud would interfere with plugging the LED and photodetector into a breadboard, you will need to solder wires onto the leads of the components to lengthen the leads. The leads need to be long enough to reach the breadboard from a shroud a short distance away (so that you can rest your hand comfortably on the bench top with your finger over the phototransistor, without touching the breadboard). I recommend leads about 10 cm–20 cm long (4"–8").

If you have flexible (multi-stranded) wire, that can make positioning easier, but you will need to stiffen the free end by tinning it with solder, if you plan to put it into the breadboard. (A more durable solution is to use a crimp-on female header, but that requires a crimping tool, which is probably not worth the expense for just this course.)

Color coding the leads is very useful, so that you can connect the leads correctly, even when the LED and phototransistor are hidden inside the block.

> I recommend red and black for the anode and cathode of the LED, respectively, so that you can remember which one gets hooked up to the positive and which to the negative end. I recommend using different colors (say, green for collector and yellow for emitter) on the phototransistor, so that each wire coming from the block will have a unique easily deciphered color.
>
> Document the colors on your schematic!

Figure 25.1 shows how to extend the leads. Be sure to insulate the wires with electrical tape or heat-shrink tubing all the way to the LED or phototransistor, so that the wires don't accidentally short underneath the block—this shorting is hard to see and has been a very frequent failure mode in work by students trying to "save time" by doing a hasty, sloppy job. The time they lost in debugging was orders of magnitude larger than the time it would have taken to do a neat, thorough job of insulating the wires. The wires from the phototransistor should probably be twisted around each other, to minimize stray currents picked up by the wiring.

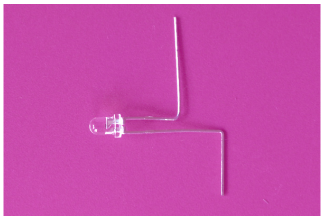

(a) Bend the leads of the phototransistor in different places, keeping the shorter lead shorter after bending to make it easier to keep track of which lead is which.

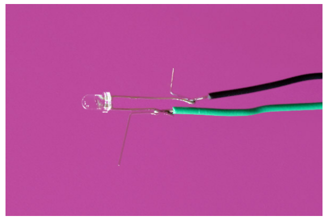

(b) Mechanically join the wires and the leads by bending the ends of the wires into little hooks with the long-nose pliers, linking them together, and crimping the hooks closed.

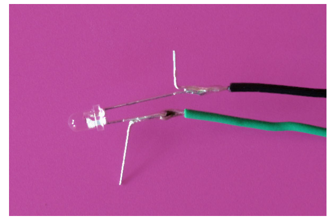

(c) Solder the joints. This can be done with the wires and phototransistor sitting on the bench top, but it helps to put the lead and the wire both into jaws of a helping hand so that they are held motionless.

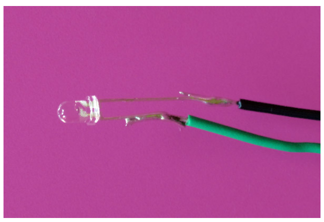

(d) Clip off the excess lead, to get a smooth connection.

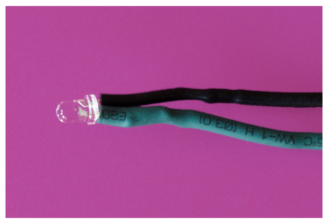

*(e) Use heat-shrink tubing or a little electrical tape around the leads and solder joints to keep the leads from shorting. Cover **all** the bare wire.*

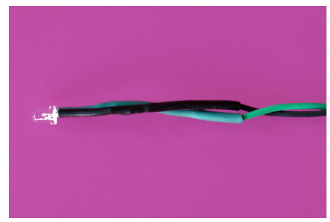

(f) Twist the pair of wires together, to reduce the pick up of stray signals.

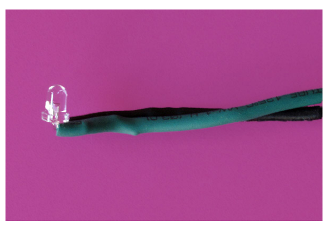

(g) Bend the phototransistor 90°, so that the wires can be taped flat to the back of a LEGO® brick.

Figure 25.1: The leads on the optoelectronic devices need to be lengthened to make connection to the breadboard easier. Use a color code for the wires so that you can tell which lead is which. I chose here to use black for whichever of the wires of the phototransistor had to be connected to the lower voltage. If you wire up an LED as well as a phototransistor, use different colors for them, so that you can tell the devices apart even when they are taped inside opaque holders.

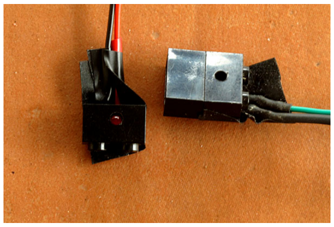

(a) Two-stud black LEGO® bricks with 1/8" holes can be used to hold 3 mm optoelectronics. An extra two-stud brick blocks stray light from entering the opening at the bottom of the phototransistor brick.

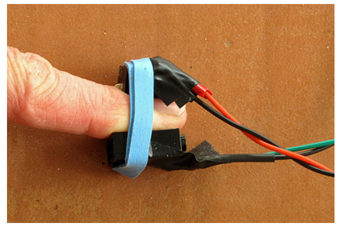

(b) A rubber band can be used with the LEGO® bricks to hold an LED and phototransistor on opposite sides of the finger (trans-illumination).

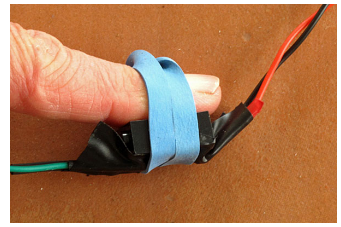

(c) If the LED brick is snapped onto the bottom of the phototransistor brick, a rubber band can be used to hold both on the same side for cis-illumination.

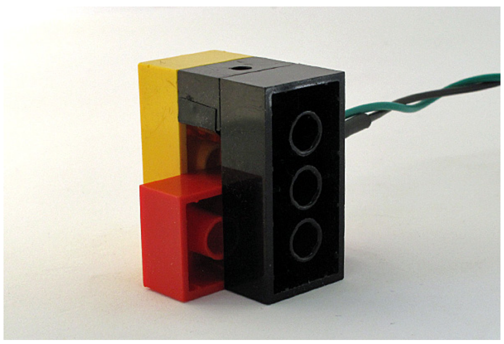

(d) Using some 2×4 LEGO® bricks to support the 1×2 brick makes the lump caused by the wires coming out the back of the phototransistor less of a problem. The brick closing off the bottom of the phototransistor brick should be black, to avoid light leaks. Similarly, the black electrical tape holding the phototransistor to the brick blocks light from reaching the back of the phototransistor.

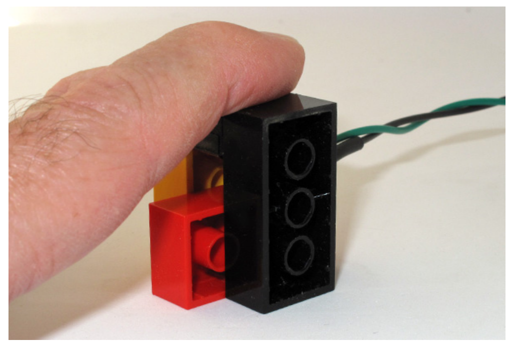

(e) A finger can be placed covering the phototransistor hole while the hand rests on the tabletop, to make keeping the pressure on the fingertip fairly easy.

Figure 25.2: The phototransistor brick can be left on the bench top and covered with a finger to use ambient light as the illumination, or strapped to a finger using rubber bands. By changing rubber bands or which finger is used, pressure can be adjusted to be between systolic and diastolic blood pressure, maximizing the signal.

25.3.4 Assembling the finger sensor

In early versions of this lab, I used wooden blocks for holding phototransistors, but I now prefer to use 1×2 black LEGO® bricks, with 1/8" or 3 mm holes drilled to accept 3 mm LEDs or phototransistors (one could also drill 13/64" holes to accept 5 mm optoelectronics parts), see Figure 25.2. Black LEGO® bricks are more opaque than wood, and it is easier to prototype finger clips with them than with the wooden blocks.

A simple block is not an ideal way to make a pulse sensor, as you need to provide a moderate pressure (between systolic and diastolic blood pressure) to get a really good pulse measurement. The correct pressure provides a slight throbbing in the finger—squeezing too hard cuts off blood flow to the finger tip and too little squeeze results in rather small fluctuations in opacity. If the pressure is too low, movement artifacts easily swamp out the small fluctuations due to pulse.

A gently squeezing clip is better than having to apply the pressure by pressing down with your finger, as the clip can provide a uniform pressure independent of slight movements. Crude clips can be cheaply prototyped with rubber bands and the LEGO® blocks.

Not everyone finds the rubber bands useful, though—a lot of students find it easier to train themselves to provide the correct pressure over a stationary block than to get the right combination of rubber bands for their finger width.

I prefer to have students drill the LEGO® bricks themselves, but the bricks can be predrilled if required. A jig to hold the bricks in the correct place in a drill press can save a lot of time on aligning the bricks for drilling, and is a fairly simple thing to make, as shown in Figure 25.3.

For any of the block designs here, the leads of the LED or phototransistor need to be bent at right angles close to the body of the device. It may be easiest to make those bends before inserting the device into the holder.

Cover the back of the phototransistor and wires with a piece of black electrical tape, both to keep them in place and to prevent stray light from reaching the detector from the back. A second black LEGO® brick snapped onto the bottom of the phototransistor brick blocks light from the bottom.

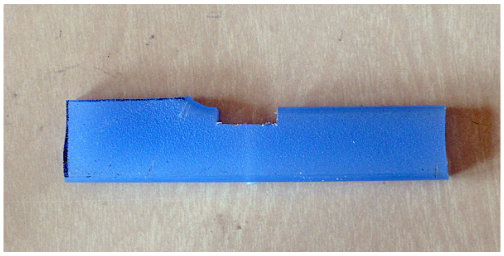

(a) This drill jig was cut out of an old, 1-cm-thick, HDPE cutting board using a scroll saw. The bottom is flat, and the top has a rectangular notch that is just the width of a 2-stud LEGO® brick (1.57 cm). The notch is only 4 mm deep, about half the thickness of the brick.

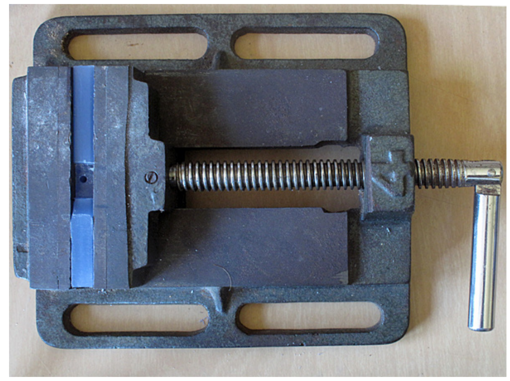

(b) The jig is the same width as the jaws of the drill-press vice, making alignment easy. In classroom use, it is worthwhile to use some painters' tape to tape the jig to the non-moving jaw of the vice, so that it doesn't move when students insert or remove the LEGO® brick.

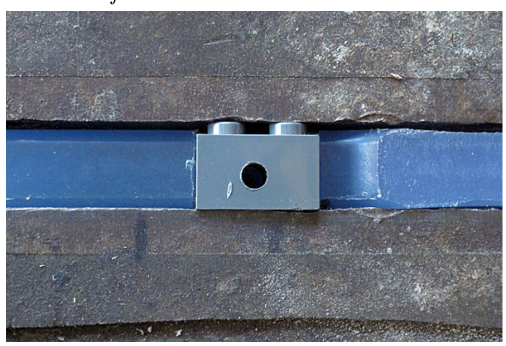

(c) The LEGO® brick is slightly taller than the thickness of the jig (11 mm vs. 10 mm), so the clamping force of the vice is entirely on the brick—the vice needs to be tightened gently, to avoid crushing the ABS plastic brick.

Figure 25.3: A jig for holding LEGO® bricks when drilling the 1/8" holes for the 3 mm phototransistor (and possibly for a 3 mm LED) makes the drilling much more repeatable. A grey brick is shown here, but black bricks should be used for their greater opacity.

25.3.5 Try it and see: low-gain pulse signal

After assembling the block, connect the phototransistor up to your logarithmic current-to-voltage converter (see Figure 23.5), rest your finger over the phototransistor, and try recording your pulse with PteroDAQ using a 600 Hz sampling rate. Sample at this high rate so that you can see how much 60 Hz or 120 Hz interference you are getting—you need to be sure that your amplifier will not start clipping as a result of the interference, which would

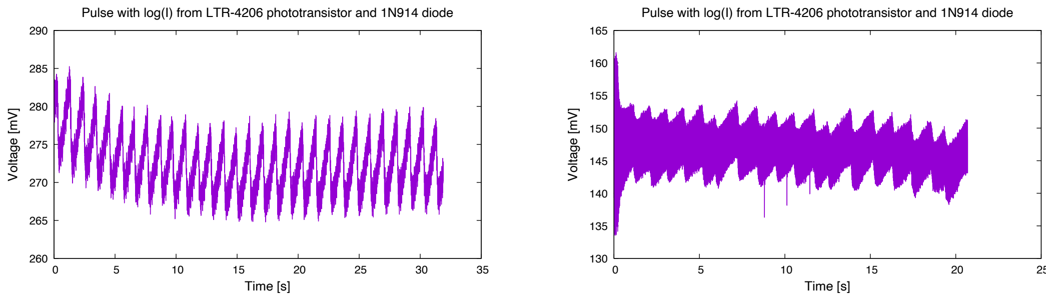

(a) With steady light from an LED desk lamp, the pulse signal is clearly visible as a saw-tooth waveform.

(b) With a fluorescent light, the pulse signal is almost buried in 120 Hz modulation of the light.

Figure 25.4: Signals from just the first stage of the pulse monitor. Note the large, but varying DC offset compared to the size of the pulse signal, as well as the size of the higher-frequency interference. The voltage being reported is the difference from V_{ref}*, so that zero current would appear as zero voltage here.*

Data collected with PteroDAQ using a Teensy LC board [46]. Figure drawn with gnuplot [33].

mask the pulse signal that you want to see. If you are using steady light (sunlight or LED illumination), you should get a signal something like the one in Figure 25.4a, but if you have fluorescent illumination, you will see a lot of interference at 120 Hz, as in Figure 25.4b.

If you get 60 Hz interference (as is likely), you should try to determine whether the interference is due to modulated light or to capacitive coupling. What simple experiments can you do to distinguish between possible sources of interference?

If the currents are very low, you may also have problems with picking up stray signals from the electrical and magnetic fields that pervade our buildings. These stray signals will mainly be 60 Hz signals (50 Hz in Europe), capacitively coupled to the wires from the phototransistor to the amplifier. You can reduce the problem by twisting the wires together, so that each wire picks up about the same amount of interference; by keeping the wires short; and by grounding yourself, since the finger-to-phototransistor capacitance is probably the largest capacitance contributing to the interference.

You can also minimize the effect of the interference by increasing the intensity of the light, as the small currents induced by capacitive coupling will be a smaller fraction of a large photocurrent than of a small photocurrent.

A source of interference that many people overlook is from the switching power supply powering a laptop—these provide a high-frequency fluctuation in the USB power-supply voltage to the PteroDAQ Teensy board, as well as electromagnetic fields from the inductors in the supply itself. When measuring small signals, it is often better to power the Teensy board from a laptop battery, without the switching power supply connected to the laptop.

If you sample at 60 Hz (or a divisor like 30 Hz), the line frequency interference will be hidden. It is still present in the signal, though, so you must either limit your gain so that even with the interference your signal is not clipped or filter out the interference before amplifying further.

The V_d signals plotted in Figure 25.4 are only 3 mV–10 mV peak-to-peak, which is too small to be seen on the PteroDAQ sparkline—you have to record them and then plot them with gnuplot.

The DC offset for V_d will vary with the brightness of the overall illuminations—moving your head or your body can shadow your finger and change the DC offset. The pulse signal is much smaller than the DC offset, and the pulse signal in Figure 25.4 is large compared with typical signals seen by students—it took me several attempts adjusting the pressure of my finger on the sensor to get such a large fluctuation in opacity.

It may be a little easier to see the fluctuations with the Analog Discovery 2—the resolution is worse than PteroDAQ on the Teensy board (335 μV per step instead of 50 μV), but the user interface makes it easier to zoom in on the signal.

Pre-lab 6.2

For each of the plots of V_d in Figure 25.4, what is the range of photocurrent measured? What is the peak-to-peak fluctuation in light passed through the finger due to the pulse? Report fluctuation in dB ($20 \log_{10}(A/B)$) or percent ($100 \frac{B-A}{(B+A)/2}$), where A and B are the lowest and highest light amounts in a pulse period.

You can maximize the pulse signal by gently squeezing your finger against the block—when the pressure on your finger is between the diastolic and systolic pressure, you will partially block the blood flow during diastole, increasing the difference in the amount of blood in different phases of the pulse. As with blood pressure measurement, the pulse signal is maximized when the pressure is close to the mean arterial pressure.

In addition to the 60 Hz interference, you may also see a large slow drift in the DC portion of the current. The second design project for this lab is to add filtering and more amplification, so that the 60 Hz interference and the DC drift are removed, while the signal for your pulse is amplified to be large enough to be easily visible.

The slow drift in the DC offset can be caused by changes in pressure of the finger, from changes in illumination (particularly if ambient light is used), or from charging of a high-pass DC-blocking filter.

25.3.6 Procedures for second stage

After determining how much additional gain you need, design and construct a second stage to your amplifier to provide this gain, including high-pass and low-pass filtering to remove DC offset and 60 Hz interference.

Pre-lab 6.3

Design a band-pass filter to pass pulses through, but block DC and 60 Hz. You want to have at least 20 dB attenuation of the 60 Hz interference (a factor of 10) relative to the peak gain, and you want to pass pulse signals as low as 20 bpm (0.33 Hz). The gain for the filter should be based on the observed signal levels out of your first stage.

As mentioned in Section 18.3, you may need more than one additional gain stage, to avoid having the DC offset of the op amps themselves result in clipping of the output.

Test this filter stage independent of the log-transimpedance amplifier. There are two tests worth performing:

- (Required) Use the waveform generator to create a sawtooth waveform (called "ramp up" or "ramp down" by WaveForms 3) at around 1 Hz, to mimic the waveform that should come out of the first stage. The opacity of the blood vessels spikes up sharply when the heart squeezes, then ramps back down as the blood seeps back out through the veins, so the light level drops sharply, then gradually recovers.

  Figure 25.5 shows a recorded pulse waveform illustrating how a ramp-down signal is a reasonable approximation to the opacity of the finger. In recordings for some people, but not this one, you can observe a "notch" in the downward slope about halfway down from the peak, called the *dicrotic notch* [17]. This dicrotic notch seems to be the result of the resistance of peripheral blood vessels, though a similar dip in pressure in the aorta (called the *incisura*) is due to closure of the aortic valve.

  Use the oscilloscope to observe the output of the filter stage. What do you observe about the shape of the output waveform? You can adjust the amplitude and offset of the waveform generator to investigate whether your filter is blocking DC correctly and has the right gain, but the gain computation is only accurate for sine waves—a ramp waveform has substantial higher frequency components which will have a different gain than the fundamental frequency. If you want to check your peak-gain computation, then use a sine wave at the frequency for which you should get peak gain.

  The best way to measure the gain of a filter is to measure both the input and the output of the filter, and take the ratio of the measurements. Use sine waves if you are testing a filter, as other waveforms have a combination of many different frequencies.

  You don't have to eyeball the measurements on a digital oscilloscope, as it can compute and report the amplitude for you. On the Analog Discovery 2 scope, you can choose "Measurements" from the "View" menu, then use the "Add" button to add "Defined Measurement", "Channel 1", "Vertical", "Amplitude" and the same measurement for "Channel 2". The ratio of the measurements is the gain of the filter for the measured input signal.

- (Bonus) The gain measurement can be fully automated, to get a measured Bode plot for the filter. Connect waveform generator W1 to the input, use the first oscilloscope channel to observe that input, and use the second oscilloscope channel to observe the output. Now use the network analyzer to plot the response of the filter over the frequency range of interest. (Warning: the network analyzer can be quite slow at low

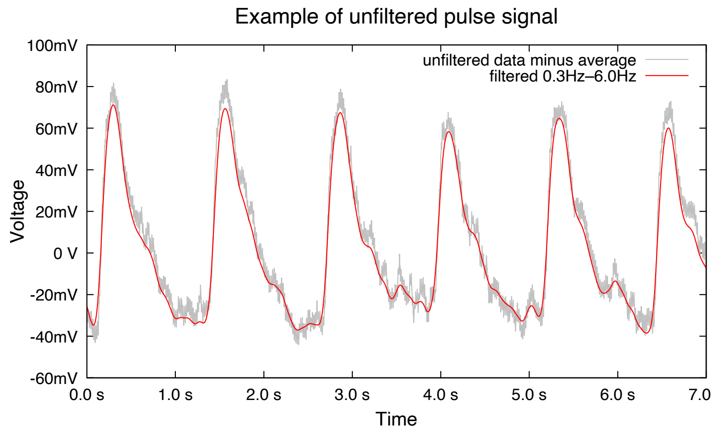

Figure 25.5: The unfiltered pulse waveform was taken in full sunlight with the patient grounded to minimize 60 Hz interference from capacitive coupling. You will probably not be able to get so little 60 Hz interference in the electrically noisy environment of the lab. The average value of the signal was subtracted off, to center the signal at 0 V.

Digitally filtering the waveform removes the remaining 60 Hz interference, but rounds the peaks a little. The underlying opacity waveform can be thought of as ramp-down sawtooth waveform, with an initial jump upwards in opacity followed by a steady decline.

frequencies, so you might want to do some quicker checks at the high end of the range of interest first.)

If you are working indoors, you will almost certainly see 60 Hz interference in your pulse waveform (assuming you sample at a high enough frequency to see it, and don't alias the 60 Hz and its harmonics to DC). This interference could come from any of several sources: capacitive coupling between the amplifier input and a conductor with 60 Hz through an electric field, inductive coupling through a magnetic field, or modulation of the light that is illuminating the finger.

Think up some ways that can distinguish between the different sources of interference: changing light sources, grounding the other plate of the capacitor, unplugging the laptop power supply, ... Do some experiments to see what source of interference seems to be the most troublesome and suggest ways to reduce the problem that could be used if you were making a commercial pulse monitor.

Record the pulse output for 20 s with the PteroDAQ data-acquisition system.

Do digital filtering (using the filter program from Lab 5) to plot a further cleaned-up pulse waveform.

25.4 Demo and write-up

For the first half of the lab, show me a single-stage amplifier that produces a DC voltage in a reasonable range when a finger is in the block, and a (small) fluctuation related to the pulse on the oscilloscope or PteroDAQ.

For the second half of the lab, show me a clear PteroDAQ trace of a heartbeat from a finger in the sensor. (That means recording using PteroDAQ and plotting with a program like gnuplot, not a screenshot of PteroDAQ!)

You will probably encounter some problems with the mechanical design of the finger-tip sensor. Describe changes you would make to improve the mechanical design. Incidentally, please refer to the LEGO® bricks as "bricks", not "legos". The Lego company has strong feelings on the matter:

> *If the LEGO trademark is used at all, it should always be used as an adjective, not as a noun. For example, say "MODELS BUILT OF LEGO BRICKS". Never say "MODELS BUILT OF LEGOs".* [65]

> The report should contain the overall gain of the whole system from light in to voltage out. The gain for your system should be reported in mV/dB.

The write-up should include schematics for both the single-stage and multi-stage amplifier, calculations used to get initial component values, adjustments to components that were needed to make the circuit work properly, and a plot of the output from the final amplifier design.

> In your schematics, make sure that you use the proper NPN phototransistor symbol—incorrectly using a PNP phototransistor symbol is a serious failure.

Use the amplifier output file to compute the pulse rate for the individual.

> Compute the pulse rate by measuring the time t precisely for exactly N periods and dividing the number of periods by the time, N/t. Because the interval between heart beats can fluctuate a lot from one pulse to the next, make the number of periods fairly large (say, $10 \le N \le 60$). *Don't get confused by fencepost errors* (see page 329).

Because time is one of the easiest things to measure precisely and accurately, you get much more precise results by taking an integer number of periods and measuring their time than by taking a time interval and estimating the number of periods.

Fluctuations in the period of the pulse can be quite large, and the period of one pulse is not independent of adjacent pulses—Figure 25.6 plots 202 pairs of pulse periods (about 4 minutes of recorded pulse). The mean period is 1.184 s (for a pulse rate of 50.66 bpm), and the standard deviation of the pulse period is 0.0928 s, but the distribution does not look Gaussian, as there are more very short periods than very long periods (skewness −0.8742 s).

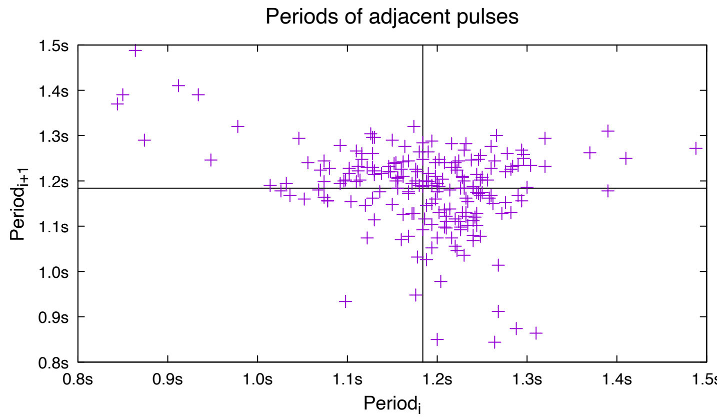

Figure 25.6: Pulse periods have a wide range and are not independent of each other. The mean period is 1.184 s (for a pulse rate of 50.66 bpm), shown by the horizontal and vertical lines on the plot. Short periods are usually followed by long periods in this data—the Pearson's r correlation between adjacent periods is −0.3068.

Figure drawn with gnuplot [33].

The range of periods is from 0.84 s to 1.49 s, but short periods are usually followed by long periods.

Averaging pulse periods (taking the time for N periods and dividing by N) results in lower variance for the estimate of the period. There is a trade-off between averaging many periods (to get a good estimate of the average pulse rate) and averaging few periods (to track changes in pulse rate as they occur). An exercise monitor may want to use a shorter window for the averaging than a pulse monitor for resting pulse rate.

41: Electrocardiograms (EKGs)

41.1 EKG basics

An electrocardiogram (abbreviated ECG or EKG—both are acceptable in English, with ECG being more common[1]) measures small differential voltages using electrodes on the patient's skin to view the electrical behavior of the heart muscles. For EKG purposes, the heart is modeled as an electrical dipole that changes its orientation and magnitude as waves of depolarization pass through the muscles of the heart.

The single-channel EKG usually uses 3 electrodes: the pair of electrodes that carry the differential signal and an extra body electrode that provides a reference voltage to keep the differential pair between the power rails of the instrumentation amp. The body electrode is an *output* from our circuitry to set the common-mode voltage for the differential *input* electrodes. For our purposes, it suffices to generate a V_{ref} voltage with a unity-gain buffer, as we've done for other amplifiers, for the body electrode.

The medically standard 12-lead EKG has ten electrodes: one on each ankle and wrist and six horizontally across the chest [90]. The designation "12-lead" means that there are 12 differential channels being recorded, not that there are 12 wires—six of the channels are in the vertical plane, based on the four limb electrodes, the other six are from the six chest electrodes. The standard placement requires that the patient be lying down, with arms and legs relaxed to avoid picking up electromyogram (EMG) signals from skeletal muscles in the arms and legs.

For monitoring people moving around (as in a stress EKG), the limb leads are moved in—the wrist electrodes are moved to the chest near the shoulders and the ankle electrodes are moved either to the abdomen just above the pelvic bone or to the lower back (to reduce EMG interference from the hip and thigh muscles).

In Lab 13, you won't be connecting the electrodes to your wrists and ankles, as keeping your arms relaxed is very difficult to do when you are adjusting the circuit, turning recording on or off, or typing up results. Instead, we'll use the electrode locations commonly used for emergency EKGs: in the hollows just below the collar bones avoiding the pectoral muscles and shoulder muscles, and on the abdomen just above and inside the left iliac crest (the bone at the top of the pelvis). You can then choose two of the electrodes as the signal electrodes (for leads I, II, or III) and use the third one as a body electrode, as shown in Figure 41.1. With this placement of the electrodes, only the pectoral muscles are likely to provide interference with the EKG signal, and they are fairly easy to keep relaxed when you are sitting.

[1] I use EKG, the older acronym that comes from the German name, in part because my first electrocardiogram was ordered by a Prussian pediatrician.

(a) Lead I is LA–RA. *(b) Lead II is LL–RA.* *(c) Lead III is LL–LA.*

Figure 41.1: This figure shows the approximate electrode placement for the three electrodes. The two "arm" electrodes go just below the collarbone, trying to avoid any muscle that might add extraneous signals. The "left-leg" electrode goes just above the iliac crest (the top of the pelvic bone). If you are just recording Lead I (LA − RA), then the placement of the LL electrode is not very important—almost anywhere on the body will do. The LL electrode is important if you want to record Lead II (LL − RA) or Lead III (LL − LA).

> We will be looking at lead I (left arm minus right arm, often abbreviated LA−RA) or lead II (left leg minus right arm, LL−RA) which have the classic EKG signal shown in Figure 41.2. The 200 ms by 0.5 mV grid of Figure 41.2 is standard for EKG traces and is usually printed as a 5 mm by 5 mm square. In fact, many of the guides for interpreting EKGs refer to the voltage and time both in *cm*, assuming that doctors and nurses can't deal with millivolts and milliseconds. The standard scaling is 1 mV/cm and 400 ms/cm, which is probably not the scaling here.

The vertical spike of the large R part of the wave is typically around 0.5 mV–1 mV, depending mainly on the quality of the electrode-skin contact. The initial P wave corresponds to depolarization of the atrium, which starts the push of blood from the atria to the ventricles. The R spike corresponds to the depolarization of the ventricles, which begins the squeeze that pushes blood out of the heart. The T wave corresponds to repolarization of the ventricles. See Ashley and Niebauer's explanation for more details of the physiological events the various parts of the ECG correspond to [6].

We don't want to DC-couple the whole amplifier, as the ±1 mV signal we're interested in is often accompanied by a ±300 mV DC offset, due to different half-cell potentials at the skin-electrode contact. The DC blocking is usually done after the first stage of amplification by the instrumentation amp, and the initial gain is limited to avoid saturating the instrumentation amplifier with the DC offset.

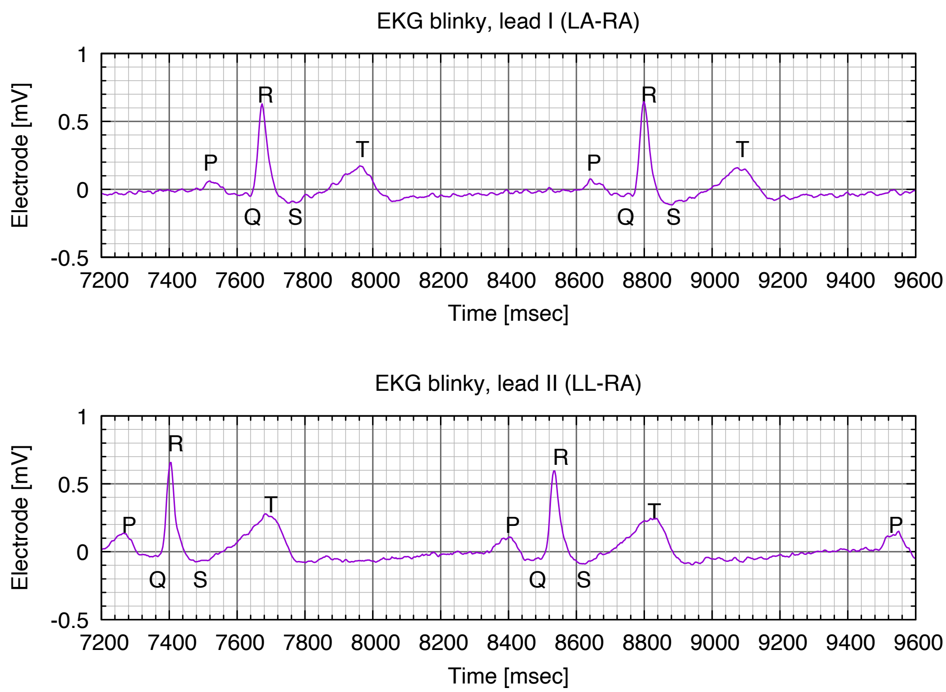

Figure 41.2: An EKG trace of Kevin Karplus's heartbeat. Despite the same time scales, these traces were recorded at different times and correspond to different heartbeats. The period in both the traces is about 1125 ms for a pulse rate of 0.89 Hz or 53.3 beats/minute. The gnuplot commands to create a grid like that seen here are given in Figure 41.3.

Data collected with PteroDAQ using a Teensy LC board [46]. Figure drawn with gnuplot [33].

Cardiac monitors usually have a passband of 0.5 Hz to 40 Hz or even narrower (1 Hz–30 Hz), and diagnostic 12-lead EKGs usually have a wider range (0.05 Hz–100 Hz or 0.05 Hz–150 Hz) [14, 86].

The wider frequency response provides a detailed look at the heart signals and is necessary for some sorts of diagnosis (such as elevated ST response), but the narrower frequency response of cardiac monitors filters out high-frequency noise from skeletal muscles and low-frequency artifacts from breathing and other motions, and so is more suitable when looking just at cardiac rhythm or cardiac rate.

Capacitive coupling of ambient electrical interference to the signal wires can be reduced by twisting the wires together and surrounding the twisted pair with a conductive shield, biased to the common-mode voltage of the wires. This common-mode bias can be generated from the instrumentation amplifier itself, by splitting the R_{gain} resistor in two resistors, each

```
set margins 9,6,1,1

set xtics 200
set mxtics 5
set ytics 0.5
set mytics 5
set style line 1 linetype 1 linecolor rgb 'grey40' lw 1
set style line 2 linetype 1 linecolor rgb 'grey70' lw 0.5
set grid  xtics ytics mxtics mytics ls 1, ls 2

# depends on xrange & yrange: number of boxes high/boxes wide
set size ratio 3./12.
```

Figure 41.3: These are the gnuplot commands used to set up the grid and change the shape of the plots for Figure 41.2. The "size" command does not guarantee that the grid will come out square—you'll have to tweak the size or the xrange and yrange to make the grid square.

half of R_{gain}, and using a unity-gain buffer whose input is connected to the middle of the R_{gain} resistance. This shielding is not absolutely needed for an EKG—we do unshielded leads in Lab 13, but if you are trying to amplify much smaller signals, such as from an electroencephalogram (EEG), then shielding becomes more important.

Common-mode noise can also be reduced by using an active bias circuit, instead of just biasing the body electrode to V_{ref}. Supposedly, feeding back the common-mode voltage with a gain of about -40 as the body bias will reduce the interference substantially, but I have not had much luck with this technique. Although the active bias circuit did indeed reduce the common-mode noise substantially, the signal on the bias wire was capacitively coupled back into the EKG signal leads, creating an interference signal that appeared as a differential voltage, not just a common-mode voltage. The result was a much, much noisier output signal than when I just used V_{ref} on the bias electrode. For the technique to work well, it may be necessary to shield the EKG signal pair (with the shield set to the common-mode voltage) and run the bias wire outside the shield, to minimize coupling to the signal pair.

41.2 Safety

Medical equipment, such as an EKG machine, is required to be isolated from the line voltage, to avoid the dangers of electric shock. Blood and muscle are good conductors of electricity, but skin is not (think of your body as a bag of salt water—the bag is not conductive, but everything inside is). Cuts in the skin eliminate the resistance of the skin, making the body much more conductive.

AC voltages are more dangerous than DC ones, in part because of capacitive coupling through the skin, and in part because muscles and nerves respond more to varying voltages than to constant ones. It only takes a small current through the heart to cause problems: about 30 mA at 60 Hz or 300 mA–500 mA DC can cause fibrillation [124]. About 10 mA at 60 Hz through the arm can cause violent muscular contractions that make it impossible for the person to release a wire.

To prevent such problems, it is common to put a large resistor (like 1 MΩ) between EKG electrodes and the EKG amplifier.

All electrode leads should be connected through resistors to make sure that any current from the EKG to the body stays under 50 μA (the limit according to the ANSI/AAMI ES1-1993 standard, as reported by Company-Bosch and Hartmann [14]), even if the wires are accidentally connected to the highest voltages in the EKG circuitry. The American Heart association calls for a stricter limit, of a maximum of 10 μA "between any patient electrode and either power line ground or the accessible part of the electrocardiograph . . . even in the presence of a single fault" [62].

Because the instrumentation amplifier for the EKG has a very high input impedance, this extra resistance in the signal leads does not result in noticeable changes in the voltage observed.

Adding 1 MΩ to the lead that biases the body, however, could cause problems with not providing enough current to keep the electrode voltages within the range needed by the instrumentation amplifier. The body and signal leads can have fairly large capacitive coupling to the 60 Hz ambient electrical fields, and the resulting voltage may be large enough to make the common-mode signal fall outside the range your instrumentation amplifier can handle.

Think of your body as the output node of a voltage divider, with a capacitor to a 60 Hz source and the bias-lead resistor to virtual ground—if the capacitance and resistance are both large, then the 60 Hz voltage on the output will be large. Reducing the resistance reduces the signal on your body.

Reducing the resistance of the bias lead can result in safety concerns, particularly if there is a path through your power supply to the AC power lines. The pre-lab exercise in Lab 13 has you compute the lowest resistance you can get away with. You can reduce the hazard by powering your EKG circuit from a battery (such as a laptop battery) with no connection to AC power. Battery operation also reduces power-supply noise from being injected into your electronics.

To protect circuitry or patients from higher voltages, it is common to add protection diodes to a circuit, as shown in Figure 41.4.

Two of the circuits use ordinary diodes to limit the voltage range, either to the power-supply range or to a narrow range around a reference voltage. The other two use special *Zener* diodes that conduct in the reverse of the normal direction when the back bias exceeds a well-defined *Zener* voltage.

It is common for medical equipment to use optoisolators to separate the parts connected to the patient from the parts connected to the computers and power lines. An *optoisolator* is a paired LED and phototransistor (or photodiode) that communicates information via photons, rather than electrons. They provide excellent protection against stray current paths. Some optoisolators can be used for transferring analog signals, but many have a nonlinear response optimized for digital signals.

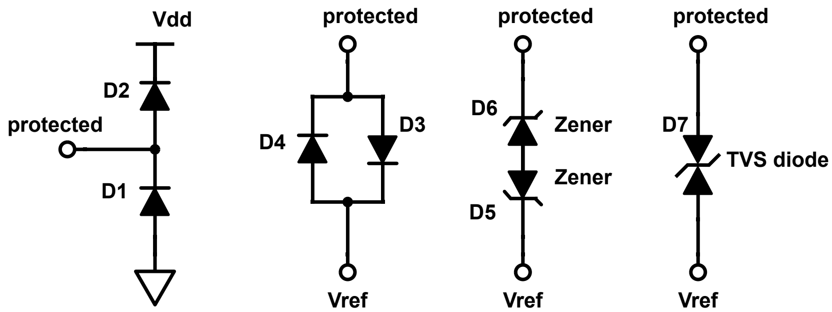

Figure 41.4: Four diode-clamp circuits for protecting a node from excessive voltages. D_1 and D_2 limit the voltage swing to be from a little below ground to a little above $V_{\rm dd}$, with the amount outside the voltage range determined by the current and the diode characteristics. This sort of protection is very common for inputs to integrated circuits.

D_3 and D_4 limit the voltage swing to a narrow window around $V_{\rm ref}$, with the amount of swing determined by the current and diode characteristics. For the tiny currents of EKGs, this clamping method may eliminate the desired signal.

D_5 and D_6 limit the voltage to a wide window around $V_{\rm ref}$ ($V_{\rm ref} \pm V_{\rm Zener}$) by using a pair of Zener diodes.

D_7 is roughly equivalent to the circuit using D_5 and D_6, but using a single transient-voltage-suppression *diode designed for clamping voltages in a specific range.*

Figure drawn with Digi-Key's Scheme-it [18].

Unfortunately, optoisolators are not good for providing power to the electronics on the patient side of the isolation barrier. For that, there are two main solutions: batteries and isolated power supplies. Battery power is a cheap solution, but has reliability issues—batteries always seem to fail when they are needed most, and replacing batteries is both expensive and a nuisance. Isolated power supplies use small transformers to transfer power through magnetic fields, rather than through direct conduction of electrons. Isolated power supplies are slightly more expensive than other power supply designs (used for electronics that people do not come in contact with). Power supplies designed and tested to meet medical standards usually cost about twice as much as other power supplies with similar ratings, but even cheap wall warts usually meet minimal isolation resistance standards.

41.3 Action potentials

Where does the signal that we're seeing on an EKG come from? There are two ways to answer this: one based on what happens in individual cells and one based on the waves of activity in the heart.

> We can look at a cell membrane as a capacitor: the lipid bilayer is a good insulator, allowing no ions to pass it, and both the cytosol inside the cell and the extracellular fluid are good conductors, containing a lot of salt. Ion pumps can move charge from one side of the membrane to the other, charging the capacitor, and ion channels can let ions flow back, discharging the capacitor.

In resting states, cardiac muscle cells pump sodium and calcium ions out of the cell and potassium ions into the cell, resulting in a positive charge on the outside relative to the inside of about 90 mV. When sodium channels are opened, the sodium ions rush in, depolarizing the membrane (discharging the capacitor). This individual-cell activity is referred to as an *action potential*.

The sodium channels then close and the cell enters a plateau phase, where potassium ions flow out and calcium ions flow in. It is during this phase that the cardiac muscles contract, triggered by the calcium. The sodium action potential lasts less than a millisecond, but the calcium action potential can last for hundreds of milliseconds [108].

Then the calcium channels close and the cell rapidly repolarizes as the potassium continues to flow out. The ion pumps then gradually restore the initial balance of ions, with an excess of sodium and calcium outside the cell and an excess of potassium inside the cell. You can find a more detailed explanation of the entire process in the Wikipedia article on cardiac action potentials [114].

The action potentials refer to voltages across a cell membrane, but our EKG electrodes have no access to the inside of cells. When a cell depolarizes, the sudden influx of positively charged sodium ions increases the charge inside the cell, but decreases the charge outside the cell. As the depolarization wave travels across the heart, the tissue in front of the wave is positively charged relative to the tissue behind the wave.

> The dipole moment between the tissue in front of and behind the wave is what the EKG is detecting, so a depolarization wave traveling towards the + electrode and away from the − electrode will result in an upward (positive) spike in the EKG waveform.

The large ventricular wave traveling down the septum between the ventricles produces the R spike—the septum is usually oriented somewhat to the left and downwards, resulting in a positive R spike on leads I and II. The repolarization wave has the opposite polarity, but travels in the opposite direction, so it also results in a positive spike (the T spike). One of the best explanations of heart dipole and repolarization waves I've seen (including an animation of the depolarization wave) is on Natalie's Casebook [72].

The online *Textbook of Cardiology* has a very detailed explanation of the relationship between action potentials and EKG waveforms on the "Cardiac Arrhythmias" page [61]. A more basic explanation can be found on the Wikipedia EKG page [126], but the explanation by Ashley and Niebauer may be clearer and more informative [6].

42: Lab 13: EKG

Bench equipment: oscilloscope, soldering station, fume extractor

Student parts: MCP6004 op amp, resistors, capacitors, breadboard, op-amp protoboard, screw terminals, EKG electrodes, clip leads, PteroDAQ

42.1 Design goal

For this lab you will design and solder a one-channel electrocardiogram circuit, not using a prebuilt instrumentation amplifier, but making your own instrumentation amplifier out of op amps.

> This is not a medical-grade EKG, as it does not include electrical isolation, electrostatic protection for the electronics, or calibration signals. It is also only a single channel, not the 12 channels of a standard EKG.

You don't get to use the INA126P instrumentation-amp chip for this lab. You have to design and construct your own instrumentation amplifier out of MCP6004 op amps. Because you aren't using the INA126P, you don't have its output range limitations, but the much more relaxed ones of the MCP6004 chips.

42.2 Pre-lab assignment

All electrode leads should be connected through resistors to make sure that any current from the EKG to the body stays under $50\,\mu$A (the limit according to the ANSI/AAMI ES1-1993 standard, as reported by Company-Bosch and Hartmann [14]), even if the wires are accidentally connected to the highest voltages in the EKG circuitry. The American Heart association calls for a stricter limit, of a maximum of $10\,\mu$A "between any patient electrode and either power line ground or the accessible part of the electrocardiograph . . . even in the presence of a single fault" [62].

Pre-lab 13.1

Choose the size for your current-limiting resistors on the EKG electrode wires. What is the smallest resistance you can use on the reference electrode, to ensure that the ANSI/AAMI limits are met? (The signal wires can use much higher resistances if you wish.)

In Lab 12, we did not have to worry about the half-cell potential of our electrodes, since the two electrodes were very similar and shared a common electrolyte, so the half-cell potentials cancelled. With EKG electrodes, each electrode has its own electrolyte gel, but the concentration of electrolytes in the gel will vary depending on such things as how sweaty the skin is there and how salty the sweat is. Because of this difference in electrolytes between electrodes, EKG electrodes can have as much as a 200 mV–300 mV difference in their half-cell potentials.

> That means that the DC offset of your input signal could be anywhere from −300 mV to +300 mV, and this offset is in the differential voltage (not just the common-mode voltage). Your instrumentation amplifier will amplify that DC offset, and so you can't have a large gain in your first stage.

Some designs have high-pass filtering before the first stage, but those designs are difficult to get right, because of the high impedance of the voltage source—filtering after first-stage amplification is usually easier.

If you use a high-pass filter to block the DC offset, you need to consider what cutoff frequency it should use. If the cutoff frequency is too high, you can lose part of your signal, but if the cutoff frequency is too low, then the EKG will take a long time to recover from any sudden changes in the DC offset. Such sudden changes in DC offset are common in EKG signals as a result of movement (sometimes called *motion artifacts*).

As a rough rule of thumb, it takes about 5 RC time constants for a simple RC filter to recover from a sudden change, so the lower bound for your corner frequency is set by how long you are willing to tolerate a loss of signal due to a motion artifact.

Because your first-stage instrumentation amplifier can't have a large gain, you will need a second-stage amplifier, after eliminating the amplified DC offset.

Pre-lab 13.2

Figure out what gain you need to achieve to record your EKG signal with PteroDAQ. Your input signals will probably be up to 2 mV peak-to-peak, but may have a 300 mV DC offset (in either direction). The DC offset is an offset to the differential signal, not just to the common mode.

Split the gain into different stages, making sure that the DC offset won't saturate any stage.

You will be using your own amplifier design made from MCP6004 op amps, not the INA126P, so the limitations of the INA126P amplifier are irrelevant to this design problem.

You will most likely also need a low-pass filter in the second stage, to prevent clipping due to amplifying 60 Hz interference. A low-pass filter as an anti-aliasing filter is a good idea in any case—choose your corner frequency based on your sampling rate.

Pre-lab 13.3

Draw a block diagram for your EKG, complete with signal levels.

Because you need a second-stage amplifier, you probably can't use a 3-op-amp instrumentation amp, as you'll run out of op amps (you only get 4 on the protoboard). Your instrumentation amplifier should use resistors you have—the 40 kΩ used internally by the INA126P is not a standard value in any of the common resistor series (see Table 4.2), so don't use that value!

Pre-lab 13.4

Design your own instrumentation amplifier using op amps. Consider the design in Section 40.2, which uses only two op amps, allowing you to have two left for creating V_{ref} or second-stage amplification.

Pre-lab 13.5

Have a complete, clear schematic (preferably drawn with a tool like Scheme-it [18]) before coming to the lab.

You need to have three electrodes connected in your complete schematic, as described in Section 41.1, but many students in the past only showed two in their initial design.

Pre-lab 13.6

Plot the gain of your amplifier as a function of frequency, taking into account any low-pass or high-pass filtering in the amplifier. What is the gain at 1 Hz? at 60 Hz?

Think about the layout of your printed-circuit board—don't just throw wires and components around randomly. For example, all the electrode wires should be together, so that they can be twisted into a bundle to reduce 60 Hz pickup. Similarly, the power wires and the output wires that communicate with the microcontroller board should also all be together. There are two places on the printed-circuit board for screw terminals (or right-angle male header pins)—how will you use them effectively?

How will you assign the op amps in your schematic to the op amps in the layout? Although the four op-amp positions on the MCP6004 chip are electrically interchangeable, your wiring will be much simpler and shorter with some assignments than with others (if you use all four op amps, there are $4! = 24$ different possible assignments). Similarly choosing the placement of your resistors can lead to either short, clean wiring or messy tangles.

Pre-lab 13.7

Design a layout for the circuit on the op-amp protoboard worksheet [50].

42.3 Procedures

You can get an idea of the signal you want to amplify by recording an EKG signal without any amplification—connect the two sense electrodes to a differential oscilloscope channel of the Analog Discovery 2 and the bias electrode to ground. (Or connect the sense electrodes to the differential inputs of PteroDAQ and use a couple of resistors (100 kΩ to 220 kΩ) to make a voltage divider for the bias electrode.)

The small signal should be just barely visible, though it may be buried in 60 Hz interference and quantization noise. After recording the signal with 1200 Hz sampling, I was able to recover the EKG signal with digital filtering (low-pass at 100 Hz and notch filter rejecting 58 Hz to 62 Hz) [49]. I don't recommend trying to use just the least-significant bits of an analog-to-digital converter like this, but it is amazing how much information can be recovered by digital filtering.

> Make a breadboard version of your EKG circuit, to make sure that everything works as designed.

You can test the low-gain first stage of your EKG with the Analog Discovery 2, but to test the high-gain second stage or both stages together, you may want to make a small test fixture. The function generators are not good at producing very small signals, but you can use a voltage divider to reduce the output of the function generator. If you want to be able to control both the differential DC offset and the common-mode, you can use two function generator channels to drive the voltage divider, as shown in Figure 42.1. Make sure that the two channels are synchronized and 180° out of phase.

> Disconnect your voltage divider from the instrumentation amplifier when using the EKG electrodes—the 1 kΩ resistor is such a low impedance that it would effectively short-circuit the electrodes!

If you want to use the network analyzer with the test fixture of Figure 42.1, you run into a small problem—the network analyzer of the Analog Discovery 2 controls only Wavegen 1 not Wavegen 2. You can set Wavegen 2 to be a constant voltage for the network analyzer, but then the common-mode voltage varies as well as the differential voltage (the common-mode voltage only varies by half as much as the differential voltage, but you can't tell whether you are amplifying the differential signal or the common-mode signal).

If you connect both inputs of your instrumentation amplifier to Wavegen 1, you can measure the common-mode rejection of your amplifier design. There is no need to produce very small input signals for this test, as your amplifier should have gain much less than 1 for common-mode signals (the INA126PA has a common-mode rejection of at least 74 dB, and your home-brew amplifier should be able to get a rejection of at least 33 dB).

> Make a wiring harness by twisting three color-coded wires together and attaching an alligator clip to each. You need to twist the wires together to reduce AC pickup by the leads. Wires should be about 100 cm (3') long with the last 20 cm (8") before the alligator clips not twisted together, so that the leads can spread out to the electrodes.

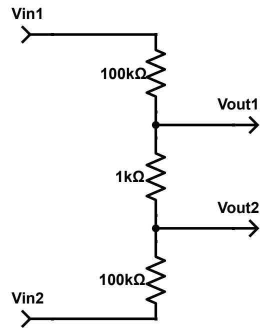

Figure 42.1: You can make a dummy EKG signal that has only a few millivolts of differential signal with an arbitrary common mode.

With the Analog Discovery 2, you can use two synchronized function generators, with $V_{\text{in}1}$ *connected to one and* $V_{\text{in}2}$ *connected to the other. If the amplitudes and offsets of both are identical, with the second channel having a phase of 180°, then the output of this divider will be a differential signal with common mode equal to the average of the offsets and differential amplitude 2/201 of the input amplitudes. An amplitude of* $\pm 100.5\,mV$ *on each function-generator channel will result in an amplitude of* $\pm 1\,mV$ *on the differential output, and an offset of 1.65 V on each channel will result in a 1.65 V common mode for the differential signal.*

If you have only a single function generator, then you can connect the other input to a constant voltage. The common mode will not be constant, but will be a signal with half the input amplitude. Such a test signal is still usable, but it does not allow you to see whether you are amplifying the differential signal or the common-mode signal.

Figure drawn with Digi-Key's Scheme-it [18].

Prep your skin by cleaning the spots where you want the electrodes. Using some Scotch tape to peel off a layer of dead skin (part of the *stratum corneum*) is supposedly effective in reducing the resistance of the skin. Other methods include gentle abrasion with fine sandpaper [81], shaving, and non-alcohol wipes [85]. I've had good luck just washing the skin and drying with vigorous rubbing with a towel. Do dry the skin thoroughly before attempting to attach the electrode—not only does this remove any loose dead skin cells, but the electrodes don't stick well to wet skin.

> Remember that it is safer to have the EKG circuit completely isolated from the power lines. If you are using PteroDAQ or a USB oscilloscope, running off a laptop battery is safer than using a power supply.

Not only is running off of batteries safer, but some laptop power supplies also create a lot of electrical noise. Running off of batteries can reduce noise problems substantially.

Once your breadboard version is working, populate and solder your PC board. Don't take apart your breadboard—you may want to use it to confirm that your electrodes are working correctly.

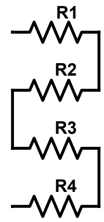

Figure 42.2: A series chain of resistors can be folded up to make the ends of resistors that are connected adjacent in the layout.

Figure drawn with Digi-Key's Scheme-it [18].

Here are some hints for laying out your PC board:

- Don't solder in electrode wires or output connections—that's what the screw terminals are for. If you don't want to use screw terminals for the power and output connections, you can solder male headers sticking out the bottom of the board to plug into a breadboard (like the male headers we put on the microcontroller board).
- Pay attention to the wiring that is already present on the back of the printed-circuit board—don't short out your safety resistors!
- When deciding where to place resistors, try to put the shared nodes of series connections next to each other to keep wiring short. A long series chain can be folded up so that shared nodes are adjacent, as shown in Figure 42.2.
- If you run out of the spaces dedicated for resistors you can use other holes on the board, but don't run resistors for long distances—the bare wires can short. Use holes either 0.1" apart (flying resistors) or 0.4" apart (flush to board). The feedback connections for op amps are conveniently 0.1" apart, which is attractive for flying resistors.
- Color code your schematic and your wiring, as described in Section 13.3.
- Check the wiring carefully—a lot of errors can be trivially debugged just by counting how many things are connected together for each node. If a resistor has no connections at one end or has too many connections, something is obviously wrong.

I highly recommend that everyone solder their own EKG board (not just one per pair of partners), so that they have something that they can demonstrate to friends and relatives after the end of the course.

42.4 Demo and write-up

Using the soldered EKG board, record several seconds of heartbeat using at least a 240 Hz sampling rate (lower sampling rates may miss the narrow R spike, unless you have done some low-pass filtering to broaden the spike.)

Plot the data with properly labeled x and y axes. Be sure that your y-axis shows the voltage at the electrodes, *not* the voltage after amplification. I expect to see EKG signals that are around 1 mV, not a volt! Remember to put what lead was used for the recording in the title of each plot.

Use a digital filter (with a modification of the script from Figure 20.7) to eliminate DC drift and 60 Hz interference. If you are trying to get cardiac-monitor bandwidth, you can use a band-pass filter with cutoffs of 0.5 Hz and 50 Hz, but if you want diagnostic EKG bandwidth, you will need a wider band-pass filter and an additional band-stop filter that cuts out a narrow band around 60 Hz (say 58 Hz to 62 Hz) and possibly yet another filter that cuts out a band around 120 Hz, if your band-pass filter includes frequencies that high.

To add an additional band-stop filter after the band-pass filter in Figure 20.7, make new variables for the new cutoff frequencies, and copy the lines using `signal.iirfilter` and `signal.sosfiltfilt` for the band-pass filter. Change the variables in the `iirfilter` to the new cutoff frequencies and change the `btype` from “bandpass” to “bandstop”. Change the input of the `sosfiltfilt` call from the `values` to `bandpass`, the already computed output of the band-pass filter, and save the result in a new variable `bandstop`. In the for-loop that outputs values, use the `bandstop` array rather than the `bandpass` array. You may want to play with the first parameter of `iirfilter`, which determines the order of the filter—a higher order makes for sharper drop-off at the cutoff frequencies, but too high an order can result in ringing artifacts and numerical instability.

Remember that both cutoff frequencies of a digital band-pass or band-stop filter must be below the Nyquist frequency, so if you want a filter with a cutoff frequency of 62 Hz, your sampling frequency needs to be above 124 Hz (and should be even higher, because cutoff frequencies very near 0 Hz or the Nyquist frequency make for numerically unstable filters). I recommend a sampling rate of 300 Hz to 1200 Hz for EKG signals if you want to use the diagnostic EKG bandwidth, and 120 Hz to 600 Hz if you want to use cardiac-monitor bandwidth.

Plot the EKG signals both before and after applying the digital filtering.

Compute the pulse rate from the recording. This should be a precise computation, based on the time between peaks, not an approximate one based on peaks in a fixed time interval. Watch out for fencepost errors!

The report should, of course, include the block diagram and schematics for the one-channel EKG, along with any design notes about why the design ended up the way that it did.

B: Study sheet

There is very little to memorize in this class. Here are the few concepts that are worth having instantly available in your memory. Your studying should not be memorizing these few formulas, but using them repeatedly to solve design and analysis problems.

B.1 Physics

$$Q = CV$$
$$I(t) = \frac{\mathrm{d}Q(t)}{\mathrm{d}t}$$
$$V = IR, \text{Ohm's law}$$

B.2 Math

$$j = \sqrt{-1}$$
$$|x + jy| = \sqrt{x^2 + y^2}$$
$$e^{j\theta} = \cos(\theta) + j\sin(\theta)$$
$$\frac{\mathrm{d}e^{j\omega t}}{\mathrm{d}t} = j\omega e^{j\omega t}$$
$$\omega = 2\pi f$$

where ω is angular frequency in radians/sec, and f is frequency in Hz

B.3 Op amps

For the generic op-amp amplifier in Figure B.1, only the approximations when open-loop gain $A \to \infty$ are worth memorizing:

Inverting: $V_{\text{out}} - V_p \approx \frac{Z_{\text{f}}}{Z_{\text{i}}}(V_p - V_m)$

Non-inverting: $V_{\text{out}} - V_m \approx \frac{Z_{\text{f}}+Z_{\text{i}}}{Z_{\text{i}}}(V_p - V_m)$

Transimpedance (set $Z_{\text{i}} = 0$ and look at I_m from amplifier into input node V_m): $V_{\text{out}} - V_p \approx Z_{\text{f}} I_m$

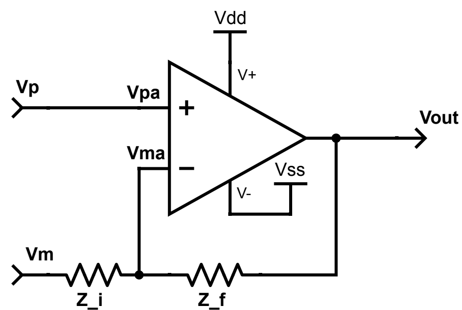

Figure B.1: Generic negative-feedback op-amp circuit, used for inverting and non-inverting amplifiers. The negative feedback tries to keep the plus and minus inputs of the op amp at the same voltage: $V_p = \frac{V_m Z_{\text{f}} + V_{\text{out}} Z_{\text{i}}}{Z_{\text{f}}+Z_{\text{i}}}$.

B.4 Impedance

$$v(t) = i(t)Z$$
$$Z = R, \text{resistor}$$
$$Z = \frac{1}{j\omega C}, \text{capacitor, angular frequency} = \omega$$
$$Z = j\omega L, \text{inductor, angular frequency} = \omega$$
$$Z_{\text{series}} = Z_1 + Z_2$$
$$Z_{\text{parallel}} = Z_1 \parallel Z_2 = \frac{1}{\frac{1}{Z_1} + \frac{1}{Z_2}} = \frac{Z_1 Z_2}{Z_1 + Z_2}$$
$$\text{gain} = \frac{V_{\text{out}}}{V_{\text{in}}} = \frac{Z_{\text{down}}}{Z_{\text{up}} + Z_{\text{down}}}, \text{voltage divider}$$
$$2\pi f = \omega = \frac{1}{RC}, \text{corner frequency for RC}$$
$$2\pi f = \omega = \frac{1}{\sqrt{LC}}, \text{resonant frequency for LC}$$

Gain of a simple RC circuit (one R, one C) is $\sqrt{2}/2$ at the corner frequency.

References

[1] Agilent Technologies.
User's Guide: Agilent 33120A 15MHz Function/Arbitrary Waveform Generator.
6th ed. Publication number 33120-90006. Agilent Technologies, Inc., Mar. 2002.
URL: https://www.keysight.com/us/en/assets/9018-04436/user-manuals/9018-04436.pdf
(visited on 27 Sept. 2021) (cit. on pp. 237, 278).

[2] American Medical Association and American Heart Association.
The One Graphic You Need for Accurate Blood Pressure Reading. 8 Feb. 2019.
URL: https://www.ama-assn.org/delivering-care/hypertension/one-graphic-you-need-accurate-blood-pressure-reading (visited on 2 Nov. 2019) (cit. on p. 344).

[3] *Analog Discovery 2 Reference Manual.* Digilent.
URL: https://digilent.com/reference/test-and-measurement/analog-discovery-2/reference-manual
(visited on 27 Sept. 2021) (cit. on pp. 244, 440).

[4] Android. *3.5mm Headset: Accessory Specification.* 2 Aug. 2022.
URL: https://source.android.com/docs/core/accessories/headset/plug-headset-spec (visited on 27 Aug. 2022) (cit. on p. 491).

[5] Art of Problem Solving. *LaTeX:About.* URL:
https://artofproblemsolving.com/wiki/index.php/LaTeX:About (visited on 15 June 2014)
(cit. on p. 89).

[6] Euan A. Ashley and Josef Niebauer. In Ashley and Niebauer. *Cardiology Explained.*
London: Remedica, 2004. Chap. 3, Conquering the ECG.
URL: https://www.ncbi.nlm.nih.gov/books/NBK2214/ (visited on 29 Aug. 2022)
(cit. on pp. 606, 611).

[7] Charles F. Babbs. "Oscillometric Measurement of Systolic and Diastolic Blood Pressures Validated in a Physiologic Mathematical Model."
BioMedical Engineering OnLine 11.56 (2012). DOI: 10.1186/1475-925X-11-56.
URL: https://biomedical-engineering-online.biomedcentral.com/articles/10.1186/1475-925X-11-56
(visited on 27 Sept. 2022) (cit. on p. 328).

[8] Peter J. Baxandall. "Negative-Feedback Tone Control."
Wireless World (Oct. 1952), pp. 402–405 (cit. on p. 491).

[9] Akshay Bhat. *Stabilize Transimpedance Amplifier Circuit Design.*
Application Note 5129. Maxim Integrated. 3 Feb. 2012.
URL: https://www.maximintegrated.com/en/design/technical-documents/app-notes/5/5129.html
(visited on 5 Sept. 2020) (cit. on p. 374).

[10] BitScope. *BitScope Model 10.*
URL: https://bitscope.com/product/BS10/ (visited on 7 Apr. 2015) (cit. on p. 244).

[11] Russ Brit. "Skin Tone Not a New Issue for Heart Monitors." *MarketWatch* (30 Apr. 2015). URL: https://www.marketwatch.com/story/skin-tone-not-a-new-issue-for-heart-monitors-2015-04-30 (cit. on p. 359).

[12] Lovely Chhabra et al. "Mouse Heart Rate in a Human: Diagnostic Mystery of an Extreme Tachyarrhythmia." *Indian Pacing and Electrophysiology Journal* 12.1 (2012), pp. 32–35. ISSN: 0972-6292 (cit. on p. 328).

[13] Eric Coates. *Learn about Electronics: Module 4.2 Amplifier Controls*. 9 Dec. 2020. URL: https://www.learnabout-electronics.org/Amplifiers/amplifiers42.php (visited on 29 Aug. 2022) (cit. on p. 491).

[14] Enrique Company-Bosch and Eckart Hartmann. "ECG Front-End Design is Simplified with MicroConverter." *Analog Dialogue* 37 (Nov. 2003). URL: https://www.analog.com/en/analog-dialogue/articles/ecg-front-end-design-simplified.html (visited on 13 June 2019) (cit. on pp. 607, 609, 613).

[15] CUI Inc. *CMA-4544PF-W Electret Condenser Microphone*. CUI Inc. 17 Jan. 2020. URL: https://www.cuidevices.com/product/resource/cma-4544pf-w.pdf (visited on 5 Sept. 2020) (cit. on p. 420).

[16] CUI Inc. *CME-1538-100LB Electret Condenser Microphone*. CUI Inc. 24 Apr. 2020. URL: https://www.cuidevices.com/product/resource/cme-1538-100lb.pdf (visited on 5 Sept. 2020) (cit. on p. 429).

[17] Deranged Physiology. *Normal Arterial Waveforms*. URL: https://derangedphysiology.com/main/cicm-primary-exam/required-reading/cardiovascular-system/Chapter%20760/normal-arterial-line-waveforms (visited on 20 Mar. 2021) (cit. on p. 409).

[18] Digi-Key. *Scheme-it: Free Online Schematic Drawing Tool*. 2014. URL: https://www.digikey.com/schemeit/ (visited on 6 Sept. 2020) (cit. on pp. 38, 45, 52, 53, 55, 56, 59–62, 73, 96, 101, 102, 179, 186, 194, 199, 202, 208, 211, 218, 228, 231, 235, 242, 250, 257–259, 263, 264, 279, 288, 295, 302, 307, 313, 315–317, 335, 348, 351, 353, 356, 362, 367, 368, 372, 378, 381, 386, 394, 395, 420, 425, 433, 434, 436, 440, 447, 452, 454, 469, 494, 495, 498, 500, 503, 504, 506, 507, 513, 520, 521, 523, 525, 526, 528, 543, 545, 550, 551, 553, 556, 557, 560, 565, 581, 583, 584, 600, 602, 610, 615, 617, 618, 651).

[19] Digilent. *Analog Discovery 2 100MSPS USB Oscilloscope, Logic Analyzer and Variable Power Supply*. URL: https://digilent.com/shop/analog-discovery-2-100ms-s-usb-oscilloscope-logic-analyzer-and-variable-power-supply/ (visited on 27 Sept. 2021) (cit. on pp. 176, 266, 267, 269, 296, 318, 320, 373–375, 397, 424, 448, 456, 471, 503, 510, 518, 520, 558, 559, 561–564, 566).

[20] Digilent. *WaveForms Reference Manual*. URL: https://digilent.com/reference/software/waveforms/waveforms-3/reference-manual (visited on 27 Sept. 2021) (cit. on p. 245).

[21] *Draw.IO*. Moved to https://app.diagrams.net/, but the draw.io link redirects. URL: https://www.draw.io (visited on 16 July 2017) (cit. on pp. 69, 70, 78, 96, 99, 234, 297, 300, 303, 552, 554, 600, 651).

[22] *Electrical Conductivity of Aqueous Solutions*. URL: https://diverdi.colostate.edu/all_courses/CRC%20reference%20data/electrical%20conductivity%20of%20aqueous%20solutions.pdf (visited on 27 Sept. 2021) (cit. on p. 576).

[23] John A. Evans and William A. Whitelaw. "The Assessment of Maximal Respiratory Mouth Pressures in Adults." *Respiratory Care* 54.10 (Oct. 2009), pp. 1348–1359. URL: https://rc.rcjournal.com/content/54/10/1348 (visited on 18 Nov. 2022) (cit. on p. 323).

[24] Everlight. *PD204-6C Technical Data Sheet: 3mm Silicon PIN Photodiode T-1*. No longer available from Everlight directly. 2005. URL: https://media.digikey.com/pdf/Data%20Sheets/Everlight%20PDFs/PD204-6C.pdf (visited on 30 Jan. 2023) (cit. on p. 352).

[25] Nick Faulkner. *The Blame Game: Things were Done, Mistakes were Made*. 2 May 2013. URL: https://nickfalkner.com/2013/05/02/the-blame-game-things-were-done-mistakes-were-made/ (visited on 5 July 2020) (cit. on p. 87).

[26] Foxboro. *Conductivity Ordering Guide*. 3 Oct. 1999. URL: https://pdf4pro.com/cdn/conductivity-ordering-guide-supersedes-august-1-1986-43e9eb.pdf (visited on 5 Sept. 2020) (cit. on p. 576).

[27] Freescale Semiconductor. *200 kPa On-Chip Temperature Compensated Silicon Pressure Sensors MPX2202 Series*. Oct. 2012. URL: https://www.nxp.com/docs/en/data-sheet/MPX2202.pdf (visited on 27 Aug. 2022) (cit. on p. 335).

[28] Freescale Semiconductor. *K20 Sub-Family Data Sheet*. K20P81M100SF2. Freescale Semiconductor (NXP). Feb. 2013. URL: https://www.nxp.com/docs/en/data-sheet/K20P81M100SF2.pdf (visited on 5 Sept. 2020) (cit. on p. 67).

[29] Freescale Semiconductor. *Kinetis KL26 Sub-Family 48 MHz Cortex-M0+ Based Microcontroller*. KL26P121M48SF4 Rev. 5. Freescale Semiconductor (NXP). Aug. 2014. URL: https://www.nxp.com/docs/en/data-sheet/KL26P121M48SF4.pdf (visited on 14 Nov. 2018) (cit. on pp. 67, 234).

[30] Freescale Semiconductor. *Kinetis KL26 Sub-Family Reference Manual*. KL26P121M48SF4RM Rev. 3.2. Freescale Semiconductor (NXP). Oct. 2013 (cit. on p. 270).

[31] Freescale Semiconductor (NXP). *FRDM-KL25Z Quick Reference Card*. URL: https://www.nxp.com/docs/en/supporting-information/FRDM-KL25Z-Quick-Reference-Card.pdf (visited on 15 Nov. 2018) (cit. on p. 621).

[32] L.A. Geddes et al. "Characterization of the Oscillometric Method for Measuring Indirect Blood Pressure." *Annals of Biomedical Engineering* 10 (1982), pp. 271–280 (cit. on p. 328).

[33] *gnuplot homepage*. Use http, as gnuplot refuses https connection. URL: http://www.gnuplot.info/ (visited on 19 Jan. 2021) (cit. on pp. 14, 43, 52, 64, 76, 96, 150, 162, 165–167, 176, 182, 187, 189, 204, 206, 209, 210, 212, 213, 241, 254–256, 259–261, 266, 267, 269–271, 292, 296, 306, 318, 320, 325, 329–331, 352, 355, 356, 358, 365, 366, 371, 373–375, 384, 390, 394, 397, 407, 412, 419, 422, 424, 439, 448, 456, 471, 495, 503, 510, 514, 518, 520, 540, 542, 548, 554, 558, 559, 561–564, 566, 577, 607, 651).

[34] Google. *Google Books Ngram Viewer*. URL: https://books.google.com/ngrams/ (visited on 18 Nov. 2022) (cit. on pp. 116, 129, 132).

[35] Google. *Google Books Ngram Viewer (bandpass)*. URL: https://books.google.com/ngrams/graph?content=bandpass%2Cband-pass%2Cband+pass (visited on 22 Mar. 2018) (cit. on p. 117).

[36] Google. *Google Books Ngram Viewer (textbook)*. URL: https://books.google.com/ngrams/graph?content=text+book%2Ctextbook%2Ctext-book&year_start=1750 (visited on 14 Mar. 2021) (cit. on p. 116).

[37] Google. *Google Books Ngram Viewer (try and)*. URL: https://books.google.com/ngrams/graph?content=try+to%2Ctry+and (visited on 26 Mar. 2018) (cit. on p. 125).

[38] Google. *Google Books Ngram Viewer (x axis)*. URL: https://books.google.com/ngrams/graph?content=x+axis%2C+x-axis%2C+y+axis%2C+y-axis&year_start=1800&year_end=2020 (visited on 12 June 2019) (cit. on p. 117).

[39] Sverre Grimnes and Ørjan Grøttem Martinsen. *Bioimpedance and Bioelectricity Basics*. 3rd ed. Academic Press, 2014. ISBN: 978-0124114708 (cit. on p. 579).

[40] Hanwei Electronics Co. Ltd. *Technical Data MQ-3 Gas Sensor*. 2017. URL: https://www.sparkfun.com/datasheets/Sensors/MQ-3.pdf (cit. on p. 68).

[41] Tom Harris and Wesley Fenlon. *How Light Emitting Diodes Work*. 31 Jan. 2002. URL: https://electronics.howstuffworks.com/led.htm (visited on 11 Dec. 2018) (cit. on p. 349).

[42] Frederic G. Hirsch et al. "The Electrical Conductivity of Blood I. Relationship to Erythrocyte Concentration." *Blood* 5.11 (1950), pp. 1017–1035. DOI: 10.1182/blood.V5.11.1017.1017. URL: https://ashpublications.org/blood/article/5/11/1017/7012/THE-ELECTRICAL-CONDUCTIVITY-OF-BLOOD-I (visited on 5 Sept. 2020) (cit. on p. 575).

[43] Hobby Projects. *About Oscilloscope: Oscilloscope Trigger Controls*. 2011. URL: https://hobbyprojects.com/oscilloscope_tutorial/oscilloscope_trigger_controls.html (visited on 5 Dec. 2019) (cit. on p. 247).

[44] Adam E. Hyatt, Paul D. Scanlon, and Masao Nakamura. *Interpretation of Pulmonary Function Tests: A Practical Guide*. 3rd ed. Lippincott Williams & Wilkins, 2011 (cit. on p. 323).

[45] *Draw Freely—Inkscape*. URL: https://inkscape.org/ (visited on 25 Nov. 2018) (cit. on pp. 81, 96, 103, 452, 651).

[46] Abraham Karplus and Kevin Karplus. *PteroDAQ*. URL: https://github.com/karplus/PteroDAQ (visited on 17 Apr. 2021) (cit. on pp. 167, 241, 259, 270, 271, 325, 329–331, 355, 356, 366, 407, 607).

[47] Kevin Karplus. *Compensation in Impedance Analyzer*. 18 July 2020. URL: https://gasstationwithoutpumps.wordpress.com/2020/06/05/compensation-in-impedance-analyzer/ (visited on 5 June 2020) (cit. on p. 468).

[48] Kevin Karplus. *Descaffolding*. 12 Feb. 2013. URL: https://gasstationwithoutpumps.wordpress.com/2013/02/12/descaffolding/ (cit. on p. viii).

[49] Kevin Karplus. *EKG without amplifier*. 18 Dec. 2017. URL: https://gasstationwithoutpumps.wordpress.com/2017/12/18/ekg-without-amplifier/ (visited on 5 Dec. 2019) (cit. on p. 616).

[50] Kevin Karplus. *Op amp protoboard rev 0 worksheet*. 8 Aug. 2015. URL: https://users.soe.ucsc.edu/~karplus/bme51/pc-boards/op-amp-proto-rev0/op-amp-proto-rev0-worksheet-fullpage.pdf (cit. on pp. 487, 615).

[51] Kevin Karplus. *Resonance for Non-linear Impedance*. 9 June 2020. URL: https://gasstationwithoutpumps.wordpress.com/2020/06/09/resonance-for-non-linear-impedance/ (visited on 22 June 2021) (cit. on p. 457).

[52] Kevin Karplus. *Showing is better than telling, but not by much*. 14 Apr. 2013. URL: https://gasstationwithoutpumps.wordpress.com/2013/04/14/showing-is-better-than-telling-but-not-by-much/ (cit. on p. vii).

[53] KEMET Corporation. *Introduction to Capacitor Technologies: What is a Capacitor?* 2013. URL: https://www.kemet.com/content/dam/kemet/lightning/documents/ec-content/technical-archive/What-is-a-Capacitor.pdf (visited on 10 Mar. 2023) (cit. on p. 170).

[54] KEMET Corporation. *Single-Ended Aluminum Electrolytic Capacitors: ESK +85*°C. 17 Nov. 2020. URL: https://content.kemet.com/datasheets/KEM_A4004_ESK.pdf (visited on 29 May 2021) (cit. on pp. 174–176).

[55] Philip Kime, Moritz Wemheuer, and Philipp Lehman. *The biblatex Package*. Version 3.19. 5 Mar. 2023. URL: https://mirrors.ibiblio.org/CTAN/macros/latex/contrib/biblatex/doc/biblatex.pdf (visited on 24 Apr. 2023) (cit. on p. 138).

[56] Kingbright. *Phototransistor WP3DP3BT Datasheet*. Version 5. 7 Oct. 2016. URL: https://www.kingbrightusa.com/images/catalog/SPEC/wp3dp3bt.pdf (visited on 24 Nov. 2018) (cit. on p. 138).

[57] Henry Knipe and Rachael Nightingale. *Ultrasound Transducer*. Sept. 2017. URL: https://radiopaedia.org/articles/ultrasound-transducer (cit. on p. 301).

[58] Henry Knipe and Patricia O'Gorman. *Propagation Speed*. URL: https://radiopaedia.org/articles/propagation-speed (visited on 5 Sept. 2020) (cit. on p. 300).

[59] Knowles Acoustics. *Ultra-Mini SiSonic™ Microphone Specification With MaxRF Protection.* 27 Mar. 2013. URL: https://media.digikey.com/pdf/Data%20Sheets/Knowles%20Acoustics%20PDFs/SPU0410HR5H-PB.pdf (visited on 27 Aug. 2022) (cit. on p. 419).

[60] Donald E. Knuth. *The TEXbook.* Addison-Wesley Professional, 1984. ISBN: 978-0201134483 (cit. on p. 88).

[61] Sébastien Krul. *Cardiac Arrhythmias.* URL: https://www.textbookofcardiology.org/wiki/Cardiac_Arrhythmias (visited on 5 Sept. 2020) (cit. on p. 611).

[62] Michael M. Laks et al. "Recommendations for Safe Current Limits for Electrocardiographs: A Statement for Healthcare Professionals From the Committee on Electrocardiography, American Heart Association." *Circulation* 93 (1996), pp. 837–839. DOI: 10.1161/01.CIR.93.4.837. URL: https://www.ahajournals.org/doi/full/10.1161/01.cir.93.4.837 (visited on 5 Sept. 2020) (cit. on pp. 609, 613).

[63] Leslie Lamport. *LATEX: A Document Preparation System.* 2nd ed. Addison-Wesley Professional, 1994. ISBN: 978-0201529838 (cit. on pp. 88, 89, 136).

[64] W. Marshall Leach. In W. Marshall Leach, Jr. *Introduction to Electroacoustics and Audio Amplifier Design.* Second Edition—Revised Printing. Kendall/Hunt, 2001. Chap. The Class-D Amplifier. URL: https://leachlegacy.ece.gatech.edu/ece4435/f01/ClassD2.pdf (visited on 27 Aug. 2022) (cit. on p. 544).

[65] LEGO Group. *Fair Play.* URL: https://www.lego.com/en-us/legal/notices-and-policies/fair-play/ (visited on 20 Mar. 2021) (cit. on p. 411).

[66] *matplotlib.* URL: https://matplotlib.org/ (visited on 25 Nov. 2018) (cit. on pp. 35, 162, 651).

[67] Maxim Integrated. *Penetration Depth Guide for Biosensor Applications.* 30 Mar. 2018. URL: https://www.maximintegrated.com/en/design/technical-documents/app-notes/6/6433.html (visited on 27 Nov. 2019) (cit. on p. 359).

[68] Mayo Clinic. *Ankle-brachial Index.* URL: https://www.mayoclinic.org/tests-procedures/ankle-brachial-index/about/pac-20392934 (visited on 9 Dec. 2018) (cit. on p. 326).

[69] David Meeker. *Finite Element Method Magnetics: Woofer—Analysis of a Woofer Motor.* URL: https://www.femm.info/wiki/Woofer (visited on 5 Sept. 2020) (cit. on p. 452).

[70] Dan Meyer. *Asilomar #4: Be Less Helpful.* 2009. URL: https://blog.mrmeyer.com/2009/asilomar-4-be-less-helpful/ (cit. on p. viii).

[71] Microchip. *MCP6001/R/U/2/4, 1 MHz, Low-Power Op Amp.* 14 May 2019. URL: https://ww1.microchip.com/downloads/en/DeviceDoc/MCP6001-1R-1U-2-4-1-MHz-Low-Power-Op-Amp-DS20001733L.pdf (visited on 5 Sept. 2020) (cit. on pp. 305, 371).

[72] James Moffatt. *E C G Basics*.
URL: https://www.nataliescasebook.com/tag/e-c-g-basics (visited on 19 Oct. 2018) (cit. on p. 611).

[73] S. Mohan et al. "Simple Accurate Expressions for Planar Spiral Inductances." *IEEE Journal of Solid-State Circuits* 34.10 (Oct. 1999), pp. 1419–1424. URL: https://web.stanford.edu/~boyd/papers/inductance_expressions.html (visited on 5 July 2020) (cit. on p. 446).

[74] Matt A. Morgan and Mirjan M. Nadrljanski. *Ultrasound Frequencies*. Sept. 2017. URL: https://radiopaedia.org/articles/ultrasound-frequencies (cit. on p. 300).

[75] *MQ303A Alcohol Sensor*. 2017. URL: https://raw.githubusercontent.com/SeeedDocument/Grove-Alcohol_Sensor/master/res/MQ303A.pdf (cit. on p. 68).

[76] Randall Munroe. *.NORM Normal File Format*.
URL: https://xkcd.com/2116/ (visited on 6 July 2019) (cit. on p. 104).

[77] Randall Munroe. *Curve Fitting*. URL: https://xkcd.com/2048/ (visited on 6 July 2019) (cit. on p. 110).

[78] Munir H. Nayfeh and Morton K. Brussel. *Electricity and Magnetism*. Copyright 1985. Dover Publications, 2015. ISBN: 978-0-486-78971-2 (cit. on p. 578).

[79] Michael R. Neuman. "Biopotential Electrodes." In *The Biomedical Engineering Handbook: Second Edition*. Ed. by Joseph D. Bronzino. Boca Raton: CRC Press LLC, 2000. Chap. 48. URL: https://web.archive.org/web/20220330061138if_/http://www.fis.uc.pt/data/20062007/apontamentos/apnt_134_5.pdf (visited on 29 Aug. 2022) (cit. on p. 579).

[80] Nexperia. *Understanding Power MOSFET Data Sheet Parameters*. 6 July 2020. URL: https://assets.nexperia.com/documents/application-note/AN11158.pdf (visited on 27 Aug. 2022) (cit. on p. 511).

[81] Craig D. Oster. *Proper Skin Prep Helps Ensure ECG Trace Quality*. 2005. URL: https://multimedia.3m.com/mws/media/358372O/proper-skin-prep-ecg-trace-quality-white-paper.pdf (visited on 5 Dec. 2018) (cit. on p. 617).

[82] Overleaf. *Documentation*.
URL: https://www.overleaf.com/learn (visited on 14 Nov. 2018) (cit. on p. 89).

[83] Oxford University Press. *Oxford Learner's Dictionaries*. 2019.
URL: https://www.oxfordlearnersdictionaries.com/us (visited on 6 July 2019) (cit. on p. 117).

[84] Richard Palmer. *DC Parameters: Input Offset Voltage (V_{IO})*. SLOA059. Texas Instruments. Mar. 2001.
URL: https://www.ti.com/lit/an/sloa059/sloa059.pdf (visited on 27 Aug. 2022) (cit. on p. 293).

[85] Philips Healthcare. *Improving ECG Quality*. Sept. 2008.
URL: https://www.documents.philips.com/doclib/enc/fetch/577817/577869/Improving_ECG_Quality_Application_Note_(ENG).pdf (visited on 24 Apr. 2023) (cit. on p. 617).

[86] Physio Control. *DIAG Frequency Response Necessary to Reproduce ST Segment.* 2009. URL: https://emtlife.com/attachments/diag-frequency-response-necessary-to-reproduce-st-segment-3304790-a-3-pdf.1603 (visited on 18 Nov. 2022) (cit. on p. 607).

[87] PJRC. *Teensy Pin Assignments.* URL: https://www.pjrc.com/teensy/pinout.html (visited on 15 Nov. 2018) (cit. on p. 621).

[88] John Prymak et al. "Why That 47 uF Capacitor Drops to 37 uF, 30 uF, or Lower." In *Proceedings CARTS USA 2008, 28th Symposium for Passive Electronics.* Newport Beach, CA: Electronics Components, Assemblies & Materials Association, Mar. 2008. URL: https://www.kemet.com/content/dam/kemet/lightning/documents/ec-content/technical-archive/Why-47-uF-capacitor-drops-to-37-uF-30-uF-or-lower.pdf (visited on 18 Nov. 2022) (cit. on pp. 173–175).

[89] Purdue Online Writing Lab. *Extended Rules for Using Commas.* 2020. URL: https://owl.purdue.edu/owl/general_writing/punctuation/commas/extended_rules_for_commas.html (visited on 27 Feb. 2021) (cit. on p. 127).

[90] Andrew Randazzo. *The Ultimate 12-Lead ECG Placement Guide (with Illustrations).* 1 June 2016. URL: https://www.primemedicaltraining.com/12-lead-ecg-placement/ (visited on 5 Dec. 2018) (cit. on p. 605).

[91] Arie Van Rhijn. *LMV1012: Integrated Circuits for High Performance Electret Microphones.* SNAA114. Texas Instruments. 2011. URL: https://www.ti.com/lit/wp/snaa114/snaa114.pdf (cit. on p. 415).

[92] *Salts Weight vs. Volume Conversion.* URL: https://convert-to.com/types-of-salts-weight-volume-converters (visited on 24 Apr. 2023) (cit. on p. 586).

[93] Warren Schultz. *Interfacing Semiconductor Pressure Sensors to Microcomputers.* URL: https://www.nxp.com/webapp/Download?colCode=AN1318&location=null (visited on 18 Nov. 2022) (cit. on pp. 334, 335).

[94] Stanford Medicine. *Measuring and Understanding the Ankle Brachial Index (ABI).* URL: https://stanfordmedicine25.stanford.edu/the25/ankle-brachial-index.html (visited on 9 Dec. 2018) (cit. on p. 326).

[95] Milan Stork and Jiri Jilek. "Cuff Pressure Waveforms: Their Current and Prospective Application in Biomedical Instrumentation." In *Biomedical Engineering, Trends in Electronics, Communications and Software.* Ed. by Anthony N. Laskovski. InTech, 8 Jan. 2011. ISBN: 978-953-307-475-7. DOI: 10.5772/13475. URL: https://www.intechopen.com/chapters/12908 (visited on 27 Aug. 2022) (cit. on p. 328).

[96] Robert Talbert. *Examples and the Light Bulb.* 25 Mar. 2013. URL: https://www.chronicle.com/blognetwork/castingoutnines/2013/03/25/examples-and-the-light-bulb/ (cit. on p. vii).

[97] Texas Instruments.
INAx126 MicroPower Instrumentation Amplifier, Single and Dual Versions.
Dec. 2015. URL: https://www.ti.com/lit/ds/symlink/ina126.pdf (visited on 8 Dec. 2018) (cit. on pp. 293, 603).

[98] Texas Instruments.
LM339, LM239, LM139, LM2901 Quad Differential Comparators. Version Rev. U. Nov. 2018. URL: https://www.ti.com/lit/ds/symlink/lm139.pdf (visited on 8 Dec. 2019) (cit. on p. 524).

[99] *Theory and Practice of pH Measurement.* PN 44-6033/rev. D.
Emerson Process Management. Dec. 2010.
URL: https://www.emerson.com/documents/automation/manual-theory-practice-of-ph-measurement-rosemount-en-70736.pdf (visited on 5 Sept. 2020) (cit. on p. 299).

[100] Robert H. Tyler et al. "Electrical Conductivity of the Global Ocean."
Earth, Planets, and Space 69.156 (14 Nov. 2017). DOI: 10.1186/s40623-017-0739-7.
URL: https://earth-planets-space.springeropen.com/articles/10.1186/s40623-017-0739-7 (visited on 27 Aug. 2022) (cit. on p. 586).

[101] Vishay Siliconix. *Power MOSFET Basics: Understanding Gate Charge and Using it to Assess Switching Performance.* Application Note 608A.
Vishay Siliconix. 16 Feb. 2016.
URL: https://www.vishay.com/docs/73217/an608a.pdf (visited on 24 Dec. 2018) (cit. on p. 511).

[102] Vyaire Medical. *MicroGard®.*
URL: https://www.vyaire.com/products/microgard-ii-pft-filter (visited on 5 Sept. 2020) (cit. on p. 324).

[103] J. Watson. "The Tin Oxide Gas Sensor and its Applications."
Sensors and Actuators 5 (1984), pp. 29–42 (cit. on p. 68).

[104] Grant Wiggins. *Autonomy and the need to back off by design as teachers.*
12 Feb. 2013. URL: https://grantwiggins.wordpress.com/2013/02/12/autonomy-and-the-need-to-back-off-by-design-as-teachers/ (cit. on p. viii).

[105] Wikibooks. *LaTeX.* URL: https://en.wikibooks.org/wiki/LaTeX (visited on 15 June 2014) (cit. on p. 89).

[106] wikiHowStaff. *How to Solder Aluminum.* 18 June 2019.
URL: https://www.wikihow.com/Solder-Aluminum (visited on 24 Aug. 2019) (cit. on p. 280).

[107] Wikipedia. *A-weighting.*
URL: https://en.wikipedia.org/wiki/A-weighting (visited on 19 Dec. 2016) (cit. on p. 480).

[108] Wikipedia. *Action potential.*
URL: https://en.wikipedia.org/wiki/Action_potential (visited on 26 May 2015) (cit. on p. 611).

[109] Wikipedia. *Blood pressure.*
URL: https://en.wikipedia.org/wiki/Blood_pressure (visited on 7 Sept. 2014) (cit. on p. 326).

[110] Wikipedia. *Callendar-Van Dusen equation.*
URL: https://en.wikipedia.org/wiki/Callendar%E2%80%93Van_Dusen_equation (visited on 29 Aug. 2022) (cit. on p. 65).

[111] Wikipedia. *Capacitance.*
URL: https://en.wikipedia.org/wiki/Capacitance (visited on 4 July 2014) (cit. on p. 170).

[112] Wikipedia. *Capacitive coupling.*
URL: https://en.wikipedia.org/wiki/Capacitive_coupling (visited on 7 Apr. 2015) (cit. on p. 246).

[113] Wikipedia. *Carbonic Acid.*
URL: https://en.wikipedia.org/wiki/Carbonic_acid (visited on 28 Dec. 2018) (cit. on p. 577).

[114] Wikipedia. *Cardiac Action Potential.*
URL: https://en.wikipedia.org/wiki/Cardiac_action_potential (visited on 16 Dec. 2017) (cit. on p. 611).

[115] Wikipedia. *Class-D amplifier.*
URL: https://en.wikipedia.org/wiki/Class_d_amplifier (visited on 17 Feb. 2013) (cit. on p. 569).

[116] Wikipedia. *Color blindness.*
URL: https://en.wikipedia.org/wiki/Color_blindness (visited on 6 Sept. 2015) (cit. on p. 241).

[117] Wikipedia. *Concert Pitch.*
URL: https://en.wikipedia.org/wiki/Concert_pitch (visited on 2 July 2017) (cit. on p. 12).

[118] Wikipedia. *Conductivity (electrolytic).*
URL: https://en.wikipedia.org/wiki/Conductivity_(electrolytic) (visited on 18 Dec. 2015) (cit. on pp. 575, 577).

[119] Wikipedia. *Decoupling capacitor.*
URL: https://en.wikipedia.org/wiki/Bypass_capacitor (visited on 4 July 2014) (cit. on p. 215).

[120] Wikipedia. *Diode modelling.*
URL: https://en.wikipedia.org/wiki/Diode_modelling (visited on 28 Sept. 2016) (cit. on p. 364).

[121] Wikipedia. *Direct coupling.*
URL: https://en.wikipedia.org/wiki/Direct_coupling (visited on 7 Apr. 2015) (cit. on p. 246).

[122] Wikipedia. *Electret.*
URL: https://en.wikipedia.org/wiki/Electret (visited on 16 June 2014) (cit. on p. 413).

[123] Wikipedia. *Electret microphone.*
URL: https://en.wikipedia.org/wiki/Electret_microphone (visited on 16 June 2014) (cit. on p. 413).

[124] Wikipedia. *Electric shock.*
URL: https://en.wikipedia.org/wiki/Electric_shock (visited on 26 May 2015) (cit. on p. 608).

[125] Wikipedia. *Electrical resistance.*
URL: https://en.wikipedia.org/wiki/Electrical_resistance (visited on 15 June 2014) (cit. on p. 40).

[126] Wikipedia. *Electrocardiography.*
URL: https://en.wikipedia.org/wiki/EKG (visited on 7 Mar. 2013) (cit. on p. 611).

[127] Wikipedia. *Electrolytic capacitor.*
URL: https://en.wikipedia.org/wiki/Electrolytic_capacitor (visited on 25 June 2014) (cit. on p. 173).

[128] Wikipedia. *Engineering notation.*
URL: https://en.wikipedia.org/wiki/Engineering_notation (visited on 29 Mar. 2018) (cit. on p. 90).

[129] Wikipedia. *Inductance.*
URL: https://en.wikipedia.org/wiki/Inductance (visited on 6 Aug. 2014) (cit. on p. 448).

[130] Wikipedia. *Inductance: Self-inductance of thin wire shapes.*
URL: https://en.wikipedia.org/wiki/Inductance#Self-inductance_of_thin_wire_shapes (visited on 1 Aug. 2020) (cit. on p. 445).

[131] Wikipedia. *ITU-R 468 Noise Weighting.*
URL: https://en.wikipedia.org/wiki/ITU-R_468_noise_weighting (visited on 13 Oct. 2016) (cit. on p. 418).

[132] Wikipedia. *JFET.* URL: https://en.wikipedia.org/wiki/JFET (visited on 29 Oct. 2017) (cit. on p. 416).

[133] Wikipedia. *Lanczos Resampling.*
URL: https://en.wikipedia.org/wiki/Lanczos_resampling (visited on 22 Oct. 2018) (cit. on p. 161).

[134] Wikipedia. *Magnetic core.*
URL: https://en.wikipedia.org/wiki/Magnetic_core (visited on 6 Aug. 2014) (cit. on p. 446).

[135] Wikipedia. *Microphone.*
URL: https://en.wikipedia.org/wiki/Microphone (visited on 16 June 2014) (cit. on p. 413).

[136] Wikipedia. *Microphone: Condenser microphone.* URL: https://en.wikipedia.org/wiki/Microphone#Capacitor_microphone (visited on 27 Aug. 2022) (cit. on p. 413).

[137] Wikipedia. *MIDI Tuning Standard.*
URL: https://en.wikipedia.org/wiki/MIDI_tuning_standard (visited on 2 July 2017) (cit. on p. 15).

[138] Wikipedia. *Molar conductivity.*
URL: https://en.wikipedia.org/wiki/Molar_conductivity (visited on 19 Dec. 2015) (cit. on p. 575).

[139] Wikipedia. *Near-infrared Window in Biological Tissue.*
URL: https://en.wikipedia.org/wiki/Near-infrared_window_in_biological_tissue (visited on 29 Sept. 2016) (cit. on p. 359).

[140] Wikipedia. *Nernst Equation*. URL: https://en.wikipedia.org/wiki/Nernst_equation (visited on 19 Dec. 2015) (cit. on p. 580).

[141] Wikipedia. *Nyquist-Shannon sampling theorem*. URL: https://en.wikipedia.org/wiki/Nyquist-Shannon_sampling_theorem (visited on 19 June 2014) (cit. on p. 161).

[142] Wikipedia. *Operational amplifier*. URL: https://en.wikipedia.org/wiki/Operational_amplifier (visited on 19 Dec. 2016) (cit. on p. 480).

[143] Wikipedia. *Operational amplifier applications*. URL: https://en.wikipedia.org/wiki/Operational_amplifier_applications (visited on 19 Dec. 2016) (cit. on p. 480).

[144] Wikipedia. *Oscilloscope*. URL: https://en.wikipedia.org/wiki/Oscilloscope (visited on 20 July 2014) (cit. on p. 243).

[145] Wikipedia. *Phasor*. URL: https://en.wikipedia.org/wiki/Phasor (visited on 7 Aug. 2014) (cit. on p. 178).

[146] Wikipedia. *Photodiode*. URL: https://en.wikipedia.org/wiki/Photodiode (visited on 4 Feb. 2013) (cit. on p. 351).

[147] Wikipedia. *Photometry (optics)*. URL: https://en.wikipedia.org/wiki/Photometry_(optics) (visited on 23 Aug. 2019) (cit. on p. 349).

[148] Wikipedia. *Pulse-width Modulation*. URL: https://en.wikipedia.org/wiki/Pulse-width_modulation (visited on 27 Aug. 2022) (cit. on p. 540).

[149] Wikipedia. *Purified water*. URL: https://en.wikipedia.org/wiki/Purified_water (visited on 3 Jan. 2017) (cit. on p. 577).

[150] Wikipedia. *Relaxation Oscillator*. URL: https://en.wikipedia.org/wiki/Relaxation_oscillator (visited on 27 Aug. 2022) (cit. on pp. 257, 274).

[151] Wikipedia. *Response of Silicon Photodiode*. URL: https://en.wikipedia.org/wiki/File:Response_silicon_photodiode.svg (visited on 27 Aug. 2022) (cit. on p. 359).

[152] Wikipedia. *Sallen-Key topology*. URL: https://en.wikipedia.org/wiki/Sallen%E2%80%93Key_topology (visited on 7 Nov. 2018) (cit. on p. 395).

[153] Wikipedia. *Shockley diode equation*. URL: https://en.wikipedia.org/wiki/Shockley_diode_equation (visited on 28 Sept. 2016) (cit. on p. 364).

[154] Wikipedia. *Silver chloride electrode*. URL: https://en.wikipedia.org/wiki/Silver_chloride_electrode (visited on 19 Apr. 2015) (cit. on p. 580).

[155] Wikipedia. *Sinc Filter*.
URL: https://en.wikipedia.org/wiki/Sinc_filter (visited on 22 Oct. 2018) (cit. on p. 161).

[156] Wikipedia. *Skin Effect: Impedance of a Round Wire*. URL:
https://en.wikipedia.org/wiki/Skin_effect#Impedance_of_round_wire (visited on 1 Aug. 2020) (cit. on p. 445).

[157] Wikipedia. *Sound Level: Sound Pressure Level*.
URL: https://en.wikipedia.org/wiki/Sound_pressure#Sound_pressure_level (visited on 27 Aug. 2022) (cit. on pp. 417, 418, 480).

[158] Wikipedia. *Spectral sensitivity*.
URL: https://en.wikipedia.org/wiki/Spectral_sensitivity (visited on 3 May 2014) (cit. on p. 349).

[159] Wikipedia. *Stainless steel*.
URL: https://en.wikipedia.org/wiki/Stainless_steel (visited on 19 Apr. 2015) (cit. on p. 579).

[160] Wikipedia. *Standard hydrogen electrode*.
URL: https://en.wikipedia.org/wiki/Standard_hydrogen_electrode (visited on 19 Dec. 2015) (cit. on p. 578).

[161] Wikipedia. *Steinhart-Hart equation*.
URL: https://en.wikipedia.org/wiki/Steinhart-Hart_equation (visited on 18 June 2014) (cit. on p. 63).

[162] Wikipedia. *Sunlight*.
URL: https://en.wikipedia.org/wiki/Sunlight (visited on 26 Sept. 2016) (cit. on p. 363).

[163] Wikipedia. *Surgical steel*.
URL: https://en.wikipedia.org/wiki/Surgical_steel (visited on 19 Apr. 2015) (cit. on p. 579).

[164] Wikipedia. *Thermistor*.
URL: https://en.wikipedia.org/wiki/Thermistor (visited on 4 July 2014) (cit. on p. 63).

[165] Wikipedia. *Thévenin's Theorem*. URL:
https://en.wikipedia.org/wiki/Th%C3%A9venin%27s_theorem (visited on 27 Aug. 2022) (cit. on p. 56).

[166] S.H. Wilson et al. "Predicted Normal Values for Maximal Respiratory Pressures in Caucasian Adults and Children." *Thorax* 39.7 (July 1984), pp. 535–538 (cit. on p. 323).

[167] Wolfram Alpha LLC. *Wolfram|Alpha: Computational Intelligence*.
URL: https://www.wolframalpha.com/ (visited on 1 Feb. 2021) (cit. on p. 202).

[168] Al Wright.
Printed Circuit Board Surface Finishes—Advantages and Disadvantages.
URL: https://www.epectec.com/articles/pcb-surface-finish-advantages-and-disadvantages.html (visited on 18 July 2017) (cit. on p. 28).

[169] Würth Elektronik.
151034BS03000 WL-TMRW THT Mono-color Round Waterclear. 2 Nov. 2018.
URL: https://www.we-online.com/components/products/datasheet/151034BS03000.pdf (visited on 29 Jan. 2023) (cit. on p. 350).

Index

Colophon

This book was typeset by the author using LaTeX and the Computer Modern family of fonts.

Schematics were created using Scheme-it, a free schematic editor provided by Digi-Key Corporation [18]. Block diagrams were created using https://www.draw.io [21]. Graphs were created primarily using gnuplot [33], with one graph (Figure 9.2b) using Matplotlib [66]. Other line art was created with hand-written SVG code, with a few uses of Inkscape [45]. Photographs were taken by the author, using a Canon G10 digital camera.

The design of the title page is based on a LaTeX template by Peter Wilson (herries.press@earthlink.net).