Lab Manual for Physics 223

Benjamin Crowell
Fullerton College

www.lightandmatter.com

Contents

4   Contents

1 Absolute Zero

Note to the students: As (hopefully) announced previously, you need to wear close-toed shoes, and it is up to you to buy your own safety glasses, although the physics department has a few.

Note to the lab technician: Please put the alcohol and acetone in a freezer overnight, then put them in the ice chest right before lab. Please also purchase dry ice. When you put out the waste disposal container, please don't leave the cap with it, because we don't want students to screw on the cap; if they did, evaporating dry ice could make it explode like a bomb.

Apparatus

electric heating pad
oven mitts
latex tubing
ice chest to keep liquids cool
acetone (1.5 liter for the whole class)
alcohol
mineral oil
waste disposal container
dry ice (9 lb)
tongs
hammer
funnels
gas pressure sensor
temperature probe
125 ml Erlenmyer flask
600 ml beaker 2/group
safety goggles
fold-top sandwich bags

Introduction

If heat is a form of random molecular motion, then it makes sense that there is some minimum temperature at which the molecules aren't moving at all. With fancy equipment, physicists have gotten samples of matter to within a fraction of a degree above absolute zero, but they have never actually reached absolute zero (and the laws of thermodynamics actually imply that they never can). Nevertheless, we can determine how cold absolute zero is without even getting very close to it. Kinetic theory tells us that heat is composed of random molecular motion, and

temperature is interpeted as a measure of the average kinetic energy per molecule; the zero of the *absolute* temperature scale occurs when all molecular motion is eliminated. Suppose we heat up a gas so that the typical speeds of the atoms are doubled. The kinetic energy depends on v^2, so the result is that the temperature is quadrupled:

$$T \to 4T$$

In this lab, we'll be heating and cooling air while it is sealed inside a flask with a fixed volume. We won't actually be quadrupling the absolute temperature in this lab, but just to get the idea, let's pretend that we were. The hot gas exerts more pressure on the inside of the flask, for two reasons: (1) the molecules are moving twice as fast, so when they hit the sides of the flask, each impact is twice as hard; (2) because the molecules are moving twice as fast, they also take less time to cross from one side of the flask to the other, so the collisions occur twice as frequently. The result is that the pressure is quadrupled:

$$P \to 4P$$

Based on these arguments, we conclude that in general, the pressure of a gas maintained at constant volume is proportional to its absolute temperature:

$$P \propto T$$

In this lab, you'll measure the volume of a sample of air at temperatures between about -70 and 150 degrees C, and determine where absolute zero lies by extrapolating to the temperature at which it would have had zero pressure.

Because absolute zero is very far below room temperature, this is a long extrapolation. Extrapolating a long way like this tends to be inaccurate unless you can get data covering a large range, so that the slope is well determined. For this reason, we want to get a set of temperatures that goes as high and as low as possible.

Observations

The following important rules serve to keep facts separate from opinions and reduce the chances of getting a garbled copy of the data:

(1) Take your raw data in pen, directly into your lab notebook. This is what real scientists do. The point is to make sure that what you're writing down is a first-hand record, without mistakes introduced by recopying it. (If you don't have your two lab notebooks yet, staple today's raw data into your notebook when you get it.)

(2) Everybody should record their own copy of the raw data. Do not depend on a "group secretary."

(3) If you do calculations during lab, keep them on a separate page or draw a line down the page and keep calculations on one side of the line and raw data on the other. This is to distinguish facts from inferences. (I will deduct 25% from your grade if you mix calculations and raw data.)

(4) Never write numbers without units. Without units, a number is meaningless. There is a big difference bewteen "Johnny is six" and "Johnny is six feet." (I will deduct 25% from your grade if you write numbers without units.)

Because this is the first meeting of the lab class, there is no prelab writeup due at the beginning of the class. Instead, you will discuss your results with your instructor at various points.

- *Hot mineral oil* Thoroughly dry the beaker in which you'll heat the mineral oil; if there are drops of water mixed into the oil, the oil will spatter. Measure and plan the volume of oil you will use. If you use too little, it won't cover the whole Erlenmeyer flask that holds the sample of air. If you use too much, it will overflow and make a nasty mess when you dunk the flask in it. Start heating the mineral oil. Keep an eye on the temperature. You should heat it up to about 150 °C; above that, it starts to smoke.

While you're waiting for the mineral oil to heat up, plug the temperature and pressure sensors into CH1 and CH2 on the LabPro interface. Use the Logger Pro software to view the temperature readout. Put the stopper in the Erlenmyer flask in order to make a sealed sample of air. There is an extra port on the stopper with a blue stopcock; make sure the stopcock is closed, so that the sample is sealed. Connect the flask to the pressure sensor using the latex tubing. The temperature probe goes in the liquid, not the air.

The whole lab is predicated on the ability to maintain the same sample of gas at a range of temperatures. Therefore if you have a leak, you have to do the whole lab over. Make sure the bayonet connector with the stopcock is firmly shoved into the hole in the rubber stopper. If you notice that the pressure doesn't change as the temperature changes, it means you have a leak.

A practical difficulty in this lab is that if the flask is initially sealed, and then heated to a higher temperature, the stopper tends to pop out due to the higher pressure. To keep this from happening, we want to start off with a sample of air that is hot and at atmospheric pressure; then all the other pressures will be at lower than atmospheric pressure, which will tend to suck the stopper down into the flask rather than popping it out. To accomplish this, take the oil off the burner, open the stopcock on the flask, immerse the flask in the oil, wait a little bit for the air inside the flask to heat up, and then close the stopcock again. The Erlenmyer flask wants to bob up out of the oil, so use some tape to hold it down. Take pressure and temperature data.

Although it's undesirable that the small amount of air in the tubing won't be at exactly the same temperature as the rest of the air, we can't avoid this, because the mineral oil is hot enough to melt the tubing.

When you're done with the mineral oil, wash the beaker with soap and water.

We'll next do a series of measurements at lower temperatures:

- *Room-temperature tap water* Make sure that the pressure drops by about a third when you come down to this temperature. If it doesn't, you probably have an air leak.

- *Ice-cold alcohol*

- *Acetone/dry ice slurry:* Use the hammer to knock off a piece of dry ice. Remove the piece using the tongs, stick it in a baggie, and crush it up some more with a hammer. Add the dry ice to the acetone (nail polish remover) slowly; if you do it rapidly, it can fizz violently. Mix the dry ice and acetone to make a slush. Acetone is flammable, so avoid creating any sparks or flames. This mixture cannot be dumped down the drain when you're done; keep it so that it can be disposed of properly. To reduce the amount of waste disposal, you can reuse another group's slurry. You should be able to get the temperature down to about -60 to -80 celsius. If you only get to -20, you're doing something wrong.

Analysis

Graph the temperature and pressure against each other. Does the graph appear to be linear? If so, extrapolate to find the temperature at which the pressure would be zero.

Error analysis and propagation of errors are discussed in Appendices 2 and 3, which you should read if you haven't had a previous lab course that did these topics.

If your data are nice and linear, then your main source of error will be random errors, and you should then determine error bars for your value of absolute zero using the techniques discussed in Appendix 4. The appendix discusses finding the *slope* of a line, although in this lab it's actually the x- or y-intercept that you want; the technique is analogous, however. The easiest way to estimate the error bars on the points is to use the typical amount of scatter of the points about the best-fit line. For example, if the systematic trend of the data is linear, but the points generally lie an average of about $5\,^\circ\text{C}$ away from the line, then the error bars are approximately $5\,^\circ\text{C}$.

Compare your result with the accepted value and give a probabilistic interpretation as in the example on p. 75.

Notes For Next Week

(1) Next week, when you turn in your writeup for this lab, you also need to turn in a prelab writeup for the next lab. The prelab questions are listed at the end of the description of that lab in the lab manual. Never start a lab without understanding the answers to all the prelab questions; if you turn in partial answers or answers you're unsure of, discuss the questions with your instructor or with other students to make sure you understand what's going on.

(2) You should exchange phone numbers with your lab partners for general convenience throughout the semester. You can also get each other's e-mail addresses by logging in to Spotter and clicking on "e-mail."

Rules and Organization

Collection of raw data is work you share with your lab partners. Once you're done collecting data, you need to do your own analysis. E.g., it is not okay for two people to turn in the same calculations, or on a lab requiring a graph for the whole group to make one graph and turn in copies.

You'll do some labs as formal writeups, others as informal "check-off" labs. As described in the syllabus, they're worth different numbers of points, and you have to do a certain number of each type by the end of the semester.

The format of formal lab writeups is given in appendix 1 on page 70. The raw data section must be contained in your bound lab notebook. Typically people word-process the abstract section, and any other sections that don't include much math, and stick the printout in the notebook to turn it in. The calculations and reasoning section will usually just consist of hand-written calculations you do in your lab notebook. You need two lab notebooks, because on days when you turn one in, you need your other one to take raw data in for the next lab. You may find it convenient to leave one or both of your notebooks in the cupboard at your lab bench whenever you don't need to have them at home to work on; this eliminates the problem of forgetting to bring your notebook to school.

For a check-off lab, the main thing I'll pay attention to is your abstract. The rest of your work for a check-off lab can be informal, and I may not ask to see it unless I think there's a problem after reading your abstract.

2   Two-Source Interference

Apparatus

Goals

> Observe how a 2-source interference pattern of
> water waves depends on the distance between
> the sources.

Observations

A car or a pool ball travels along a well-defined path,
but a wave doesn't. Not only that, but waves don't
simply collide and bounce like pool balls. They over-
lap and add on to one another, so that they can
reinforce or cancel. This seems like it would be ex-
tremely complicated to calculate.

Life isn't that bad. It turns out that all of the most
important ideas about light as a wave can be seen

in one simple experiment, shown in the first figure.[1]
A wave comes up from the bottom of the page, and
encounters a wall with two slits chopped out of it.
The result is a fan pattern, with strong wave motion
coming out along directions like X and Z, but no
vibration of the water at all along lines like Y. The
reason for this pattern is shown in the second figure.
The two parts of the wave that get through the slits
create an overlapping pattern of ripples. To get to
a point on line X, both waves have to go the same
distance, so they're in step with each other, and re-
inforce. But at a point on line Y, due to the unequal
distances involved, one wave is going up while the
other wave is going down, so there is cancellation.
The angular spacing of the fan pattern depends on
both the wavelength of the waves, λ, and the dis-
tance between the slits, d.

The ripple tank is tank that sits about 30 cm above
the floor. You put a little water in the tank, and
produce waves. There is a lamp above it that makes
a point-like source of light, and the waves cast pat-
terns of light on a screen placed on the floor. The
patterns of light on the screen are easier to see and
measure than the ripples themselves.

In reality, it's not very convenient to produce a double-
slit diffraction pattern exactly as depicted in the first
figure, because the waves beyond the slits are so
weak that they are difficult to observe clearly. In-
stead, you'll simply produce synchronized circular
ripples from two sources driven by a motor.

Put the tank on the floor. Plug the hole in the side of
the tank with the black rubber stopper. If the plastic
is dirty, clean it off with alcohol and kimwipes. Wet
the four yellow foam pads, and place them around
the sides of the tank. Pour in water to a depth of
about 5-7 mm. Adjust the metal feet to level the
tank, so that the water is of equal depth throughout
the tank. (Do not rotate the wooden legs them-

---

[1] The photo is from the textbook PSSC Physics, which has
a blanket permission for free use after 1970.

selves, just the feet.) If too many bubbles form on the plastic, wipe them off with a ruler.

Make sure the straight-filament bulb in the light source is rotated so that when you look in through the hole, you are looking along the length of the filament. This way the lamp acts like a point source of light above the tank. To test that it's oriented correctly, check that you can cast a perfectly sharp image of the tip of a pen.

The light source is intended to be clamped to the wooden post, but I've found that that works very poorly, since the clamp doesn't hold it firmly enough. Instead, clamp the light source to the lab bench's lip or its leg. Turn it on. Put the butcher paper on the floor under the tank. If you make ripples in the water, you should be able to see the wave pattern on the screen.

The wave generator consists of a piece of wood that hangs by rubber bands from the two L-shaped metal hangers. There is a DC motor attached, which spins an intentionally unbalanced wheel, resulting in vibration of the wood. The wood itself can be used to make straight waves directly in the water, but in this experiment you'll be using the two little L-shaped pieces of metal with the yellow balls on the end to make two sources of circular ripples. The DC motor runs off of the DC voltage source, and the more voltage you supply, the faster the motor runs.

Start just by sticking one little L-shaped arm in the piece of wood, and observing the circular wave pattern it makes. Now try two sources at once, in neighboring holes. Pick a speed (frequency) for the motor that you'll use throughout the experiment — a fairly low speed works well. Measure the angular spacing of the resulting diffraction pattern for several values of the spacing, d, between the two sources of ripples.

Use the methods explained in Appendix 5 and look for any kind of a power law relationship for the dependence of the angular spacing on d.

3   Standing Waves

Apparatus

string
weights (in lab benches)
1 g slotted weights (in tray)
pulley
vibrator
paperclips
metersticks
butcher paper
scissors
weight holders
digital balances

Goals

Observe the resonant modes of vibration of a string.

Find how the speed of waves on a string depends on the tension in the string.

Introduction

The Greek philosopher Pythagoras is said to have been the first to observe that two plucked strings sounded good together when their lengths were in the proportion of two small integers. (This is assuming the strings are of the same material and under the same tension.) For instance, he thought a pleasant combination of notes was produced when one string was twice the length of the other, but that the combination was unpleasant when the ratio was, say, 1.4 to 1 (like the notes B and F). Although different combinations of notes are used in different cultures and different styles of music, there is at least some scientific justification for Pythagoras' statement. We now know that a plucked string does not just vibrate at a single frequency but simultaneously at a whole series of frequencies f_1, $2f_1$, $3f_1$,... These frequencies are called the harmonics. If one string is twice the length of the other, then its lowest harmonic is at half the frequency of the other string's, and its harmonics coincide with the odd-numbered harmonics of the other string. If the ratio is 1.4 to 1, however, then there is essentially no regular relationship between the two sets of frequencies, and many of the harmonics lie close enough in frequency to produce unpleasant beats.

Setup

The apparatus allows you to excite vibrations at a fixed frequency of f (twice the frequency of the alternating current from the wall that runs the vibrator). Since the point of the lab is to determine a proportionality, any constant factor, such as f, can be discarded.

The tension in the string can be controlled by varying the weight.

You may find it helpful to put a strip of white butcher paper behind the black string for better visual contrast.

It's important to get the vibrator set up properly along the same line as the string, not at an angle.

If there's a loud buzzing, try moving the vibrator so that it doesn't touch the bracket holding it.

Observations

Observe as many patterns ("modes") of vibration as you can. Each mode can be labeled by N, the number of humps or half-wavelengths. You will probably not be able to observe the fundamental ($N = 1$) because it would require too much weight. In each case, you will want to fine-tune the weight to get as close as possible to the middle of the resonance, where the amplitude of vibration is at a maximum. When you're close to the peak of a resonance, an easy way to tell whether to add or remove weight is by gently pressing down or lifting up on the weights with your finger to see whether the amplitude increases or decreases.

For large values of N, you may find that you need to use a paperclip instead of the weight holder, in order to make the mass sufficiently small. Keep in mind, however, that you won't really improve the quality of your data very much by taking data for very high values of N, since the 1-gram precision with which

you can locate these resonances results in a poor relative precision compared to a small weight.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

If you were assigned problem 2-12 in Modern Physics, *skip the prelab.*

P1 How is the tension in the string, T, related to the mass of the hanging weight?

P2 The figure below shows the $N = 1$, 2, and 3 patterns of vibration. Suppose the length of the string is one meter. In each case, find the wavelength.

P3 Generalize your numerical results from P2 to give a general equation for λ in terms of N and L, the length of the string. Check its units, and check that it recovers the special cases done numerically in P2.

P4 How can the velocity of the waves be determined if you know the frequency, f, the length of the string, L, and the number of humps, N?

Self-Check

Do your analysis in lab.

Analysis

Use the graphing technique given in appendix 5 to see if you can find a power-law relationship between the velocity of the waves in the string and the tension in the string. Note that you can omit constant factors without affecting the exponent. (Do not just try to find the correct power law in the textbook, because besides observing the phenomenon of resonance, the point of the lab is to prove experimentally what the power-law relationship is, and to test whether this is always a good approximation in real life.)

4   Resonances of Sound

Apparatus

wave generator (Pasco PI-8127, in lab benches) 1/group
speaker (Thornton) 1/group
100 mL graduated cylinder 1/group
Linux computers with digital oscilloscope software installed (see note above)
flexible whistling tube 1
tuning fork marked with frequency, mounted on a wooden box 1
aluminum rod, 3/4-inch dia, about 1 m long 2
wood block 1/group
thermometer 1

Goals

> Find the resonant frequencies of the air inside a cylinder by two methods.

> Measure the speeds of sound in air and in aluminum.

Introduction

In the womb, your first sensory experiences were of your mother's voice, and soon after birth you learned to distinguish the particular sounds of your parents' voices from those of strangers. The human ear-brain system is amazingly sophisticated in its ability to classify vowels and consonants, recognize people's voices, and analyze musical sound. Until the 19th-century investigations of Helmholtz, the whole process was completely mysterious. How could we so easily tell a cello from a violin playing the same note? A radio station in Chicago has a weekly contest in which jazz fanatics are asked to identify instrumentalists simply by their distinctly individual timbres — how is this possible?

Helmholtz found (using incredibly primitive nonelectronic equipment) that part of the answer lay in the relative strengths of the overtones. The psychological sensation of pitch is related to frequency, e.g., 440 Hz is the note "A." But a saxophonist playing the note "A" is actually producing a rich spectrum of frequencies, including 440 Hz, 880 Hz, 1320 Hz, and many other multiples of the lowest frequency, known as the fundamental. The ear-brain system perceives all these overtones as a single sound because they are

all multiples of the fundamental frequency. (The Javanese orchestra called the gamelan sounds strange to westerners partly because the various gongs and cymbals have overtones that are not integer multiples of the fundamental.)

One of the things that would make "A" on a clarinet sound different from "A" on a saxophone is that the 880 Hz overtone would be quite strong for the saxophone, but almost entirely missing for the clarinet. Although Helmholtz thought the relative strengths of the overtones was the whole story when it came to musical timbre, actually it is more complex than that, which is why electronic synthesizers still do not sound as good as acoustic instruments. The timbre depends not just on the general strength of the overtones but on the details of how they first build up (the attack) and how the various overtones fade in and out slightly as the note continues.

Why do different instruments have different sound spectra, and why, for instance, does a saxophone have an overtone that the clarinet lacks? Many musical instruments can be analyzed physically as tubes that have either two open ends, two closed ends, or one open end and one closed end. The overtones correspond to specific resonances of the air column inside the tube. A complete treatment of the subject is given in your textbook, but the basic principle is that the resonant standing waves in the tube must have an antinode (point of maximum vibration) at any closed end of the tube, and a node (point of zero vibration) at any open end.

Using the Wave Generator

The wave generator works like the amplifier of your stereo, but instead of playing a CD, it produces a sine wave whose frequency and amplitude you can control. By connecting it to a speaker, you can convert its electrical currents to sound waves, making a pure tone. The frequency of the sine wave corresponds to musical pitch, and the amplitude corresponds to loudness.

Setup

Plug the speaker into the wave generator. The banana plugs go in the plugs marked OUTPUT. The side of the banana plug converter with the tiny tab

marked GND should be the one that goes into the black output plug. Set the frequency to something audible. Turn the amplitude knob up until you hear a sound.

The wave generator and the speaker are not really designed to work together, so if you leave the volume up very high for a long time, it is possible to blow the speaker or damage the wave generator. Also, the sine waves are annoying when played continuously at loud volumes!

Preliminary Observations

Observations

This lab has three parts, A, B, and C. It is not really possible for more than one group to do part A in the same room, both because their sounds interfere with one another and because the noise becomes annoying for everyone. Your instructor will probably have three groups working on part A at one time, one group in the main room, one in the small side room, and one in the physics stockroom. Meanwhile, the other groups will be doing parts B and C.

A Direct Measurement of Resonances by Listening

Set up the graduated cylinder so its mouth is about 3 mm from the center of the speaker grille. Find as many frequencies as possible at which the cylinder resonates. When you sweep through those frequencies, the sound becomes louder. To make sure you're really hearing a resonance of the cylinder, make sure to repeat each observation with the cylinder removed, and make sure the resonance goes away. For each resonance, take several measurements of its frequency — if you are careful, you can pin it down to within ±10 Hz or so. You can probably speed up your search significantly by calculating approximately where you expect the resonances to be, then looking for them.

B Electronic Measurement of Resonances of an Air Column

The resonances of the air column in a cylinder can also be excited by a stream of air flowing over an opening, as with a flute. In this part of the lab, you will excite resonances of a long, flexible plastic tube by grabbing it at one end and swinging it in a circle. The frequency of the sound will be determined electronically. Note that your analysis for these resonances will be somewhat different, since the tube is open at both ends, and it therefore has different

patterns of resonances from the graduated cylinder, which was only open at one end.

To measure the frequency, you will use a computer to analyze the sound. The Linux computers are the ones with the right hardware and software. As a warmup before attempting the actual measurements with the whistling tube, try the following. First, start up the program if nobody else has already done so. It is called "scope," and you can run it by double-clicking on its icon on the desktop. In real time, the program will monitor the sound coming into the microphone, and display a graph of loudness versus frequency. Try whistling. The frequency at which you whistled should show up as a prominent peak.

Next you need to verify that you can actually measure a known frequency and reproduce its value. Besides letting you practice using the software, this is important because I've seen some cheap computer sound input chips that produce frequencies that are off by large amounts, like 10%. Put the microphone near the wooden box that the tuning fork is mounted on and hit the tuning fork gently with the rubber mallet. To get an accurate frequency measurement, you need to zoom in on the peak. To do this, click on the peak to get an extreme close-up. To zoom back out, click on the graph again. (The Zoom In button doesn't let you get close enough, so don't use it for this purpose.) When you get the graph you want to see, you can freeze it by clicking on the Freeze/Go button.

Once you have done these warmups, you are ready to analyze the sound from the whistling tube. You only need to analyze data from one frequency, although if you're not sure which mode you produced, it may be helpful to observe the pattern of the frequencies. (If you guess wrong about which mode it was, you'll find out, because the value you extract for the speed of sound will be way off.)

C The Speed of Sound in Aluminum

The speed of sound in a solid is much faster than its speed in air. In this part of the lab, you will extract the speed of sound in aluminum from a measurement of the lowest resonant frequency of a solid aluminum rod. You will use the computer for an electronic measurement of the frequency, as in part B.

Grab the rod with two fingers exactly in the middle, hold it vertically, and tap it on the lab bench. You will hear two different notes sounding simultaneously. A quick look at their frequencies shows that they are not in a 2:1 ratio as we would expect based on our experiences with symmetric wave pat-

terns. This is because these two frequencies in the rod are actually two different types of waves. The higher note is produced by longitudinal compression waves, which means that an individual atom of aluminum is moving up and down the length of the rod. This type of wave is analogous to sound waves in air, which are also longitudinal compression waves. The lower note comes from transverse vibrations, like a vibrating guitar string. In the transverse vibrations, atoms are moving from side to side, and the rod as a whole is bending.

If you listen carefully, you can tell that the transverse vibration (the lower note) dies out quickly, but the longitudinal mode keeps going for a long time. That gives you an easy way to isolate the longitudinal mode, which is the one we're interested in; just wait for the transverse wave to die out before you freeze the graph on the computer.

Identifying the mode of vibration
The rod is symmetric, so we expect its longitudinal wave patterns to be symmetric, like those of the whistling tube. The rod is different, however, because whereas we can excite a variety of wave patterns in the tube by spinning it at different speeds, we find we only ever get one frequency from the rod by tapping it at its end: it appears that there is only one logitudinal wave pattern that can be excited strongly in the rod by this method. The problem is that we then need to infer what the pattern is.

Since you hold the rod at its center, friction should very rapidly damp out any mode of vibration that has any motion at the center. Therefore there must be a node at the center. We also know that at the ends, the rod has nothing to interact with but the air, and therefore there is essentially no way for any significant amount of wave energy to leak out; we therefore expect that waves reaching the ends have 100% of their energy reflected. Since energy is proportional to the square of amplitude, this means that a wave with unit amplitude can be reflected from the ends with an amplitude of either $R = +1$ (100% uninverted reflection) or -1 (100% inverted). In the $R = -1$ case, the reflected wave would cancel out the incident wave at the end of the rod, and we would have a node at the end, as in lab 3. In the $R = +1$ case, there would be an antinode. But when you tap the end of the rod on the floor, you are evidently exciting wave motion by moving the end, and it would not be possible to excite vibrations by this method if the vibrations had no motion at the end. We therefore conclude that the rod's pattern of vi-

bration must have a node at the center, and antinode at the ends.

There is an infinite number of possible wave patterns of this kind, but we will assume that the pattern that is excited strongly is the one with the longest wavelength, i.e., the only node is at the center, and the only antinodes are those at the ends.

If you feel like it, there are a couple of possible tests you can try to do to check whether this is the right interpretation. One is to see if you can detect any other frequencies of longitudinal vibration that are excited weakly. Another is to predict where the other nodes would be, if there were more than one, and then see if the vibration is killed by touching the rod there with your other hand; if there is a node there, touching it should have no effect.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

P1 *Skip this question if you were assigned problem 5-5 in* Modern Physics. Find an equation to predict the frequencies of the resonances in parts A and B. Note that they will not be the same equations, since one tube is symmetric and the other is asymmetric. For the symmetric one, you could simply staple your prelab from lab 3 onto the back. For the asymmetric one, go through the same process that you were explicitly led through in the prelab for lab 3. This process starts with drawing the first few wave patterns. It's not necessary to define a variable N analogous to the one defined in lab 3, and in fact it's not obvious how one could define such a thing as the "number of humps" in the asymmetric case. Instead of giving a single formula with an N in it, it's fine to list expressions for the first few frequencies, showing the pattern.

Self-Check

Extract the speed of sound in air from either part A or part B, without error analysis, and make sure you get something reasonable. We don't necessarily expect it to be exactly the same as a standard value, because it depends on conditions, especially the temperature.

Analysis

Make a graph of wavelength versus period for the resonances of the graduated cylinder, check whether it looks like it theoretically should, and if so, find the speed of sound from its slope, with error bars, as discussed in appendix 4.

Analyze part B either by using the same technique (if you took data for multiple frequencies) or just by solving algebraically for the speed of sound.

The effective length of the cylinder in part A should be increased by 0.4 times its diameter to account for the small amount of air beyond the end that also vibrates. For part B, where the whistling tube is open at both ends, you should add 0.8 times its diameter.

When estimating error bars from part B, you may be tempted to say that it must be perfectly accurate, since it's being done by a computer. Not so! The graph only has a certain frequency resolution, and in addition, the peak may be a little ragged.

Extract the speed of sound in aluminum from your data in part C, including error bars.

5   Refraction and Images

Apparatus

plastic box
propanol (1 liter/group, to be reused)
laser
spiral plastic tube and fiber optic cable for demonstrating total internal reflection
ruler
protractor
butcher paper
funnel

incident medium

transmitting medium

Goals

Test whether the index of refraction of a liquid is proportional to its density.

Observe the phenomena of refraction and total internal reflection.

Locate a virtual image in a plastic block by ray tracing, and compare with the theoretically predicted position of the image.

Introduction

Without the phenomenon of refraction, the lens of your eye could not focus light on your retina, and you would not be able to see. Refraction is the bending of rays of light that occurs when they pass through the boundary between two media in which the speed of light is different.

Refraction occurs for the following reason. Imagine, for example, a beam of light entering a swimming pool at an angle. Because of the angle, one side of the beam hits the water first, and is slowed down. The other side of the beam, however, gets to travel in air, at its faster speed, for longer, because it enters the water later — by the time it enters the water, the other side of the beam has been limping along through the water for a little while, and has not gotten as far. The wavefront is therefore twisted around a little, in the same way that a marching band turns by having the people on one side take smaller steps.

Quantitatively, the amount of bending is given by Snell's law:

$$n_i \sin \theta_i = n_t \sin \theta_t,$$

where the subscript i refers to the incident light and

incident medium, and t refers to the transmitted light and the transmitting medium. This relation can be taken as defining the quantities n_i and n_t, which are known as the indices of refraction of the two media. Note that the angles are defined with respect to the normal, i.e., the imaginary line perpendicular to the boundary.

Also, not all of the light is transmitted. Some is reflected — the amount depends on the angles. In fact, for certain values of n_i, n_t, and θ_i, there is no value of θ_t that will obey Snell's law ($\sin \theta_t$ would have to be greater than one). In such a situation, 100% of the light must be reflected. This phenomenon is known as total internal reflection. The word internal is used because the phenomenon only occurs for $n_i > n_t$. If one medium is air and the other is plastic or glass, then this can only happen when the incident light is in the plastic or glass, i.e., the light is trying to escape but can't. Total internal reflection is used to good advantage in fiber-optic cables used to transmit long-distance phone calls or data on the internet — light traveling down the cable cannot leak out, assuming it is initially aimed at an angle close enough to the axis of the cable.

Although most of the practical applications of the phenomenon of refraction involve lenses, which have curved shapes, in this lab you will be dealing almost exclusively with flat surfaces.

Preliminaries

Check whether your laser's beam seems to be roughly parallel.

Observations

A Index of refraction of alcohol

The index of refraction is sometimes referred to as the optical density. This usage makes sense, because when a substance is compressed, its index of refraction goes up. In this part of the lab, you will test whether the indices of refraction of different liquids are proportional to their mass densities. Water has a density of 1.00 g/cm^3 and an index of refraction of 1.33. Propanol has a density of 0.79 g/cm^3. You will find out whether its index of refraction is lower than water's in the same proportion. The idea is to pour some alcohol into a transparent plastic box and measure the amount of refraction at the interface between air and alcohol.

Make the measurements you have planned in order to determine the index of refraction of the alcohol. The laser and the box can simply be laid flat on the table. Make sure that the laser is pointing towards the wall.

B Total internal reflection

Try shining the laser into one end of the spiral-shaped plastic rod. If you aim it nearly along the axis of the cable, none will leak out, and if you put your hand in front of the other end of the rod, you will see the light coming out the other end. (It will not be a well-collimated beam any more because the beam is spread out and distorted when it undergoes the many reflections on the rough and curved inside the rod.)

There's no data to take. The point of having this as part of the lab is simply that it's hard to demonstrate to a whole class all at once.

C A virtual image

Pour the alcohol back into the container for reuse, and pour water into the box to replace it.

Pick up the block, and have your partner look sideways through it at your finger, touching the surface of the block. Have your partner hold her own finger next to the block, and move it around until it appears to be as far away as your own finger. Her brain achieves a perception of depth by subconsciously comparing the images it receives from her two eyes. Your partner doesn't actually need to be able to see her own finger, because her brain knows how to position her arm at a certain point in space. Measure the distance d_i, which is the depth of the image of your finger relative to the front of the block.

Next we will use the laser to simulate the rays com-

image of finger

Light rays spreading out from the finger by diffuse reflection. The emerging rays all appear to have come from a point inside the block.

laser

Simulating one of the rays using the laser.

ing from the finger, as shown in the figure. Shine the laser at the point where your finger was originally touching the block, observe the refracted beam, and draw it in. Repeat this whole procedure several times, with the laser at a variety of angles. Finally, extrapolate the rays leaving the block back into the block. They should all appear to have come from the same point, where you saw the virtual image. You'll need to photocopy the tracing so that each person can turn in a copy with his or her writeup.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

You should have already done the laser safety checklist for the relativity lab.

P1 Laser beams are supposed to be very nearly parallel (not spreading out or contracting to a focal point). Think of a way to test, roughly, whether this is true for your laser.

P2 Plan how you will determine the index of refraction in part A.

P3 *Skip this question if you were assigned problem 5-22 in* Modern Physics*.* You have complete freedom to choose any incident angle you like in part A. Discuss what choice would give the highest possible precision for the measurement of the index of refraction.

Analysis

Using your data for part A, extract the index of refraction of propanol, with error bars. Test the hypothesis that the index of refraction is proportional to the density in the case of water and propanol.

Using trigonometry and Snell's law, make a theoretical calculation of d_i. You'll need to use the small-angle approximation $\sin\theta \approx \tan\theta \approx \theta$, for θ measured in units of radians. (For large angles, i.e. viewing the finger from way off to one side, the rays will not converge very closely to form a clear virtual image.)

Explain your results in part C and their meaning.

Compare your three values for d_i : the experimental value based on depth perception, the experimental value found by ray-tracing with the laser, and the theoretical value found by trigonometry.

6   Geometric Optics

Apparatus

optical bench 1/group
converging lens, $f = 50$ cm, 5 cm diam
1/group converging lens, $f = 5$ cm 1/group
lamp and arrow-shaped mask 1/group
frosted glass screen 1/group

Goals

Observe a real image formed by a convex lens, and determine its focal length.

Construct a telescope and measure its angular magnification.

Introduction

The credit for invention of the telescope is disputed, but Galileo was probably the first person to use one for astronomy. He first heard of the new invention when a foreigner visited the court of his royal patrons and attempted to sell it for an exorbitant price. Hearing through second-hand reports that it consisted of two lenses, Galileo sent an urgent message to his benefactors not to buy it, and proceeded to reproduce the device himself. An early advocate of simple scientific terminology, he wanted the instrument to be called the "occhialini," Italian for "eye-thing," rather than the Greek "telescope."

His astronomical observations soon poked some gaping holes in the accepted Aristotelian view of the heavens. Contrary to Aristotle's assertion that the heavenly bodies were perfect and without blemishes, he found that the moon had mountains and the sun had spots (the marks on the moon visible to the naked eye had been explained as optical illusions or atmospheric phenomena). This put the heavens on an equal footing with earthly objects, paving the way for physical theories that would apply to the whole universe, and specifically for Newton's law of gravity. He also discovered the four largest moons of Jupiter, and demonstrated his political savvy by naming them the "Medicean satellites" after the powerful Medici family. The fact that they revolved around Jupiter rather than the earth helped make more plausible Copernicus' theory that the planets did not revolve around the earth but around the sun.

Galileo's ideas were considered subversive, and many people refused to look through his telescope, either because they thought it was an illusion or simply because it was supposed to show things that were contrary to Aristotle.

The figure on the next page shows the simplest refracting telescope. The object is assumed to be at infinity, so a real image is formed at a distance from the objective lens equal to its focal length, f_o. By setting up the eyepiece at a distance from the image equal to its own focal length, f_E, light rays that were parallel are again made parallel.

The point of the whole arrangement is angular magnification. The small angle α_1 is converted to a large α_2. It is the small angular size of distant objects that makes them hard to see, not their distance. There is no way to tell visually whether an object is a thirty meters away or thirty billion. (For objects within a few meters, your brain-eye system gives you a sense of depth based on parallax.) The Pleiades star cluster can be seen more easily across many light years than Mick Jagger's aging lips across a stadium. People who say the flying saucer "looked as big as an aircraft carrier" or that the moon "looks as big as a house" don't know what they're talking about. The telescope does not make things "seem closer" — since the rays coming at your eye are parallel, the final virtual image you see is at infinity. The angular magnification is given by

$$M_A = \alpha_2/\alpha_1$$

(to be measured directly in this lab)

$$M_A = f_o/f_E$$

(theory).

Observations

A   Focal length of the lenses

In this part of the lab, you'll accurately determine the focal lengths of the two lenses being used for the telescope. They are poorly quality-controlled, and I've found that the labeled values are off by as much as 10%. If you're only doing part A of the lab, then you will only determine the focal length of one lens, which is given to you as an unknown.

A refracting telescope

Start with the short-focal-length lens you're going to use as your eyepiece. Use the lens to project a real image on the frosted glass screen. For your object, use the lamp with the arrow-shaped aperture in front of it. Make sure to lock down the parts on the optical bench, or else they may tip over and break the optics!

For the long-focal-length lens you're going to use as your objective, you will probably be unable to do a similar determination on a one-meter optical bench. Improvising a similar setup without the bench will still give you a much more accurate value than the one written on the label.

A careful measurement here pays off later by making the focus in part B much easier to find. This is especially true for the longer-focal-length lens. To improve the quality of your result, do the kind of thing they do at the optometrist — "which is better, 1 or 2?" Have several people do independent determinations of the best focus.

B   The telescope

Use your optical bench and your two known lenses to build a telescope. Since the telescope is a device for viewing objects at infinity, you'll want to take it outside.

The best method for determining the angular magnification is to observe the same object with both eyes open, with one eye looking through the telescope and one seeing the object without the telescope. Good precision can be obtained, for example, by looking at a large object like a coke machine, and determining that a small part of it, whose size you can measure with a ruler, appears, when magnified, to cover some larger part of it, which you can also measure. The figures on p. 25 show a simulation of what the superimposed images should look like and of how it would look if the telescope is not yet adjusted quite correctly.

Your brain is not capable of focusing one eye at one distance, and the other at another distance. Therefore it's important to get your telescope adjusted precisely so that the image is at infinity. You can do this by focusing your naked eye on a distant object, and then moving the objective until the image pops into focus in the other eye. Theoretically this would be accomplished simply by setting the lenses at the distance shown in the diagram, but in reality, a small amount of further adjustment is necessary, because of the uncertainty in the measured focal lengths.

A good quick test of the focus is to pick someone who's nearsighted and see if they can focus on the image without their glasses on. If they can, then the image is not at infinity, because nearsighted people can't focus on an image that's at infinity.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

P1   *Skip this question if you were assigned problem 11-17 from Modern Physics,.* In part A, do you want the object to be closer to the lens than the lens' focal length, exactly at a distance of one focal length, or farther than the focal length? What about the screen?

P2   Plan what measurements you will make in part A and how you will use them to determine the lenses' focal length.

P3   *Skip this question if you are only doing part A of the lab.* It's disappointing to construct a telescope with a very small magnification. Given a selection of lenses, plan how you can make a telescope with the greatest possible angular magnification.

Analysis

Determine the focal length of the two lenses, with error bars.

Find the angular magnification of your telescope from your data, with error bars, and compare with theory. Do they agree to within the accuracy of the measurement? Give a probabilistic interpretation, as in the example in appendix 2. See the example at the end of appendix 3 of how to test for equality between two numbers that both have error bars.

The goal: magnified and unmagnified images superimposed and in focus.

Your aim is wrong.

The telescope's image is not at infinity. Your brain focuses your eye on the dog.

The telescope's image is not at infinity. Your brain focuses your eye on the tree.

These images can be viewed in a web browser at lightandmatter.com/lec/telescope . See the file NOTES in the source code for photo credits.

7   Wave Optics

Apparatus

helium-neon laser
1/group optical bench with posts & holders 1/group
high-precision double slits 1/group
rulers
meter sticks
tape measures
butcher paper
black cloths for covering light sources

Goals

Observe evidence for the wave nature of light.

Determine the wavelength of the red light emitted by your laser, by measuring a double-slit diffraction pattern. (The part of the spectrum that appears red to the human eye covers quite a large range of wavelengths. A given type of laser, e.g., He-Ne or solid-state, will produce one very specific wavelength.)

Determine the approximate diameter of a human hair, using its diffraction pattern.

Introduction

Isaac Newton's epitaph, written by Alexander Pope, reads:

Nature and Nature's laws lay hid in night.

God said let Newton be, and all was light.

Notwithstanding Newton's stature as the greatest physical scientist who ever lived, it's a little ironic that Pope chose light as a metaphor, because it was in the study of light that Newton made some of his worst mistakes. Newton was a firm believer in the dogma, then unsupported by observation, that matter was composed of atoms, and it seemed logical to him that light as well should be composed of tiny particles, or "corpuscles." His opinions on the subject were so strong that he influenced generations of his successors to discount the arguments of Huygens and Grimaldi for the wave nature of light. It was not until 150 years later that Thomas Young demonstrated conclusively that light was a wave.

Young's experiment was incredibly simple, and could probably have been done in ancient times if some savvy Greek or Chinese philosopher had only thought of it. He simply let sunlight through a pinhole in a window shade, forming what we would now call a coherent beam of light (that is, a beam consisting of plane waves marching in step). Then he held a thin card edge-on to the beam, observed a diffraction pattern on a wall, and correctly inferred the wave nature and wavelength of light. Since Roemer had already measured the speed of light, Young was also able to determine the frequency of oscillation of the light.

Today, with the advent of the laser, the production of a bright and coherent beam of light has become as simple as flipping a switch, and the wave nature of light can be demonstrated very easily. In this lab, you will carry out observations similar to Young's, but with the benefit of hindsight and modern equipment.

Observations

A   Determination of the wavelength of red light

Set up your laser on your optical bench. You will want as much space as possible between the laser and the wall, in order to let the diffraction pattern spread out as much as possible and reveal its fine details.

Tear off two small scraps of paper with straight edges. Hold them close together so they form a single slit. Hold this improvised single-slit grating in the laser beam and try to get a single-slit diffraction pattern. You may have to play around with different widths for the slit. No quantitative data are required. This is just to familiarize you with single-slit diffraction.

Make a diffraction pattern with the double-slit grating. See what happens when you hold it in your hand and rotate it around the axis of the beam.

The diffraction pattern of the double-slit grating consists of a rapidly varying pattern of bright and dark bars, with a more slowly varying pattern superimposed on top. (See the figure two pages after this page.) The rapidly varying pattern is the one that is numerically related to the wavelength, λ, and the distance between the slits, d, by the equation

$$\Delta\theta = \lambda/d,$$

This spacing, Δθ, is the one you want to measure. It is related to d, the center-to-center distance between the slits, by Δθ=λ/d.

If your screen is too close to the slits, the finer pattern may be invisible, and this spacing may be all you see. This larger angular spacing is related to the width of the slits, not to d.

A double-slit diffraction pattern.

where θ is measured in radians. To make sure you can see the fine spacing, put your slits across the room from the wall. To make it less likely that someone will walk through the beam and get the beam in their eye, put some of the small desks under the beam. The slit patterns we're using actually have three sets of slits, with the following dimensions:

	w (mm)	d (mm)
A	.12	.6
B	.24	.6
C	.24	1.2

The small value of d is typically better, for two reasons: (1) it produces a wider diffraction pattern, which is easier to see; (2) it's easy to get the beam of the laser to cover both slits.

If your diffraction pattern doesn't look like the one in the figure on the following page, typically the reason is that you're only hitting one slit with the beam (in which case you get a single-slit diffraction pattern), or you're not illuminating the two slits equally (giving a funny-looking pattern with little dog-bones and things in it).

As shown in the figure below, it is also possible to have the beam illuminate only *part* of each slit, so that the slits act effectively as if they had a smaller value of d. The beam spreads as it comes out of the laser, so you can avoid this problem by putting it fairly far away from the laser (at the far end of the optical bench).

good   sing-slit diffraction   acts like a smaller *d*

Think about the best way to measure the spacing of the pattern accurately. Is it best to measure from a bright part to another bright part, or from dark to dark? Is it best to measure a single spacing, or take several spacings and divide by the number to find what one spacing is?

Determine the wavelength of the light, in units of nanometers. Make sure it is in the right range for red light.

Check that the $\Delta\theta$ you obtain is in the range predicted in prelab question P1. In the past, I've seen cases where groups got goofy data, and I suspect that it was because they were hitting a place on the slits where there was a scratch, bump, or speck of dust.

B   Diameter of a human hair

Pull out one of your own hairs, hold it in the laser beam, and observe a diffraction pattern. It turns out that the diffraction pattern caused by a narrow obstruction, such as your hair, has the same spacing as the pattern that would be created by a single slit whose width was the same as the diameter of your hair. (This is an example of a general theorem called Babinet's principle.) Measure the spacing of the diffraction pattern. (Since the hair's diameter is the only dimension involved, there is only one diffraction pattern with one spacing, not superimposed fine and coarse patterns as in part A.) Determine the diameter of your hair. Make sure the value you get is reasonable, and compare with the order-of-magnitude guess you made in your prelab writeup.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

Read the safety checklist.

P1   Look up the approximate range of wavelengths that the human eye perceives as red. With $d = 0.6$ mm, predict the lowest and highest possible values of $\Delta\theta$ that could occur with red light.

P2   It is not practical to measure $\Delta\theta$ directly using a protractor. Suppose that a lab group finds that 27 fringes extend over 29.7 cm on their butcher paper, which is on a wall 389 cm away from the slits. They calculate $\Delta\theta = \tan^{-1}(29.7 \text{ cm}/(27 \times 389 \text{ cm}))$ $= 2.83 \times 10^{-3}$ rad. Simplify this calculation using a small-angle approximation, and show that the resulting error is negligible.

P3   Make a rough order-of-magnitude guess of the diameter of a human hair.

Analysis

Determine the wavelength of the light and the diameter of the hair, with error bars.

8 The Michelson Interferometer

Apparatus

Michelson interferometer 1/group
Na discharge tube 1/group
tools inside drawer ....1 set/group
2 × 4 piece of wood 3/group
helium-neon laser 3
filters
Vernier calipers

Goals

Observe the basic operation of a Michelson interferometer.

Investigate the feasibility of doing other measurements with the instrument.

Introduction

The Michelson interferometer is a device for measuring the wavelength of light, used most famously in the Michelson-Morley experiment of 1887, which was later interpreted as disproving the existence of the luminiferous aether and supporting Einstein's theory of special relativity.

As shown in the figure, the idea is to take a beam of light from the source, split it into two perpendicular beams, send it to two mirrors, and then recombine the beams again. If the two light waves are in phase when recombined, they will reinforce, but if they are out of phase, they will cancel. Since the two waves originated from the splitting of a single wave,

the only reason they would be out of phase would be if the lengths of the two arms of the apparatus were unequal. Mirror A is movable, and the distance through which it moves can be controlled and measured extremely accurately using a micrometer connected to the mirror via a lever. If mirror A is moved by distance equal to a quarter of a wavelength of the light, the total round-trip distance traveled by the wave is changed by half a wavelength, which switches from constructive to destructive interference, or vice versa. Thus if the mirror is moved by a distance d, and you see the light go through n complete cycles of appearance and disappearance, you can conclude that the wavelength of the light was $\lambda = 2d/n$.

Classic uses of a small Michelson interferometer are to measure small differences in indices of refraction, compare physical length standards to atomic ones, and compare atomic standards to one another.

Apparatus

To make small and accurate adjustments of the mirror easier to do, the micrometer is connected to it through a lever that reduces the amount of movement by a factor k, which is supposed to be exactly equal to 5; the micrometer reads the bigger distance $D = kd$ that it actually travels itself, so the wavelength is $\lambda = 2D/kn$.

When you look through the device, you see a certain field of view. Within that field of view are mirror A and the image B′ of mirror B, which is nearly superimposed on A. The path-length difference between the two waves can be thought of as the distance between A and B′. The field of view covers a certain small but finite range of angles, forming a narrow cone with its vertex at your eye. For rays forming a small angle α relative to the axis, the effect is to multiply the path-length difference by a factor of $\cos\alpha$. If the two arms of the interferometer are exactly equal, then the path-length difference between them is zero, and multiplying it by $\cos\alpha$ has no effect. But when the arms are unequal, the condition for constructive or destructive interference becomes dependent on α, and the result is that you see a bull's-eye pattern of concentric circular fringes. When you turn the knob, these fringes expand or contract, appearing from or disappearing into the center. The more equal the two arms, the

smaller the number of fringes.

The setup also includes a compensating plate, described in the Wikipedia article on the Michelson interferometer.

Setup

Turn on the sodium discharge tube, and let it warm up until it's yellow.

Remove the drawer from the box, and take out the tool kit. Unscrew the screws on the bottom of the box that lock the interferometer to the floor of the box, and *very carefully* take the instrument out of the box. Screw the two aluminum legs into the bottom of the interferometer, and lay a piece of wood flat under the third leg, which is a threaded rod; this makes the apparatus level.

Place the discharge tube near the entrance window of the apparatus. If you look through the viewing window, you will see the image of the tube itself, reflected through the mirrors. To make this into a uniform circle of light, place the ground glass screen (inside the bag of tools) in the bracket at the entrance window.

Mirror B needs to be perfectly perpendicular to mirror A, and its vertical plane needs to be matched to mirror A's. This is adjusted using the knobs on mirror B, one for vertical adjustment and one for horizontal. A rough initial adjustment can be done by aligning the two images of the circular entrance window. You can then hang the metal pointer (from the bag of tools) on the top of the ground glass screen, and do a better adjustment so that the two images of the pointer's tip coincide. You should now see a set of very fine concentric circular interference fringes, centered on a point outside of the field of view. The final, fine adjustment is obtained by bringing the center of this pattern to the center of the field of view.

A Basic Operation

To get an idea of the basic operation of the instrument, you'll perform a rough measurement of the wavelength of the light emitted by a discharge tube. This is not really a task for which the spectrometer is well suited for giving high-precision results, but you should be able to approximately verify the known wavelength.

The micrometer has a millimeter scale running from 0 to 25 mm, with half-millimeter divisions on the bottom. To take a reading on it, first read the number of millimeters and half-millimeters based on where the edge of the cylindrical rotating part lies on this scale. Then add on the reading from the vernier scale that runs around the circumference of the rotating part, which runs from 0.000 to 0.500 mm. You should be able to estimate to the nearest thousandth of a millimeter (tenth of a vernier division).

While looking at the interference fringes, turn the knob on the micrometer. You will see them either expand like smoke rings, or contract and disappear into the center, depending on which way you turn the knob. Rotate the knob while counting at least 10 or 20 fringes, and record the two micrometer readings before and after. The difference between these is D. It helps if you prop your head on the table, and move the micrometer knob smoothly and continuously. Moving your head disturbs the pattern, and halting the micrometer knob tends to cause backlash that confuses the count of fringes by plus or minus one.

For the reasons described in part C, the contrast fades in and out as you move mirror B. This is probably what limits the number of fringes you can count.

Check that you get approximately the right wavelength (check on Wikipedia for the one you're using).

Exploring Other Possibilities with the instrument

In spring and fall of 2013, I had my students try a variety of things with the interferometers. I want to find a measurement that isn't too difficult and that uses the instrument for a purpose that it's well suited to. The best prospects seem to be the following. Your group can pick either one, but I think B is easier and probably more of a classic application.

B Unknown index of refraction

The Michelson interferometer is unusual in its ability to work with white light. This is because it's possible to get into the position where the lengths of the two arms are equal, so that constructive interference occurs at the center of the bull's eye for all wavelengths. At any other position, we get a random mixture of constructive and destructive interference for all the different wavelengths that are present. In fact, this allows the equal-arm condition to be determined extremely accurately. It's hard to think of any other technique that allows two large distances to be made equal to within micrometers or nanometers!

Using this technique, one can accurately determine the unknown index of refraction of thin piece of some material such as a glass microscope slide. The idea is to equalize the arms, then insert the unknown and readjust the interferometer in order to restore the interference pattern. In the new condition, it's not the lengths of the two arms that are equal but their *optical* lengths, i.e., the time it takes light to travel along them.

In fall 2013 I had one group that tried this technique and seemed to succeed in measuring the index of refraction of a thin glass microscope slide; the index of refraction came out to be in a reasonable range for glass, although different types of glass cover a wide range. Unforseen difficulties made spring 2014 a disaster. In the future I would like to see if most or all of the groups can get this to work, and reproduce each other's results for slides made of the same type of glass. If that works, then I would like to refine the technique to work with something else besides the microscope slides that would either have some intrinsic interest or allow a high-precision measurement. A possibility is to use cuvettes with distilled water and a salt solution, and measure the difference in index of refraction between the two.

C Spacing of the sodium doublet

Essentially the only way scientists have of getting detailed, high-precision information about the structure of atoms and nuclei is by measuring the set of wavelengths of light they emit. As you'll see later in the course, these wavelengths can be related to the sizes of the "quantum jumps" between different energy levels of the atom. Each of these wavelengths is referred to as a "line," because in some types of spectrometers that's what they look like, and the whole set of lines is called the spectrum of that atom. An unfortunate fact of life for the spectroscopist is that the spectrum often contains "doublets," meaning lines that are very close together and hard to distinguish, and in fact the yellow line in the sodium spectrum that you measured in part A is really a doublet whose members differ in wavelength by a fraction of a nanometer.

The Michelson interferometer is well adapted to a certain sneaky trick for measuring the difference in wavelength between two lines in a doublet, even when the difference is extremely small, as it is for sodium. The idea is that if the lines are roughly equal in intensity, then the bull's-eye pattern you see is actually two bull's-eye patterns superimposed on top of one another. If you move the micrometer to a random position, then it's a matter of chance how well these two patterns agree. They could happen to agree perfectly, in which case the bull's eye would look just as distinct as if there were only one line, but they could also disagree perfectly, so that there would no contrast at all between light and dark. The ancient lab manual says: "Loosen the carriage lock screw and move the carriage by hand. Note that the fringes pass alternately from a condition of high contrast to one of almost complete disappearance. With the micrometer screw set near one of its extreme limits and the carriage at one of the conditions of almost complete disappearance, tighten the carriage lock screw." Then you're supposed to move to another condition of minimum contrast and take another micrometer reading.

The analysis then works like this. Let the wavelengths constituting the doublet be be λ_1 and λ_2, differing by $\Delta\lambda = \lambda_2 - \lambda_1$. Let d_1 and d_2 be two successive path-length differences at which there is maximum contrast. Traveling a distance Δd between these two values of d, we pass through $N + 1$ cycles worth of the shorter λ_1 and N cycles for the longer λ_2, so $2\Delta d = (N + 1)\lambda_1 = N\lambda_2$. Algebra gives $\Delta\lambda = \lambda_1\lambda_2/2\Delta d$. The fractional precision with which we can measure Δd equals the fractional precision with which you can measure $\Delta\lambda$, which is a huge win, because $\Delta\lambda$ is very small.

Attempt to measure $\Delta\lambda$ for the yellow sodium doublet.

9 Electron Diffraction

Apparatus

cathode ray tube (Leybold 555 626)
high-voltage power supply (new Leybold)
100-kΩ resistor with banana-plug connectors
Vernier calipers

Goals

Observe wave interference patterns (diffraction patterns) of electrons, demonstrating that electrons exhibit wave behavior as well as particle behavior.

Learn what it is that determines the wavelength of an electron.

Introduction

The most momentous discovery of 20th-century physics has been that light and matter are not simply made of waves or particles — the basic building blocks of light and matter are strange entities which display both wave and particle properties at the same time. In our course, we have already learned about the experimental evidence from the photoelectric effect showing that light is made of units called photons, which are both particles and waves. That probably disturbed you less than it might have, since you most likely had no preconceived ideas about whether light was a particle or a wave. In this lab, however, you will see direct evidence that electrons, which you had been completely convinced were particles, also display the wave-like property of interference. Your schooling had probably ingrained the particle interpretation of electrons in you so strongly that you used particle concepts without realizing it. When you wrote symbols for chemical ions such as Cl^- and Ca^{2+}, you understood them to mean a chlorine atom with one excess electron and a calcium atom with two electrons stripped off. By teaching you to count electrons, your teachers were luring you into the assumption that electrons were particles. If this lab's evidence for the wave properties of electrons disturbs you, then you are on your way to a deeper understanding of what an electron really is — both a particle and a wave.

side view

front view

The electron diffraction tube. The distance labeled as 13.5 cm in the figure actually varies from about 12.8 cm to 13.8 cm, even for tubes that otherwise appear identical. This doesn't affect your results, since you're only searching for a proportionality.

Method

What you are working with is basically the same kind of vacuum tube as the picture tube in your television. As in a TV, electrons are accelerated through a voltage and shot in a beam to the front (big end) of the tube, where they hit a phosphorescent coating and produce a glow. You cannot see the electron beam itself. There is a very thin carbon foil (it looks like a tiny piece of soap bubble) near where the neck joins the spherical part of the tube, and the electrons must pass through the foil before crossing over to the phosphorescent screen.

The purpose of the carbon foil is to provide an ultrafine diffraction grating — the "grating" consists of the crystal lattice of the carbon atoms themselves! As you will see in this lab, the wavelengths of the

electrons are very short (a fraction of a nanometer), which makes a conventional ruled diffraction grating useless — the closest spacing that can be achieved on a conventional grating is on the order of one micrometer. The carbon atoms in graphite are arranged in sheets, each of which consists of a hexagonal pattern of atoms like chicken wire. That means they are not lined up in straight rows, so the diffraction pattern is slightly different from the pattern produced by a ruled grating.

Also, the carbon foil consists of many tiny graphite crystals, each with a random orientation of its crystal lattice. The net result is that you will see a bright spot surrounded by two faint circles. The two circles represent cones of electrons that intersect the phosphor. Each cone makes an angle θ with respect to the central axis of the tube, and just as with a ruled grating, the angle is given by

$$\sin \theta = \lambda/d,$$

where λ is the wavelength of the wave. For a ruled grating, d would be the spacing between the lines. In this case, we will have two different cones with two different θ's, θ_1 and θ_2, corresponding to two different $d's$, d_1 and d_2. Their geometrical meaning is shown below.[1]

The carbon atoms in the graphite crystal are arranged hexagonally.

Safety

This lab involves the use of voltages of up to 6000 V. Do not be afraid of the equipment, however; there is a fuse in the high-voltage supply that limits the amount of current that it can produce, so it is not particularly dangerous. Read the safety checklist on

----

[1]See `http://bit.ly/XxoEYr` for more information.

high voltage in Appendix 6. Before beginning the lab, make sure you understand the safety rules, initial them, and show your safety checklist to your instructor. If you don't understand something, ask your instructor for clarification.

In addition to the high-voltage safety precautions, please observe the following rules to avoid damaging the apparatus:

_____ The tubes cost $1000. Please treat them with respect! Don't drop them! Dropping them would also be a safety hazard, since they're vacuum tubes, so they'll implode violently if they break.

_____ Do not turn on anything until your instructor has checked your circuit.

_____ Don't operate the tube continuously at the highest voltage values (5000-6000 V). It produces x-rays when used at these voltages, and the strong beam also decreases the life of the tube. You can use the circuit on the right side of the HV supply's panel, which limits its own voltage to 5000 V. Don't leave the tube's heater on when you're not actually taking data, because it will decrease the life of the tube.

Setup

You setup will consist of two circuits, a heater circuit and the high-voltage circuit.

The heater circuit is to heat the cathode, increasing the velocity with which the electrons move in the metal and making it easier for some of them to escape from the cathode. This will produce the friendly and nostalgia-producing yellow glow which is characteristic of all vacuum-tube equipment. The heater is simply a thin piece of wire, which acts as a resistor when a small voltage is placed across it, producing heat. Connect the heater connections, labeled F1 and F2, to the 6-V AC outlet at the back of the HV supply.

The high-voltage circuit's job is to accelerate the electrons up to the desired speed. An electron that happens to jump out of the cathode will head "downhill" to the anode. (The anode is at a *higher* voltage than the cathode, which would make it seem like it would be uphill from the cathode to the anode. However, electrons have negative charge, so they're like negative-mass water that flows uphill.) The high voltage power supply is actually two different power supplies in one housing, with a left-hand panel for one and a right-hand panel for the other. Connect the anode (A) and cathode (C) to the right-hand

panel of the HV supply, and switch the switch on the HV supply to the right, so it knows you're using the right-hand panel.

The following connections are specified in the documentation, although I don't entirely understand what they're for. First, connect the electrode X to the same plug as the cathode.[2] Also, connect F1 to C with the wire that has the 100-kΩ resistor spliced into it. The circuit diagram on page 37 summarizes all this.

Check your circuit with your instructor before turning it on!

Observations

You are now ready to see for yourself the evidence of the wave nature of electrons, observe the diffraction pattern for various values of the high voltage, and figure out what determines the wavelength of the electrons. You will need to do your measurements in the dark.

Important: As of 2018, some of our tubes are starting to die, and we will not be able to buy replacements until 2020. For this reason, please take the following steps to extend the remaining lifetimes of the working tubes. (1) Don't take too many data points. Change the voltage in steps of 1.0 kV, not smaller steps. (2) Turn the knob on the high voltage power supply all the way down to zero except when you're actually measuring a diffraction fringe. (3) Try to get all your data-taking done without leaving the heater circuit on for more than about 30 minutes.

You will measure the θ's, and thus determine the wavelength, λ, for several different voltages. Each voltage will produce electrons with a different velocity, momentum, and energy.

Hints:

> While measuring the diffraction pattern, don't touch the vacuum tube — the static electric fields of one's body seem to be able to perturb the pattern.

> It is easiest to take measurements at the highest voltages, where the electrons pack a wallop and make nice bright rings on the phosphor. Start with the highest voltages and take data at lower and lower voltages until you can't see the rings well enough to take precise data. To

get unambiguous results, you'll need to take data with the widest possible range of voltages.

Analysis

Once you have your data, the idea is to plot λ as a function of quantities such as KE, p, $1/KE$, or $1/p$. If the graph is a straight line through the origin, then the experiment supports the hypothesis that the wavelength is proportional to that quantity. You can simplify your analysis by leaving out constant factors, and P5 asks you to consider how you can rule out some of these possibilities without having to make all the graphs.

What does λ seem to be proportional to? Your data may cover a small enough range of voltage that more than one graph may look linear. However, only one will be consistent with a line that passes *through the origin*, as it must for a proportionality. This is why it is important to have your graph include the origin.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

The week before you are to do the lab, briefly familiarize yourself visually with the apparatus.

Read the high voltage safety checklist.

P1 The figure shows the vacuum tube as having a particular shape, which is a sphere with the foil and phosphor at opposite ends of a diamater. In reality, the tubes we're using now are not quite that shape. To me, they look like they may have been shaped so that the phosphor surface is a piece of a sphere centered on the foil. Therefore the arc lengths across the phosphor can be connected to diffraction angles very simply via the definition of radian measure. Plan how you will do this.

P2 If the voltage difference across which the electrons are accelerated is V, and the known mass and charge of the electron are m and e, what are the electrons' kinetic energy and momentum, in terms of V, m, and e? (As a numerical check on your results, you should find that $V = 5700$ V gives $KE = 9.1 \times 10^{-16}$ J and $p = 4.1 \times 10^{-23}$ kg·m/s.)

P3 All you're trying to do based on your graphs is

---

[2]If you look inside the tube, you can see that X is an extra electrode sandwiched in between the anode and the cathode. I think it's meant to help produce a focused beam.

POWER SUPPY (high voltage)

The circuit for the new setup.

judge which one could be a graph of a proportionality, i.e., a line passing through the origin. Because of this, you can omit any constant factors from the equations you found in P2. When you do this, what do your expressions turn out to be?

P4 Why is it not logically possible for the wavelength to be proportional to both p and KE? To both $1/p$ and $1/KE$?

P5 I have suggested plotting λ as a function of p, KE, $1/p$ and $1/KE$ to see if λ is directly proportional to any of them. Once you have your raw data, how can you immediately rule out two of these four possibilities and avoid drawing the graphs?

P6 On each graph, you will have two data-points for each voltage, corresponding to two different measurements of the same wavelength. The two wavelengths will be almost the same, but not exactly the same because of random errors in measuring the rings. Should you get the wavelengths by combining the smaller angle with d_1 and the larger angle with d_2, or vice versa?

10   **Relativity**

Apparatus

magnetic balance 1/group
meter stick 1/group
multimeter (BK or PRO-100, not HP) 1/group
laser 1/group
vernier calipers 1/group
photocopy paper, for use as a weight
DC power supply (Mastech, 30 A)
box of special cables
scissors

Goal

Measure the speed of light.

Introduction

Oersted discovered that magnetism is an interaction of moving charges with moving charges, but it wasn't until almost a hundred years later that Einstein showed why such an interaction must exist: magnetism occurs as a direct result of his theory of relativity. Since magnetism is a purely relativistic effect, and relativistic effects depend on the speed of light, any measurement of a magnetic effect can be used to determine the speed of light.

Setup

The idea is to set up opposite currents in two wires, A and B, one under the other, and use the repulsion between the currents to create an upward force on the top wire, A. The top wire is on the arm of a balance, which has a stable equilibrium because of the weight C hanging below it. You initially set up the balance with no current through the wires, adjusting the counterweight D so that the distance between the wires is as small as possible. What we care about is really the center-to-center distance (which we'll call R), so even if the wires are almost touching, there's still a millimeter or two worth of distance between them. By shining a laser at the mirror, E, and observing the spot it makes on the wall, you can very accurately determine this particular position of the balance, and tell later on when you've reproduced it.

If you put a current through the wires, it will raise

wire A. The torque made by the magnetic repulsion is now canceling the torque made by gravity directly on all the hardware, such as the masses C and D. This gravitational torque was zero before, but now you don't know what it is. The trick is to put a tiny weight on top of wire A, and adjust the current so that the balance returns to the position it originally had, as determined by the laser dot on the wall. You now know that the gravitational torque acting on the original apparatus (everything except for the weight) is back to zero, so the only torques acting are the torque of gravity on the staple and the magnetic torque. Since both these torques are applied at the same distance from the axis, the forces creating these torques must be equal as well. You can therefore infer the magnetic force that was acting.

For a weight, you can carefully and accurately cut a small rectangular piece out of a sheet of photocopy paper. In fall 2013, my students found that 500 sheets of SolCopy 20 lb paper were 2307.0 g. About 1/100 of a sheet seemed to be a good weight to use.

It's very important to get the wires A and B perfectly parallel. The result depends strongly on the small distance R between their centers, and if the wires aren't straight and parallel, you won't even have a well defined value of R.

The following technique allows R to be measured accurately. The idea is to compare the position of the laser spot on the wall when the balance is in its normal position, versus the position where the wires are touching. Using a small-angle approximation, you can then find the angle θ_r by which the reflected beam moved. This is twice the angle $\theta_m = \theta_r/2$ by which the mirror moved.[1] Once you

---

[1] To see this, imagine the following example that is unreal-

know the angle by which the moving arm of the apparatus moved, you can accurately find the air gap between the wires, and then add in twice the radius of the wires, which can be measured accurately with vernier calipers. For comparison, try to do as good a job as you can of measuring R directly by positioning the edges of the vernier calipers at the centers of the wires. If the two values of R don't agree, go back and figure out what went wrong; one possibility is that your wire is slightly bent and needs to be straightened.

You need to minimize the resistance of the apparatus, or else you won't be able to get enough current through it to cancel the weight of the staple. Most of the resistance is at the polished metal knife-edges that the moving part of the balance rests on. It may be necessary to clean the surfaces, or even to freshen them a little with a file to remove any layer of oxidation. Use the separate BK meter to measure the current — not the meter built into the power supply.

The power supply has some strange behavior that makes it not work unless you power it up in exactly the right way. It has four knobs, going from left to right: (1) current regulation, (2) over-voltage protection, (3) fine voltage control, (4) coarse voltage control. Before turning the power supply on, turn knobs 1 and 2 all the way up, and knobs 3 and 4 all the way down. Turn the power supply on. Now use knobs 3 and 4 to control how much current flows.

Analysis

The first figure below shows a model that explains the repulsion felt by one of the charges in wire A due to all the charges in wire B. This is represented in the frame of the lab. For convenience of analysis, we give the model some unrealistic features: rather than having positively charged nuclei at rest and negatively charged electrons moving, we pretend that both are moving, in opposite directions. Since wire B has zero net density of charge everywhere, it creates no electric fields. (If you like, you can verify this during lab by putting tiny pieces of paper near the wires and verifying that they do not feel any static-electrical attraction.) Since there is no electric field, the force on the charge in wire A must be purely magnetic.

The second figure shows the same scene from the

---

istic but easy to figure out. Suppose that the incident beam is horizontal, and the mirror is initially vertical, so that the reflected beam is also horizontal. If the mirror is then tilted backward by 45 degrees, the reflected beam will be straight up, $\theta_r = 90$ degrees.

one charge in wire A

unrealistic model of wire B

lab's frame of reference

point of view of the charge in wire A. This charge considers itself to be at rest, and it also sees the light-colored charges in B as being at rest. In this frame the dark-colored charges in B are the only ones moving, and they move with twice the speed they had in the lab frame. In this frame, the particle in A is at rest, so it can't feel any magnetic force. The force is now considered to be purely electric. This electric force exists because the dark charges are relativistically contracted, which makes them more dense than their light-colored neighbors, causing a nonzero net density of charge in wire B.

one charge in wire A

unrealistic model of wire B

frame of the charge in wire A

We've considered the force acting on a single charge in wire A. The actual force we observe in the experiment is the sum of all the forces acting on all such charges (of both signs). As in the slightly different example analyzed in section 11.1.1 of *Simple Nature*, this effect is proportional to the product of the speeds of the charges in the two wires, divided by c^2. Therefore the effect must be proportional to the product of the currents over c^2. In this experiment, the same current flows through wire A and then comes back through B in the opposite direction, so we conclude that the force must be proportional to I^2/c^2.

In the second frame, the force is purely electrical, and as can easily be shown by Gauss's law, the electric field of a charged wire falls off in proportion to $1/R$, where R is the distance from the wire. Electrical forces are also proportional to the Coulomb constant k.

The longer the wires, the more charges interact, so we must also have a proportionality to the length ℓ.

Putting all these factors together, we find that the force is proportional to $kI^2\ell/c^2R$. We can easily verify that the units of this expression are newtons, so the only possible missing factor is something unitless. This unitless factor turns out to be 2 by Gauss's law. The result for the repulsive force between the two wires is

$$F = \frac{k}{c^2} \cdot \frac{2I^2\ell}{R} \qquad .$$

By solving this equation you can find c. Your final result is the speed of light, with error bars. Compare with the previously measured value of c and give a probabilistic interpretation, as in the examples in appendix 2.

In your writeup, give both the values of R (laser and eyeball). The laser technique is inherently better, so that's the value you should use in extracting c, but I want to see both values of R because some groups in the past have had a bigger discrepancy than I would have expected. If you have a large discrepancy, get my attention during lab and we can see whether it might be due to a bent wire, or some other cause.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

Do the laser safety checklist, Appendix 7, tear it out, and turn it in at the beginning of lab. If you don't understand something, don't initial that point, and ask your instructor for clarification before you start the lab.

P1 Show that the equation for the force between the wires has units of newtons.

P2 Do the algebra to solve for c in terms of the measured quantities.

11   Polarization

Apparatus

laser 1/group
calcite crystal (flattest available) 1/group
polarizing films 2/group
Na gas discharge tube 1/group
photovoltaic cell and collimator 1/group

Goals

Make qualitative observations about the polarization of light.

Test quantitatively the hypothesis that polarization relates to the direction of the field vectors in an electromagnetic wave.

Introduction

It's common knowledge that there's more to light than meets the eye: everyone has heard of infrared and ultraviolet light, which are visible to some other animals but not to us. Another invisible feature of the wave nature of light is far less well known. Electromagnetic waves are transverse, i.e., the electric and magnetic field vectors vibrate in directions perpendicular to the direction of motion of the wave. Two electromagnetic waves with the same wavelength can therefore be physically distinguishable, if their electric and magnetic fields are twisted around in different directions. Waves that differ in this way are said to have different polarizations.

An electromagnetic wave has electric and magnetic field vectors that vibrate in the directions perpendicular to its direction of motion. The wave's direction of polarization is defined as the line along which the electric field lies.

Maybe we polarization-blind humans are missing out on something. Some fish, insects, and crustaceans can detect polarization. Most sources of visible light (such as the sun or a light bulb) are unpolarized. An unpolarized beam of light contains a random mixture of waves with many different directions of polarization, all of them changing from moment to moment, and from point to point within the beam.

Qualitative Observations

Before doing anything else, turn on your gas discharge tube, so it will be warmed up when you are ready to do part E.

A   Double refraction in calcite

Place a calcite crystal on this page. You will see two images of the print through the crystal.

To understand why this happens, try shining the laser beam on a piece of paper and then inserting the calcite crystal in the beam. If you rotate the crystal around in different directions, you should be able to get two distinct spots to show up on the paper. (This may take a little trial and error, partly because the effect depends on the correct orientation of the crystal, but also because the crystals are not perfect, and it can be hard to find a nice smooth spot through which to shine the beam.)

In the refraction lab, you've already seen how a beam of light can be bent as it passes through the interface between two media. The present situation is similar because the laser beam passes in through one face of the crystal and then emerges from a parallel face at the back. You have already seen that in this type of situation, when the beam emerges again, its direction is bent back parallel to its original direction, but the beam is offset a little bit. What is different here is that the same laser beam splits up into two parts, which bumped off course by different amounts.

What's happening is that calcite, unlike most substances, has a different index of refraction depending on the polarization of the light. Light travels at a different speed through calcite depending on how the electric and magnetic fields are oriented compared to the crystal. The atoms inside the crystal are packed in a three-dimensional pattern sort of like a stack of oranges or cannonballs. This packing arrangement has a special axis of symmetry, and light polarized along that axis moves at one speed, while light polarized perpendicular to that axis moves at a different

speed.

It makes sense that if the original laser beam was a random mixture of all possible directions of polarization, then each part would be refracted by a different amount. What is a little more surprising is that two separated beams emerge, with nothing in between. The incoming light was composed of light with every possible direction of polarization. You would therefore expect that the part of the incoming light polarized at, say, 45 °compared to the crystal's axis would be refracted by an intermediate amount, but that doesn't happen. This surprising observation, and all other polarization phenomena, can be understood based on the vector nature of electric and magnetic fields, and the purpose of this lab is to lead you through a series of observations to help you understand what's really going on.

## B   A polarized beam entering the calcite

A single laser beam entering a calcite crystal breaks up into two parts, which are refracted by different amounts.

The calcite splits the wave into two parts, polarized in perpendicular directions compared to each other.

We need not be restricted to speculation about what was happening to the part of the light that entered the calcite crystal polarized at a 45 °angle. You can use a polarizing film, often referred to informally as a "Polaroid," to change unpolarized light into a beam of only one specific polarization. In this part of the lab, you will use a polarizing film to produce a beam of light polarized at a 45 °angle to the crystal's internal axis.

If you simply look through the film, it doesn't look like anything special — everything just looks dimmer, like looking through sunglasses. The light reach-

The original beam contains every possible polarization.

Only two beams with two polarizations emerge from the calcite crystal.

ing your eye is polarized, but your eye can't tell that. If you looked at the film under a microscope, you'd see a pattern of stripes, which select only one direction of polarization of the light that passes through.

Now try interposing the film between the laser and the crystal. The beam reaching the crystal is now polarized along some specific direction. If you rotate the film, you change beam's direction of polarization. If you try various orientations, you will be able to find one that makes one of the spots disappear, and another orientation of the film, at a 90 °angle compared to the first, that makes the other spot go away. When you hold the film in one of these directions, you are sending a beam into the crystal that is either purely polarized along the crystal's axis or purely polarized at 90 °to the axis.

By now you have already seen what happens if the film is at an intermediate angle such as 45 °. Two spots appear on the paper in the same places produced by an unpolarized source of light, not just a single spot at the midpoint. This shows that the crystal is not just throwing away the parts of the light that are out of alignment with its axis. What is happening instead is that the crystal will accept a beam of light with any polarization whatsoever, and split it into two beams polarized at 0 and 90 °compared to the crystal's axis.

This behavior actually makes sense in terms of the wave theory of light. Light waves are supposed to obey the principle of superposition, which says that waves that pass through each other add on to each other. A light wave is made of electric and magnetic fields, which are vectors, so it is vector addition we're talking about in this case. A vector at a 45 °angle can be produced by adding two perpendicular vectors of equal length. The crystal therefore cannot respond any differently to 45-degree polarized light than it would to a 50-50 mixture of light with 0-

degree and 90-degree polarization.

The vector sum of

this   plus   this   equals   this.

The principle of superposition implies that if the 0 °and 90 °polarizations produce two different spots, then the two waves superimposed must produce those two spots, not a single spot at an intermediate location.

C Two polarizing films

So far I've just described the polarizing film as a device for producing polarized light. But one can apply to the polarizing film the same logic of superposition and vector addition that worked with the calcite crystal. It would not make sense for the film simply to throw away any waves that were not perfectly aligned with it, because a field oriented on a slant can be analyzed into two vector components, at 0 and 90 °with respect to the film. Even if one component is entirely absorbed, the other component should still be transmitted.

Based on these considerations, now think about what will happen if you look through two polarizing films at an angle to each other, as shown in the figure above. *Do not look into the laser beam!* Just look around the room. What will happen as you change the angle θ?

D Three polarizing films

Now suppose you start with two films at a 90 °angle to each other, and then sandwich a third film between them at a 45 °angle, as shown in the two figures above. Make a prediction about what will happen, and discuss your prediction with your instructor before you make the actual observation.

Quantitative Observations

E Intensity of light passing through two polarizing films

In this part of the lab, you will make numerical measurements of the transmission of initially unpolarized light transmitted through two polarizing films at an angle θ to each other. To measure the intensity of the light that gets through, you will use a photovoltaic cell, which is a device that converts light energy into an electric current. The ones we're using are of a type known as a silicon photodiode.

You will use an ammeter to measure the current flowing from the photocell when light is shining on it. This is known as the "short-circuit current," because the ammeter ideally has zero resistance, so it acts like a short. Normally when you create a short through an ammeter, it blows the fuse in the meter, but here there is about 40 kΩ of internal resistance in the silicon, which is a semiconductor. A photovoltaic cell is a complicated nonlinear device, but I've found empirically that under the conditions we're using in this experiment, the current is proportional to the power of the light striking the cell: twice as much light results in twice the current.[1]

This measurement requires a source of light that is unpolarized, constant in intensity, has a wavelength that the polaroids work with, and comes from a specific direction so it can't get to the photocell without going through the polaroids. The ambient light in the room is nearly unpolarized, but varies randomly as people walk in front of the light fixtures, etc. An incandescent lightbulb doesn't work, because it puts out a huge amount of infrared light, which the silicon cell measures but the polaroids can't work with. A laser beam is constant in intensity, but as I was creating this lab I found to my surprise that the

---

[1]It's also possible to use the same cell with a voltmeter across it, in which case we'd be measuring the "open-circuit voltage;" but the open-circuit voltage varies in a much more nonlinear way with the intensity of the light. When rooftop photovoltaic cells are used to generate power, the resistance of the load is neither zero nor infinity, and is chosen to maximize the efficiency.

light from the laser I tried was partially polarized, with a polarization that varied over time. A more suitable source of light is the sodium gas discharge tube, which makes a nearly monochromatic, unpolarized yellow light. Make sure you have allowed it to warm up for at least 15-20 minutes before using it; before it warms up, it makes a reddish light, and the polaroids do not work very well on that color.

Make measurements of the relative intensity of light transmitted through the two polarizing films, using a variety of angles θ. Don't assume that the notches on the plastic housing of the polarizing films are a good indication of the orientation of the films themselves.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

P1 Given the angle θ between the polarizing films, predict the ratio $|\mathbf{E}'|/|\mathbf{E}|$ of the transmitted electric field to the incident electric field.

P2 Based on your answer to P1, predict the ratio P'/P of the transmitted power to the incident power.

P3 Sketch a graph of your answer to P2. Superimposed on the same graph, show a qualitative prediction of how it would change if the polaroids were not 100% perfect at filtering out one component of the field.

Analysis

Discuss your qualitative results in terms of superposition and vector addition.

Graph your results from part E, and superimpose a theoretical curve for comparison. Discuss how your results compare with theory. Since your measurements of light intensity are relative, just scale the points so that their maximum matches that of the experimental data. (You might think of comparing the intensity transmitted through the two polaroids with the intensity that you get with no polaroids in the way at all. This doesn't really work, however, because in addition to acting as polarizers, the polaroids simply absorb a certain percentage of the light, just as any transparent material would.)

12   Radioactivity

Note to the lab technician: The isotope generator kits came with 250 mL bottles of eluting solution (0.9% NaCl in 0.04M HCl, made with deionized water). If we ever run out of the solution, we can make more from materials in the chem stockroom. The GM counters have 9 V batteries, which should be checked before lab.

Apparatus

isotope generator kit
Geiger-Müller (GM) counter
computer with Logger Pro software and LabPro interface
"grabber" clamp and stand
wood blocks, 25 mm thick
pieces of steel, 17 mm thick

Goal

Determine the properties of an unknown radioactive source.

Introduction

You're a science major, but even if you weren't, it would be important for you as a citizen and a voter to understand the properties of radiation. As an example of an important social issue, many environmentalists who had previously opposed nuclear power now believe that its benefits, due to reduction of global warming, outweigh its problems, such as disposal of waste. To understand such issues, you need to learn to reason about radioactivity quantitatively.

A radioactive substance contains atoms whose nuclei spontaneously decay into nuclei of a different type. Nobody has ever succeeded in finding a physical law that would predict when a particular nucleus will "choose" to decay. The process is random, but we can make quantitative statements about how quickly the process tends to happen. A radioactive substance has a certain *half-life*, defined as the time required before (on the average) 50% of its nuclei will have decayed.

Safety

The radioactive source used in this lab is very weak. It is so weak that it is exempt from government regulation, it can be sent in the mail, and when the college buys new equipment, it is legal to throw out the old sources in the trash. The following table compares some radiation doses, including an estimate of the typical dose you might receive in this lab. These are in units of microSieverts (μSv).

CT scan	$\sim 10{,}000~\mu$Sv
natural background per year	2,000-7,000 μSv
health guidelines for exposure to a fetus	1,000 μSv
flying from New York to Tokyo	150 μSv
this lab	$\sim 80~\mu$Sv
chest x-ray	50 μSv

A variety of experiments seem to show cases in which low levels of radiation activate cellular damage control mechanisms, increasing the health of the organism. For example, there is evidence that exposure to radiation up to a certain level makes mice grow faster; makes guinea pigs' immune systems function better against diphtheria; increases fertility in female humans, trout, and mice; improves fetal mice's resistance to disease; reduces genetic abnormalities in humans; increases the life-spans of flour beetles and mice; and reduces mortality from cancer in mice and humans. This type of effect is called radiation hormesis. Nobody knows for sure, but it's possible that you will receive a very tiny improvement in your health from the radiation exposure you experience in lab today.

Although low doses of radiation may be beneficial, governments, employers, and schools generally practice a philosophy called ALARA, which means to make radiation doses As Low As Reasonably Achievable. You should adhere to this approach in this lab. In general, internal exposure to radiation produces more of an effect than external exposure, so you should not eat or drink during this lab, and you should avoid getting any of the radioactive substances in an open cut. You should also reduce your exposure by not spending an unnecessarily large amount of time with your body very close to the source, e.g., you should not hold it in the palm of your hand for the entire lab period.

The source and the GM counter

You are supplied with a radioactive source packaged inside a small plastic disk about the size of the spindle that fits inside a roll of scotch tape.

Our radiation detector for this lab is called a Geiger-Müller (GM) counter. It is is a cylinder full of gas, with the outside of the cylinder at a certain voltage and a wire running down its axis at another voltage. The voltage difference creates a strong electric field. When ionizing radiation enters the cylinder, it can ionize the gas, separating negatively charged ions (electrons) and positively charged ones (atoms lacking some electrons). The electric field accelerates the ions, making them hit other air molecules, and causing a cascade of ions strong enough to be measured as an electric current.

If you look at the top side of the GM counter (behind the top of the front panel), you'll see a small window. Non-penetrating radiation can only get in through this thin layer of mica. (Gammas can go right in through the plastic housing.)

Put the bottom switch on Audio. The top switch doesn't have any effect on the data collection we'll be doing with the computer.

Poisson Statistics

Although we will not be formally estimating error bars in this lab, the following information will be helpful in understanding what's going on when you're taking data. When we have a large number of things that may happen with some small probability, the total number of them that do happen is called a Poisson random variable (accent on the second syllable). For example, the number of houses burglarized in Fullerton this year is a Poisson random variable. When you count the number of nuclear decays in a certain time interval, the result is Poisson. The helpful thing to know is that when a Poisson variable has an average value N, its statistical uncertainty is $\sqrt{N}$. So for example if your GM counter counts 100 clicks in one minute, this is 100 ± 10. Knowing this will help you to have some idea whether, for example, an apparent change in the count rate is actually too small to be statistically meaningful, or whether you need to count for longer in order to get reliable results.

Observations

A Background

Use the GM counter to observe the background radiation in the room. This radiation is probably a combination of gamma rays from naturally occurring minerals in the ground plus betas and gammas from building materials such as concrete. If you like, you can walk around the room and see if you can detect any variations in the intensity of the background.

Estimate the rate at which the GM detector counts when it is exposed only to background. Starting at this point, it is more convenient to interface the GM counter to the computer. Plug the cable into into DIG/SONIC 1 on the LabPro interface. Start LoggerPro 3 on the computer, and open the file Probes and Sensors : Radiation monitor : Counts versus time. The interface is not able to automatically identify this particular sensor, so the software will ask you to confirm that you really do have this type of sensor hooked up; confirm this by clicking on Connect.

Now when you hit the Collect button, the sensor will start graphing the number of counts it receives during successive 5-second intervals. The y axis of the graph is counts per 5 seconds, and the x axis is time.

Once you've made some preliminary observations, try to get a good measurement of the background, counting for several minutes so as to reduce the statistical errors. It is important to get a good measurement of the background rate, because in later parts of the lab you'll need to subtract it from all the other count rates you measure. For longer runs like this, it is convenient to let the software collect data for about a minute at a time, rather than 5 seconds. To get it to do this, do Experiment : Data collection : 60 seconds/sample.

B Type of radiation

Your first task in the lab is to figure out whether the source emits alpha, beta, or gamma radiation — or perhaps some mixture of these. It is up to you to decide how to do this, but essentially you want to use the fact that they are absorbed differently in matter; referring to your textbook, you'll see that the historical labels α, β, and γ were assigned purely on the basis of the differences in absorption, before anyone even knew what they really were. Note that the fact that you are able to detect the radiation at all implies that at least some of the radiation is able to penetrate the thin plastic walls of the source.

Also keep in mind that the descriptions of absorption in a textbook are generalizations that do not take into account the energy of the particles. For example, low-energy betas could be absorbed by a kleenex, whereas high-energy betas could penetrate cardboard.

There are six things you could try to prove through your measurements:

1. The source inside the plastic container emits alphas.

2. The source does not emit alphas.

3. The source emits betas.

4. The source does not emit betas.

5. The source emits gammas.

6. The source does not emit gammas.

Try to figure out which of these six statements you can either definitively prove or definitively disprove. Because you are not allowed to extract the source from the packaging, there will be some cases in which you cannot draw any definitive conclusion one way or the other.

C Distance dependence

Measure the count rate at several different distances from the source. The goal is to find the mathematical form of this function (see Analysis, below). Distances of less than about 10 cm do not work well, because the size of the GM tube is comparable to 10 cm.

D Absorption

Measure the reduction in count rate when a 25 mm thick wood block is interposed between the source and the detector, and likewise for 17 mm of steel. Keep the distance from the source to the detector constant throughout. Let the counter run for at least five minutes each time. Based on these observations, predict the count rate you would get with two 17-mm thicknesses of steel instead of one. Test your prediction.

E Decay curve

The source consists of a particular isotope of cesium; we'll refer to it as ^{NNN}Cs, since the main goal of this lab is to determine what the unknown isotope actually is. It decays to an isotope of barium, and rather than decaying to the ground state of the barium nucleus, it nearly always decays to an excited state,

which then emits the radiation you characterized in part B. Although the half-life of the parent cesium isotope is many years, the half-life of the excited isotope in the barium daughter is short enough that it can be observed during the lab period. However, if the cesium and barium are not separated, then no time variation will be observed, because the supply of barium nuclei is being continuously replenished by decay of the parent cesium.

To get around this, the source is packaged so that when a weak acid solution is forced through it, a small amount of barium is washed out. Note that the yellow tape around the circumference of the source has an arrow on it. This arrow points in the direction that you're supposed to make the solution flow. The isotope generator kit has coin-sized steel trays on which to collect a few drops of the radioactive solution.

Use the syringe to draw 1 mL of the acid solution from the 250-mL bottle (labeled "eluting solution"). Take one of the tiny coin-sized steel trays out of the isotope generator kit and lay it on the lab bench. Remove the little plugs from the top and bottom of the radioactive source. Stick the syringe in the in-flow hole, and use the plunger to force seven drops of liquid out onto the tray. Note that the amount of liquid that flows from the syringe into the source is quite a bit more than the amount that comes out into the tray. If you have solution left in the syringe at the end of lab, squeeze it back out into the 250-mL bottle.

Use the computer to collect data on the rate of decay as a function of time. About 5 or 10 seconds per sample works well.

When you're done, make sure to shut off the GM counter so that its battery doesn't get drained.

Waste disposal

To get an idea of what a non-issue radioactive waste disposal is in this lab, recall that it would be legal to throw the *entire* source in the trash — although we won't actually do that. The amount of radioactive material that you wash out in part E is a tiny

fraction of this. Furthermore, essentially all the radioactivity is gone by the end of lab. It is therefore not a problem to dump your seven drops of material down the drain at the end of class.

There is also no chemical disposal issue with this tiny amount of solution. It's a few drops of very dilute acid, equivalent to a little spritz of lemon juice.

Analysis

In part C, you should first subtract the background rate from each datum. Then make a log-log plot as described in appendix 5, and see if you can successfully describe the data using a power law. Note that, just like a human, the GM counter cannot count faster than a certain rate. This is because every time it gets a count it completely discharges its voltage, and then it has to recharge itself again. For this reason, it is possible that your data from very small distances will not agree with the behavior of the data at larger distances. The documentation for these GM counters says that they can count at up to about 3500 counts per second; this is only a very rough guide, but it gives you some idea what count rates should be expected to start departing from ideal behavior.

Estimate the half-lives of any isotopes present in the data. If you find that only one half-life is present, you can simply determine the amount of time required for the count rate to fall off by a factor of two. If the natural background count rate measured in part A is significant, you will need to subtract it from the raw data. If more than one half-life is present, try plotting the logarithm of the count rate as a function of time, and seeing if there are linear sections on the graph. Note that this is all referring to the half-lives of any decay chain that occurs *after* the cesium decays to barium. The half-life of the cesium parent nucleus is much longer, and is not measured in this experiment.

Consult the Wikipedia article "Isotopes of cesium." It has an extremely lengthy table of all the known isotopes from very light ones (with far too few neutrons to be stable) to very heavy ones (with far too many). Since the source was shipped to us through the mail, and sits on the shelf in the physics stockroom for semester after semester, you can tell that the half-life of the cesium isotope must be fairly long — at least on the order of years, not months. From this information, you should be able to narrow down the range of possibilities. (Half-lives in units of years are listed with "a," for "annum," as the unit of time.

The notations m1 and m2 mean energy states that are above the lowest-energy state.) The radioactive isotopes from this remaining list of possibilities all have their own Wikipedia articles, and these articles give the properties of the daughter nuclei (isotopes of barium and xenon), including the half-lives of any gamma-emitting states. Look for one that has a half-life that seems to match the one you measured in lab. Having tentatively identified the unknown isotope as this isotope of cesium, check against the results of part B, where you determined the type of radiation emitted.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

P1 Suppose that in part C you obtain the following data:

r (cm)	count rate (counts/2 min)
10	707
20	207
30	95

Suppose that the background rate you measure in part A is 30 counts per 2 min. Use the technique described in appendix 5 to see if the data can be described by a power law, and if so, determine the exponent.

P2 If a source emits gamma or beta radiation, then the radiation spreads out in all directions, like an expanding sphere. Based on the scaling of a sphere's surface area with increasing radius, how would you theoretically expect the intensity of the radiation to fall off with distance? Would it be a power law? If so, what power? Why would you *not* expect the same behavior for a source emitting alpha particles into air?

13   The Photoelectric Effect

Note to the lab technician: The Pasco SE-5509 Hg gas discharge tubes have too many ways of being destroyed by students,[1] so let's use the OS-9286 tubes instead (the big black ones with the fins). That means we don't need the separate power supply for the discharge tube (Pasco BEM-5007) or the track that comes with the photodiode. We need wood blocks to raise the photodiode to the same height as the discharge tube.

Apparatus

hand-held diffraction gratings
Hg gas discharge tube (Pasco OS-9286)
photodiode (Pasco SE-5509)
power supply (Pasco BEM-5001)
high-sensitivity ammeter (Pasco BEM-5004)
wood blocks 3/group

Goals

Use the photoelectric effect to test predictions of the wave and wave-particle models of light.

Introduction

The photoelectric effect, a phenomenon in which light shakes an electron loose from an object, provided the first evidence for wave-particle duality: the idea that the basic building blocks of light and matter show a strange mixture of particle and wave behaviors. At the turn of the twentieth century, physicists assumed that particle and wave phenomena were completely distinct. Young had shown that light could undergo interference effects such as diffraction, so it had to be a wave. Since light was a wave composed of oscillating electric and magnetic fields, it made sense that when light encountered matter, it would tend to shake the electrons. It was only to be expected that something like the photoelectric effect could happen, with the light shaking the electrons vigorously enough to knock them out of the atom.

---

[1] The two main modes of destruction seem to be: (1) While the discharge tube is connected to the power supply, students monkey with the red 110V/220V switch on the power supply; and (2) they don't use the power supply and connect two AC connectors together in order to connect the discharge tube directly to the wall socket.

But once the effect was observed, physicists began running into trouble interpreting how it behaved. There were various variables they could adjust, such as:

- the light's frequency (color), and

- the light's intensity (brightness).

Given these input conditions, there were outputs they could look at, including

- any time delay before electrons began to pop out,

- the rate at which the electrons then flowed (measured as a current on an ammeter), and

- the amount of kinetic energy they had.

At the time this was considered an obscure technical topic, but experimentalists began generating data, which theorists then had zero success in interpreting. Albert Einstein, better known today for the theory of relativity, was the first to come up with the radical, and correct, explanation, which involved a fundamental rewriting of the laws of physics.

Setup

The Hg gas discharge tube emits light with several different wavelengths. Turn on the discharge tube immediately, because it takes a long time to warm up.

The photodiode is a vacuum tube housed inside another box, with a small hole to allow light to come in and hit one of the electrodes (the cathode) inside the vacuum tube. On the front of the box, covering the hole, are two rotating wheels. The wheel that you can see has five colored filters. Each of these filters lets through only one of the five wavelengths of light emitted by the Hg tube, so that you can control the frequency.

The hole through which the light enters is actually a hole in a second wheel, located behind the filter wheel. By clicking that wheel into different position we can select holes of different sizes, which lets in different amounts of light. Although this is convenient, it doesn't actually control the intensity of the light

(watts per square meter) but only the total amount (watts). Therefore we'll just leave this set for the 2 mm hole.

There is an easier way to control the actual intensity of the beam, which is that you can simply change the distance between the discharge tube and the photodiode. The light from the discharge tube spreads out in a cone, so that as you vary the distance to the photodiode, the intensity falls off as the inverse square of the distance.

There is a power supply used for applying an external voltage to the photodiode. Although you don't actually want to apply that external voltage during this part of the lab, it seems that the ammeter won't work until you hook up the power supply, so you need to do that now. On the bottom right side of the power supply are two banana plugs. Connect these to the two plugs near the bottom of the photodiode. To get the polarity right, connect the positive (red) output of the power supply to the anode (marked A).

The photodiode has an output that can be connected to a extremely sensitive ammeter to measure the rate at which electrons are ejected from the cathode and absorbed at the anode. Before connecting the ammeter (labeled "DC Current Amplifier" on the front panel), set it to its most sensitive scale (10^{-13} A) and depress the button labeled "Calibration." Since the meter is not connected to anything, the current is truly zero. Use the knob to force the meter to read zero, and then let the calibration button out. Now use a BNC cable to connect the photodiode to the ammeter.

Once the discharge tube has warmed up, arrange things so that you can see a spot of its light, e.g., by letting it fall on a white piece of paper. Hold the diffraction grating up to your eye and look at the light. In the first-order ($m = \pm 1$) fringes, you should be able to see that the light contains a mixture of four discrete wavelengths of visible light. There is also an ultraviolet wavelength, which you may be able to see as well if the paper fluoresces.

Now suppose we were considering the following possible models of light: (1) pure particle model, (2) pure wave model, (3) a hybrid model in which light has both particle and wave properties. You have just observed diffraction of the light from your source. Which of the models are consistent with this observation, and which are immediately ruled out?

The wavelengths are as follows:

color	wavelength (nm)
ultraviolet	365
violet	405
blue	436
green	546
yellow	578

light

The circuit. The light creates an electric current, which trickles back through the large but finite resistance of the voltmeter.

Circuit

The figure above shows a circuit diagram of the setup. Light comes in and knocks electrons out of the curved cathode. If the voltage is turned off, there is no electric field, so the electrons travel in straight lines; some will hit the anode, creating a current. If the voltage is turned on, the electric field repels the electrons from the wire electrode, and the current is reduced or perhaps even completely eliminated.

Observations

Although the photodiode box has filters on the front, no filter is perfect, and therefore these will all let in some stray light of wavelengths other than the intended one. Therefore the room should be very dark when you do your measurements.

A Time delay

As explored in the prelab, there may be some delay between the time when the light is allowed to hit the cathode and the time when electrons begin to be ejected. If so, then we lack even a rough a priori estimate of this time. Investigate this. If the time seems to be extremely short, do what you planned in the prelab to try to make it long enough to detect. If the time is much too long, do the opposite so that you can actually observe the photoelectric effect. If you're able to get the time delay into a range where it's measurable using eyeballs and a clock, do so. If not, then try to set an upper or lower limit on it.

B Energy of the electrons

Until you do the lab, it's not obvious how much energy the electrons would have as they pop out of the cathode. It could be some fixed number, or it could depend on the conditions you choose, and it could also have multiple values for the electrons produced under a given set of conditions. In fact, we may expect a range of values for two reasons.

(1) As explored in a prelab question, the electrons will have random kinetic energies to start with, due to their thermal motion

(2) The light penetrates to some depth in the cathode, and an electron that starts at some depth will lose some amount of energy as it then comes out to the surface. The electron's direction of motion is random. If it happens to be toward the surface, then the thickness of material that it traverses will depend both on its initial depth and on its random direction of motion.

This range of energies may have an upper limit for a given set of experimental conditions. If so, then by applying a high enough voltage you should be able to eliminate the current completely. The minimum voltage required to do this would be called the stopping voltage.

Now you want to apply a voltage to the photodiode using the DC power supply, which you previously connected but didn't use. There is a button between the two LED readouts. Let this button out in order to select a range of voltages from 0 to -4.5 V.

Find out whether there is a stopping voltage, and if so, measure it for the conditions you've chosen. If there is never a sharp cutoff, you should still be able to determine some quantitative measure of the voltage that corresponds to a *typical* energy for the electrons, e.g., the voltage at which some fixed fraction of the current is eliminated. From now on in the lab manual, I'll just refer to this as "the voltage" for a given set of experimental conditions.

Determine the voltage, and compare with the estimate in the prelab of what voltage would correspond to the thermal motion of the electrons. Based on this comparison, is the amount of energy involved in the photoelectric effect much less than, much more than, or on the same order of magnitude as the thermal energy?

C Dependence of voltage on intensity

Vary the intensity of the light and determine whether and how the voltage depends on intensity.

We have observed strange results sometimes, which seem to occur because when the discharge tube is very close to the photodiode, light gets in and hits the *anode* (not just the cathode) and causes the photoelectric effect in the wrong direction. This seems to occur with the shorter wavelengths of light. You should be able to tell if this is happening because when you turn up the voltage high enough, you get a current in the opposite direction. Check for this behavior and avoid taking data under conditions that produce it.

D Dependence of voltage on frequency

Do the same thing for the frequency.

Analysis

The point of the analysis is to try to compare the observed results with our expectations based on two models of light: the pure wave model, and a model in which light is both a particle and a wave — we've seen above that a pure particle model is untenable. For conceptual simplicity, however, we may find it helpful to visualize the wave-particle model as if the light was purely particle-like, so that a beam of light would be like a stream of machine-gun bullets. This is essentially what Einstein does in his 1905 paper. He admits that this is not literally possible, but doesn't attempt a more detailed reconciliation of the particle and wave pictures, which he doesn't know how to achieve. For this reason, the title of the paper refers to the particle picture as a "heuristic," which means a kind of non-rigorous way of getting an answer without using correct fundamental principles.

Time delay: You investigated the possible time delay in the wave model in some detail in the prelab and/or homework. In a particle model, what would be your expectations about a time delay? Compare with experiment.

Energy: Based on the particle model, we would expect one "bullet" of light to give its fixed amount of energy to one electron, so that under a given set of experimental conditions, there would be a maximum possible kinetic energy for the electrons, achieved when the electron originated very close to the surface of the cathode. In the wave model it is more difficult to make a definite prediction. Did you observe that there was a definite stopping voltage, or that there was no definite cut-off in the current? Does this allow you to test either model?

Dependence of energy on intensity: In the particle model, a more intense beam of light would be one

containing a larger number of particles (per unit of cross-sectional area). In the wave model, a more intense beam would be a higher-amplitude wave. Does either model lead you to predict anything specific about the dependence of voltage on intensity? Test against experiment.

Dependence of energy on frequency: Does the (typical or maximum) energy of the electrons eV depend on the light's frequency f? If so, then in the particle model this probably means that we're observing a change in the amount of energy per particle of light. We can't just equate the electron energy eV to the energy of the particles of light, because the electrons lose a fixed amount of energy (called the work function) as they emerge from the surface of the metal. We can get around the issue of this constant offset by finding the slope of eV versus f; a constant offset on the y axis doesn't affect the slope of a graph. Let's call this constant h. Estimate its numerical value. What units does it have? There is no such universal constant in any of the classical laws of physics, so if it pops up here, it indicates that some entirely new physical theory is being probed.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

P1   Under the hypothesis that light is purely a wave, the energy of a beam of light would be smoothly and continuously accumulated by whatever it hit. Therefore it should take some amount of time before an electron could accumulate enough energy to pop out of the metal. You may have estimated this time scale in a homework problem (*Modern Physics* problem 15-6), but that calculation depends on many crude assumptions and rough estimates, so it's hard to know whether to trust it even as an order of magnitude estimate. Therefore if we want to try to observe this time delay in this experiment, it's impossible to know in advance what to look for, and it may be either too short (in which case we can't measure it) or too long (in which case it will look like the apparatus simply isn't working). Suppose that the time delay is too short to detect in the conditions you initially pick. How could you change the conditions in order to make it longer, and possibly detectable?

P2   Suppose that the cathode is at temperature T.

Then the electrons inside the cathode already have some amount of kinetic energy, due to their thermal motion. Estimate the typical amount of kinetic energy. In the photoelectric effect, an electron will absorb some additional amount of energy from the light, which is enough to pop it out of the cathode. Until doing the experiment, we do not know how much this additional amount is, but during the lab you will be able to probe this question by using a voltage V to try to stop the electrons from making it across the gap. If this additional energy was on the *same* order of magnitude as the thermal energy (which it may not be), estimate the voltage required.

P3   In this experiment, the light comes out of the discharge tube in a spreading cone. Geometrically, how should the intensity I of the light depend on the distance r? State a proportionality.

P4   Look up the wavelength of visible light and the typical distance between atoms in a solid. How do they compare? In the wave model, should we expect a particular wave-train to hit one atom, or many?

14a   Setup of the Spectrometer

Apparatus

Goals

The lab has three parts. This one, part a, is about setting up the optics of the spectrometer. This is to be done once by the instructor or lab technician. It never needs to be redone unless something gets messed up.

Introduction

Method

The apparatus is shown in the first figure below. For a given wavelength, the grating produces diffracted light at many different angles: a central zeroth-order line at $\theta = 0$, first-order lines on both the left and right, and so on through higher-order lines at larger angles. The line of order m occurs at an angle satisfying the equation $m\lambda = d\sin\theta$.

To measure a wavelength, students will move the telescope until the diffracted first-order image of the slit is lined up with the telescope's cross-hairs and then read off the angle. Note that the angular scale on the table of the spectroscope actually gives the angle labeled α in the figure, not θ.

Sources of systematic errors

There are three sources of systematic error:

> *angular scale out of alignment:* If the angular scale is out of alignment, then all the angles will be off by a constant amount.

> *factory's calibration of d:* The factory that made the grating labeled it with a certain spacing (in lines per millimeter) which can be con-

verted to d (center-to-center distance between lines). But their manufacturing process is not all that accurate, so the actual spacing of the lines is a little different from what the label says.

> *orientation of the grating:* Errors will be caused if the grating is not perpendicular to the beam from the collimator, or if the lines on the grating are not vertical (perpendicular to the plane of the circle).

Eliminating systematic errors

A trick to eliminate the error due to misalignment of the angular scale is to observe the same line on both the right and the left, and take θ to be half the difference between the two angles, i.e., $\theta = (\alpha_R - \alpha_L)/2$. Because you are subtracting two angles, any source of error that adds a constant offset onto the angles is eliminated. A few of the spectrometers have their angular scales out of alignment with the collimators by as much as a full degree, but that's of absolutely no consequence if this technique is used.

Regarding the calibration of d, the first person who ever did this type of experiment simply had to make a diffraction grating whose d was very precisely constructed. But once someone has accurately measured at least one wavelength of one emission line of one element, one can simply determine the spacing, d, of any grating using a line whose wavelength is known.

You might think that these two tricks would be enough to get rid of any error due to misorientation of the grating, but they're not. They will get rid of any error of the form $\theta \to \theta + c$ or $\sin\theta \to c\sin\theta$, but misorientation of the grating produces errors of the form $\sin\theta \to \sin\theta + c$. Part A below describes some additional adjustments that help to get rid of additional sources of error.

Theory of Operation

The second figure below shows the optics from the side, with the telescope simply looking down the throat of the collimator at $\theta = 0$. You are actually using the optics to let you see an image of the slit, not the tube itself. The point of using a telescope is that it provides angular magnification, so that a small change in angle can be seen visually.

A lens is used inside the collimator to make the light from the slit into a parallel beam. This is important, because we are using $m\lambda = d\sin\theta$ to determine the wavelength, but this equation was derived under the assumption that the light was coming in as a parallel beam. To make a parallel beam, the slit must be located accurately at the focal point of the lens. This adjustment should have already been done, but you will check later and make sure. A further advantage of using a lens in the collimator is that a telescope only works for objects far away, not nearby objects from which the reflected light is diverging strongly. The lens in the collimator forms a virtual image at infinity, on which the telescope can work.

The objective lens of the telescope focuses the light, forming a real image inside the tube. The eyepiece then acts like a magnifying glass to let you see the image. In order to see the cross-hairs and the image of the slit both in focus at the same time, the cross-hairs must be located accurately at the focal point of the objective, right on top of the image.

Adjustments

First you must check that the cross-hairs are at the focal point of the objective. If they are, then the image of the slits formed by the objective will be at the same point in space as the crosshairs. You'll be able to focus your eye on both simultaneously, and there will be no parallax error depending on the exact position of your eye. The easiest way to check this is to look through the telescope at something far away ($\gtrsim 50$ m), and move your head left and right to see if the crosshairs move relative to the image. Slide the eyepiece in and out to achieve a comfortable focus. If this adjustment is not correct, you may need to move the crosshairs in or out; this is done by sliding the tube that is just outside the eyepiece tube. (You need to use the small screwdriver to loosen the screw on the side, which is recessed inside a hole. The hole may have a dime-sized cover over it.)

The white plastic pedestal should have already been adjusted properly to get the diffraction grating oriented correctly in three dimensions, but you should check it carefully. There are some clever features built into the apparatus to help in accomplishing this. As shown in the third figure, there are three axes about which the grating could be rotated. Rotation about axis 1 is like opening a door, and this is accomplished by rotating the entire pedestal like a lazy Susan. Rotation about axes 2 (like folding down a tailgate) and 3 are accomplished using the tripod of screws underneath the pedestal. The eye-piece of the telescope is of a type called a Gauss eyepiece, with a diagonal piece of glass in it. When the grating is oriented correctly about axes 1 and 2 and the telescope is at $\theta = 0$, a beam of light that enters through the side of the eyepiece is partially reflected to the grating, and then reflected from the grating back to the eye. If these two axes are correctly adjusted, the reflected image of the crosshairs is superimposed on the crosshairs.

First get a rough initial adjustment of the pedestal by moving the telescope to 90 degrees and sighting along it like a gun to line up the grating. Now loosen the screw (not shown in the diagram) that frees the rotation of the pedestal. Put a desk lamp behind the slits of the collimator, line up the telescope with the $m = 0$ image (which may not be exactly at $\alpha = 180$ degrees), remove the desk lamp, cover the whole apparatus with the black cloth, and position a penlight so that it shines in through the hole in the side of the eyepiece. Adjust axes 1 and 2. If you're far out of adjustment, you may see part of a circle of light, which is the reflection of the penlight; start by bringing the circle of light into your field of view. When you're done, tighten the screw that keeps the pedestal from rotating. The pedestal is locked down to the tripod screws by the tension in a spring, which keeps the tips of two of the screws secure in dimples underneath the platform. Don't lower the screws too much, or the pedestal will no longer stay locked; make a habit of gently wiggling the pedestal after each adjustment to make sure it's not floating loose. Two of the spectrometers have the diagonal missing from their eyepieces, so if you have one of those, you'll have to borrow an eyepiece from another group to do this adjustment.

For the adjustment of axis 3, place a piece of masking tape so that it covers exactly half of the slits of the collimator. Put the Hg discharge tube behind the slits. The crosshairs should be near the edge of the tape in the $m = 0$ image. Move the telescope out to a large angle where you see one of the high-m Hg lines, and adjust the tripod screws so that the crosshairs are at the same height relative to the edge of the tape.

α

grating

Na or H
lamp

collimator

slit

θ

telescope

The spectrometer

slit   lens   grating   objective   crosshairs   eyepiece

collimator   telescope

Optics.

1 ("door")   2 ("tailgate")   3 ("leaning")

light

Orienting the grating.

14b   The Mass of the Electron

Apparatus

Goals

The lab is split up into parts a, b, and c. Student lab groups will do either part b or part c. This part, b, is a measurement of the mass of the electron.

Introduction

What's going on inside an atom? The question would have seemed nonsensical to physicists before the 20th century — the word "atom" is Greek for "unsplittable," and there was no evidence for subatomic particles. Only after Thomson and Rutherford had demonstrated the existence of electrons and the nucleus did the atom begin to be imagined as a tiny solar system, with the electrons moving in elliptical orbits around the nucleus under the influence of its electric field. The problem was that physicists knew very well that accelerating charges emit electromagnetic radiation, as for example in a radio antenna, so the acceleration of the electrons should have caused them to emit light, steadily lose energy, and spiral into the nucleus, all within a microsecond,.

Luckily for us, atoms do not spontaneously shrink down to nothing, but there was indeed evidence that atoms could emit light. The spectra emitted by very hot gases were observed to consist of patterns of discrete lines, each with a specific wavelength. The process of emitting light always seemed to stop short of finally annihilating the atom — why? Also, why were only those specific wavelengths emitted?

The first step toward understanding the structure of the atom was Einstein's theory that light consisted of particles (photons), whose energy was related to their frequency by the equation $E_{photon} = hf$, or substituting $f = c/\lambda$, $E_{photon} = hc/\lambda$.

According to this theory, the discrete wavelengths that had been observed came from photons with specific energies. It seemed that the atom could exist only in specific states of specific energies. To get from an initial state with energy E_i to a final state with a lower energy E_f, conservation of energy required the atom to release a photon with an energy of $E_{photon} = E_i - E_f$.

Not only could the discrete line spectra be explained, but if the atom possessed a state of least energy (called a "ground state"), then it would always end up in that state, and it could not collapse entirely. Knowing the differences between the energy levels of the atom, it was then possible to work backwards and figure out the atomic energy levels themselves.

Method

The apparatus you will use to observe the spectrum of hydrogen or nitrogen is shown in the first figure below. For a given wavelength, the grating produces diffracted light at many different angles: a central zeroth-order line at $\theta = 0$, first-order lines on both the left and right, and so on through higher-order lines at larger angles. The line of order m occurs at an angle satisfying the equation $m\lambda = d\sin\theta$.

To measure a wavelength, you will move the telescope until the diffracted first-order image of the slit is lined up with the telescope's cross-hairs and then read off the angle. Note that the angular scale on the table of the spectroscope actually gives the angle labeled α in the figure, not θ.

Eliminating systematic errors

A trick to eliminate the error due to misalignment of the angular scale is to observe the same line on both the right and the left, and take θ to be half the difference between the two angles, i.e., $\theta = (\alpha_R - \alpha_L)/2$. Because you are subtracting two angles, any source of error that adds a constant offset onto the angles is eliminated. A few of the spectrometers have their angular scales out of alignment with the collimators by as much as a full degree, but that's of absolutely no consequence if this technique is used.

Regarding the calibration of d, the first person who ever did this type of experiment simply had to make

a diffraction grating whose d was very precisely constructed. But once someone has accurately measured at least one wavelength of one emission line of one element, one can simply determine the spacing, d, of any grating using a line whose wavelength is known.

Observations

Turn on the mercury discharge tube right away, to let it get warmed up.

The second figure below shows the optics from the side, with the telescope simply looking down the throat of the collimator at $\theta = 0$. You are actually using the optics to let you see an image of the slit, not the tube itself. The point of using a telescope is that it provides angular magnification, so that a small change in angle can be seen visually.

A lens is used inside the collimator to make the light from the slit into a parallel beam. This is important, because we are using $m\lambda = d\sin\theta$ to determine the wavelength, but this equation was derived under the assumption that the light was coming in as a parallel beam. To make a parallel beam, the slit must be located accurately at the focal point of the lens. This adjustment should have already been done, but you will check later and make sure. A further advantage of using a lens in the collimator is that a telescope only works for objects far away, not nearby objects from which the reflected light is diverging strongly. The lens in the collimator forms a virtual image at infinity, on which the telescope can work.

The objective lens of the telescope focuses the light, forming a real image inside the tube. The eyepiece then acts like a magnifying glass to let you see the image. In order to see the cross-hairs and the image of the slit both in focus at the same time, the cross-hairs must be located accurately at the focal point of the objective, right on top of the image.

Setup

Skim lab 14a so you have some idea of the way the apparatus has been carefully aligned in advance by the instructor or lab technician.

A Calibration

You will use the blue line from mercury as a calibration. In theory it shouldn't matter what known line we use for calibration, but in practice there may be small aberrations in the spectrometer, and their effect is minimized by using calibration lines of nearly the same wavelength as the unknown lines to be measured.

Put the mercury tube behind the collimator. Make sure the hottest part of the tube is directly in front of the slits. You will need to use pieces of wood to get the height right. You want the tube as close to the slits as possible, and lined up with the slits as well as possible; you can adjust this while looking through the telescope at an $m = 1$ line, so as to make the line as bright as possible.

If your optics are adjusted correctly, you should be able to see the microscopic bumps and scratches on the knife edges of the collimator, and there should be no parallax of the crosshairs relative to the image of the slits.

Here is a list of the wavelengths of the most prominent visible Hg lines, in nm, to high precision.[1]

Mercury:

404.656	violet	There is a dimmer violet line nearby at 407.781 nm.
435.833	blue	
491.604	blue-green	Dim. You may also see another blue-green line that is even dimmer.
546.074	green	
	yellow	This is actually a complex set of lines, so it's not useful for calibration.

Start by making sure that you can find all of the lines lines in the correct sequence — if not, then you have probably found some first-order lines and some second-order ones. If you can find some lines but not others, use your head and search for them in the right area based on where you found the lines you did see. You may see various dim, fuzzy lights through the telescope — don't waste time chasing these, which could be coming from other tubes or from reflections. The real lines will be bright, clear and well-defined. By draping the black cloth over the discharge tube and the collimator, you can get rid of stray light that could cause problems for you or others. The discharge tubes also have holes in the back; to block the stray light from these holes, either put the two discharge tubes back to back or use one of the small "light blocks" that slide over the hole.

We will use the wavelength λ_c of the blue Hg line as a calibration. Measure its two angles α_L and α_R, and

---

[1]The table gives the wavelengths in vacuum. Although we're doing the lab in air, our goal is to find what the hydrogen or nitrogen wavelengths would have been in vacuum; by calibrating using vacuum wavelengths for mercury, we end up getting vacuum wavelengths for our unknowns as well.

check that the resulting value of θ_c is close to the approximate ones predicted in prelab question P1. The nominal value of the spacing of the grating given in that prelab question is not very accurate. Having measured θ_c, then we can sidestep the determination of the grating's spacing entirely and determine an unknown wavelength λ by using the relation

$$\lambda = \frac{\sin\theta}{\sin\theta_c}\lambda_c \qquad .$$

The angles are measured using a vernier scale, which is similar to the one on the vernier calipers you have already used in the first-semester lab course. Your final reading for an angle will consist of degrees plus minutes. (One minute of arc, abbreviated 1', is 1/60 of a degree.) The main scale is marked every 30 minutes. Your initial, rough reading is obtained by noting where the zero of the vernier scale falls on the main scale, and is of the form "xxx°0' plus a little more" or "xxx°30' plus a little more." Next, you should note which line on the vernier scale lines up most closely with one of the lines on the main scale. The corresponding number on the vernier scale tells you how many minutes of arc to add for the "plus a little more."

As a check on your results, everybody in your group should take independent readings of every angle you measure in the lab, nudging the telescope to the side after each reading. Once you have independent results for a particular angle, compare them. If they're consistent to within one or two minutes of arc, average them. If they're not consistent, figure out what went wrong.

B Spectroscopy of Hydrogen

You will study the spectrum of light emitted by the hydrogen atom, the simplest of all atoms, with just one proton and one electron. In 1885, before electrons and protons had even been imagined, a Swiss schoolteacher named Johann Balmer discovered that the wavelengths emitted by hydrogen were related by mysterious ratios of small integers. For instance, the wavelengths of the red line and the blue-green line form a ratio of exactly 20/27. Balmer even found a mathematical rule that gave all the wavelengths of the hydrogen spectrum (both the visible ones and the invisible ones that lay in the infrared and ultra-violet). The formula was completely empirical, with no theoretical basis, but clearly there were patterns lurking in the seemingly mysterious atomic spectra.

Niels Bohr showed that the energy levels of hydrogen obey a relatively simple equation,

$$E_n = -\frac{mk^2e^4}{2\hbar^2}\cdot\frac{1}{n^2}$$

where n is an integer labeling the level, k is the Coulomb constant, e is the fundamental unit of charge, $\hbar$ is Planck's constant over 2π, and m is the mass of the electron. All the energies of the photons in the emission spectrum could now be explained as differences in energy between specific states of the atom. For instance the four visible wavelengths observed by Balmer all came from cases where the atom ended up in the $n = 2$ state, dropping down from the $n = 3$, 4, 5, and 6 states.

Although the equation's sheer size may appear formidable, keep in mind that the quantity $mk^2e^4/2\hbar^2$ in front is just a numerical constant, and the variation of energy from one level to the next is of the very simple mathematical form $1/n^2$. It was because of this basic simplicity that the wavelength ratios like 20/27 occurred. The minus sign occurs because the equation includes both the electron's potential energy and its kinetic energy, and the standard choice of a reference-level for the potential energy results in negative values.

Now try swapping in the hydrogen tube in place of the mercury tube, and go through a similar process of acquainting yourself with the four lines in its visible spectrum, which are as follows:

violet   dim
purple
blue-green
red

Again you'll again have to make sure the hottest part of the tube is in front of the collimator; this requires putting books and/or blocks of wood under the discharge tube.

We will use the purple and blue-green lines to determine the mass of the electron.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

P1 The nominal (and not very accurate) spacing of the grating is stated as 600 lines per millimeter. From this information, find d, and predict the angles α_L and α_R at which you will observe the blue mercury line.

P2 Make sure you understand the first three vernier

The spectrometer

Optics.

readings in the fourth figure, and then interpret the fourth reading.

P3 For the calibration with mercury, in what sequence do you expect to see the lines on each side? Make a drawing showing the sequence of the angles as you go out from $\theta=0$.

P4 The visible lines of hydrogen come from the $3 \to 2$, $4 \to 2$, $5 \to 2$, and $6 \to 2$ transitions. Based on $E = hf$, which of these should correspond to which colors?

Self-Check

In homework problem 16-9 in *Modern Physics*, you calculated the ratio $\lambda_{blue-green}/\lambda_{purple}$. Before leaving lab, make sure that your wavelengths are consistent with this prediction, to a precision of no worse than about one part per thousand.

Analysis

Throughout your analysis, remember that this is a high-precision experiment, so you don't want to round off to less than five significant figures.

We assume that the following constants are already known:

$$e = 1.6022 \times 10^{-19} \text{ C}$$
$$k = 8.9876 \times 10^{9} \text{ N·m}^2/\text{C}^2$$
$$h = 6.6261 \times 10^{-34} \text{ J·s}$$
$$c = 2.9979 \times 10^{8} \text{ m/s}$$

The energies of the four types of visible photons emitted by a hydrogen atom equal $E_n - E_2$, where $n = 3, 4, 5$, and 6. Using the Bohr equation, we have

$$E_{photon} = A\left(\frac{1}{4} - \frac{1}{n^2}\right) \qquad ,$$

where A is the expression from the Bohr equation that depends on the mass of the electron. From the two lines you've measured, extract a value for A. If your data passed the self-check above, then you should find that these values for A agree to no worse than a few parts per thousand at worst. Compute an average value of A, and extract the mass of the electron, with error bars.

Finally, there is a small correction that should be made to the result for the mass of the electron because actually the proton isn't infinitely massive compared to the electron; in terms of the quantity m given by the equation on page 60, the mass of the electron, m_e, would actually be given by $m_e = m/(1 - m/m_p)$, where m_p is the mass of the proton, 1.6726×10^{-27} kg.

220   230   240
=221°0'

0   5   10   15   20   25   30

220   230   240
=221°30'

0   5   10   15   20   25   30

220   230   240
=221°40'

0   5   10   15   20   25   30

220   230   240
=?

0   5   10   15   20   25   30

Prelab question 2.

14c   The Nitrogen Molecule

Apparatus

Goals

The lab has three parts. Each group will only do two parts. In this lab, part c, you will use an energy sum to test a hypothesis about the energy levels of the nitrogen molecule, N_2.

Method

The apparatus you will use to observe the spectrum of hydrogen or nitrogen is shown in the first figure below. For a given wavelength, the grating produces diffracted light at many different angles: a central zeroth-order line at $\theta = 0$, first-order lines on both the left and right, and so on through higher-order lines at larger angles. The line of order m occurs at an angle satisfying the equation $m\lambda = d \sin \theta$.

To measure a wavelength, you will move the telescope until the diffracted first-order image of the slit is lined up with the telescope's cross-hairs and then read off the angle. Note that the angular scale on the table of the spectroscope actually gives the angle labeled α in the figure, not θ.

Eliminating systematic errors

A trick to eliminate the error due to misalignment of the angular scale is to observe the same line on both the right and the left, and take θ to be half the difference between the two angles, i.e., $\theta = (\alpha_R - \alpha_L)/2$. Because you are subtracting two angles, any source of error that adds a constant offset onto the angles is eliminated. A few of the spectrometers have their angular scales out of alignment with the collimators by as much as a full degree, but that's of absolutely

no consequence if this technique is used.

Regarding the calibration of d, the first person who ever did this type of experiment simply had to make a diffraction grating whose d was very precisely constructed. But once someone has accurately measured at least one wavelength of one emission line of one element, one can simply determine the spacing, d, of any grating using a line whose wavelength is known.

B Energy Sums

The nitrogen discharge tube is housed in a green plastic carousel. With the power off, rotate the carousel so that the nitrogen tube is the one that is in the active position, and then turn on the power. If you hold a diffraction grating up to your eye and look at the tube, you will see a remarkable spectrum, unlike the visible light spectrum of almost any other gas. This is because the N_2 molecule has an extremely strong bond, requiring twice the energy to break compared to otherwise similar gases such as H_2 or O_2. Whereas these other gases would break up into individual atoms under the extreme conditions present in a discharge tube, the nitrogen molecule holds together, so that you are seeing the spectrum of the molecule, not the atom. For this reason, the spectrum of nitrogen contains a large number of lines.


But these lines are not random. They occur in sets, each of which looks like a comb with an approximately equal spacing between the "teeth." The figure shows a portion of the spectrum, including three sets of lines, which I have labeled r (red), o (orange to green), and g (green).

I have spent some time trying to interpret the origin of these lines, and I believe the interpretation is something like this. Each of these lines is the emission of a photon as the molecule goes from an initial state to a final state that has less energy. The initial state has some energy because the electrons are in an excited state (labeled B by spectroscopists)

and also some energy because the molecule is vibrating, like two masses connected by a spring. The final state has the electrons in a lower-energy state (labeled A), and is also vibrating. The initial and final electronic states B and A are the same in all cases, but the vibrational states differ. An idealized quantum-mechanical vibrator turns out to have a series of energy states like a ladder with nearly evenly spaced rungs. States higher on the "ladder" are vibrating more violently — classically, they vibrate with greater amplitude. The rungs of the vibrational ladder are labeled $v = 0$, 1, 2, and so on. (Because of the Heisenberg uncertainty principle, some vibrational energy is present even in the $v = 0$ state.) I think the states in the red set are from a state v to a state $v - 3$, i.e., a change of 3 units in the vibrational quantum number. The o set would be a change of 4, and g a change of 5.

This hypothesis can be tested as follows. If it is true, we could pick out a set of energy levels like the following example:

The letters e, f, g, and h are the energy differences that would be observed as the energies of the photons. In a set like this, we would have

$$h - g = f - e,$$

since each side of the equation would be equal to the energy difference between the $v = 4$ and 5 states of ladder A. Try to find a set of lines that would be consistent with this interpretation. This may require some trial and error, but I think it may work if e is one of the lines near the middle of the r set, g is the orange line that is fourth from the short-wavelength end of the o set, f is the fifth in that set, and h is in the g set.

C Calibration

You will use the yellow line from helium as a calibration. In theory it shouldn't matter what known line we use for calibration, but in practice there may be small aberrations in the spectrometer, and their effect is minimized by using calibration lines of nearly the same wavelength as the unknown lines to be measured.

Put the helium tube behind the collimator. Make sure the hottest part of the tube is directly in front of the slits. You will need to use pieces of wood to get the height right. You want the tube as close to the slits as possible, and lined up with the slits as well as possible; you can adjust this while looking through the telescope at an $m = 1$ line, so as to make the line as bright as possible.

If your optics are adjusted correctly, you should be able to see the microscopic bumps and scratches on the knife edges of the collimator, and there should be no parallax of the crosshairs relative to the image of the slits.

Here is a list of the wavelengths of the most prominent visible He lines, in nm, to high precision.[1]

Helium:
447.148	bright blue-purple
471.314	dim blue
492.193	dim green
501.567	bright green
587.562	yellow
667.815	dim red
706.5	very dim red

Start by making sure that you can find all of the lines in the correct sequence — if not, then you have probably found some first-order lines and some second-order ones. If you can find some lines but not others, use your head and search for them in the right area based on where you found the lines you did see. You may see various dim, fuzzy lights through the telescope — don't waste time chasing these, which could be coming from other tubes or from reflections. The real lines will be bright, clear and well-defined. By draping the black cloth over the discharge tube and the collimator, you can get rid of stray light that could cause problems for you or others. The discharge tubes also have holes in the back; to block the stray light from these holes, either put the two discharge tubes back to back or use one of the small

---

[1]The table gives the wavelengths in vacuum. Although we're doing the lab in air, our goal is to find what the nitrogen wavelengths would have been in vacuum; by calibrating using vacuum wavelengths for mercury, we end up getting vacuum wavelengths for our unknowns as well.

"light blocks" that slide over the hole.

We will use the wavelength λ_c of the hellow He line as a calibration. Measure its two angles α_L and α_R, and check that the resulting value of θ_c is close to the approximate ones predicted in prelab question P1. The nominal value of the spacing of the grating given in that prelab question is not very accurate. Having measured θ_c, then we can sidestep the determination of the grating's spacing entirely and determine an unknown wavelength λ by using the relation

$$\lambda = \frac{\sin\theta}{\sin\theta_c}\lambda_c \qquad .$$

The angles are measured using a vernier scale, which is similar to the one on the vernier calipers you have already used in the first-semester lab course. Your final reading for an angle will consist of degrees plus minutes. (One minute of arc, abbreviated 1', is 1/60 of a degree.) The main scale is marked every 30 minutes. Your initial, rough reading is obtained by noting where the zero of the vernier scale falls on the main scale, and is of the form "xxx°0' plus a little more" or "xxx°30' plus a little more." Next, you should note which line on the vernier scale lines up most closely with one of the lines on the main scale. The corresponding number on the vernier scale tells you how many minutes of arc to add for the "plus a little more."

As a check on your results, everybody in your group should take independent readings of every angle you measure in the lab, nudging the telescope to the side after each reading. Once you have independent results for a particular angle, compare them. If they're consistent to within one or two minutes of arc, average them. If they're not consistent, figure out what went wrong.

Status as of August 2018

In spring of 2018, my students and I worked on measuring and interpreting this spectrum. A summary of our results is in a Google Docs spreadsheet at `goo.gl/akrbcY`. There is a basic explanation of the physics in Simple Nature section 14.2. More detailed information about my interpretation of the lines is at
`physics.stackexchange.com/a/334451/4552`.

In the notation used in the material on stackexchange, the states are labeled with a quantum number v. The energy sum based on our data come out quite nice for $v = 9$ and 8 going to 5 and 4, accurate to about 0.3%, which is reasonable for this

technique. The energy sum for $v = 10$ and 9 going to 6 and 5 also works well, but with slightly higher error, maybe partly because 10-5 is a doublet. There are two dim green lines that in my labeling system would be g0 and g-1. These may be the ones we would need in order to get a couple more energy sums.

I experimented with doing the measurements photographically. A student took a photo of a diffraction pattern using a cell phone. I first did a rough calibration against student data, then improved this calibration by doing a linear fit to my own spectrometer data. This worked fairly well, but doesn't work without actual spectrometer data, which are needed for calibration.

We have digital spectrometers which may be helpful here and are worth trying. Their resolution is supposed to be about 1 or 1.5 nm, which is an order of magnitude worse than the analog spectrometers, but may be adequate for this purpose, and they can display a spectrum as a graph, which may be help enough to make up for the lower resolution.

I couldn't resolve the green band. I asked a couple of students the next day, and they seemed to think that it was doable to resolve these lines. Possibly my tube was behaving differently than theirs, but I'm not sure I believe their data. The red and orange bands came out nice, all wavelengths being within a few tenths of a nm of Lofthus's values.

Prelab

The point of the prelab questions is to make sure you understand what you're doing, why you're doing it, and how to avoid some common mistakes. If you don't know the answers, make sure to come to my office hours before lab and get help! Otherwise you're just setting yourself up for failure in lab.

The week before you are to do the lab, briefly familiarize yourself visually with the apparatus.

P1 The nominal (and not very accurate) spacing of the grating is stated as 600 lines per millimeter. From this information, find d, and predict the angles α_L and α_R at which you will observe the yellow helium line.

P2 Make sure you understand the first three vernier readings in the fourth figure, and then interpret the fourth reading.

**P3** For the calibration with helium, in what sequence do you expect to see the lines on each side?

The spectrometer

Optics.

Make a drawing showing the sequence of the angles
as you go out from $\theta=0$.

1 ("door")   2 ("tailgate")   3 ("leaning")

Orienting the grating.

=221°0'

=221°30'

=221°40'

=?

Prelab question 2.

Appendix 1: Format of Lab Writeups

Lab reports must be three pages or less, not counting your raw data. The format should be as follows:

Title

Raw data — *Keep actual observations separate from what you later did with them.*
These are the results of the measurements you take down during the lab, hence they come first. Write your raw data directly in your lab book; don't write them on scratch paper and recopy them later. Don't use pencil. The point is to separate facts from opinions, observations from inferences.

Procedure — *Did you have to create your own methods for getting some of the raw data?*
Do not copy down the procedure from the manual. In this section, you only need to explain any methods you had to come up with on your own, or cases where the methods suggested in the handout didn't work and you had to do something different. Don't write anything here unless you think I will really care and want to change how we do the lab in the future. In most cases this section can be totally blank. Do not discuss how you did your calculations here, just how you got your raw data.

Abstract — *What did you find out? Why is it important?*
The "abstract" of a scientific paper is a *short* paragraph at the top that summarizes the experiment's results in a few sentences.

Many of our labs are comparisons of theory and experiment. The abstract for such a lab needs to say whether you think the experiment was consistent with theory, or not consistent with theory. If your results deviated from the ideal equations, don't be afraid to say so. After all, this is real life, and many of the equations we learn are only approximations, or are only valid in certain circumstances. However, (1) if you simply mess up, it is your responsibility to realize it in lab and do it again, right; (2) you will never get exact agreement with theory, because measurements are not perfectly exact — the important issue is whether your results agree with theory to roughly within the error bars.

The abstract is not a statement of what you hoped to find out. It's a statement of what you *did* find out. It's like the brief statement at the beginning of a debate: "The U.S. should have free trade with China." It's not this: "In this debate, we will discuss whether the U.S. should have free trade with China."

If this is a lab that has just one important numerical result (or maybe two or three of them), put them in your abstract, with error bars where appropriate. There should normally be no more than two to four numbers here. Do not recapitulate your raw data here — this is for your final results.

If you're presenting a final result with error bars, make sure that the number of significant figures is consistent with your error bars. For example, if you write a result as 323.54 ± 6 m/s, that's wrong. Your error bars say that you could be off by 6 in the ones' place, so the 5 in the tenths' place and the four in the hundredths' place are completely meaningless.

If you're presenting a number in scientific notation, with error bars, don't do it like this

$$1.234 \times 10^{-89} \text{ m/s} \pm 3 \times 10^{-92} \text{ m/s} \qquad ,$$

do it like this

$$(1.234 \pm 0.003) \times 10^{-89} \text{ m/s} \qquad ,$$

so that we can see easily which digit of the result the error bars apply to.

Calculations and Reasoning — *Convince me of what you claimed in your abstract.*
Often this section consists of nothing more than the calculations that you started during lab. If those calculations are clear enough to understand, and there is nothing else of interest to explain, then it is not necessary to write up a separate narrative of your analysis here. If you have a long series of similar calculations, you may just show one as a sample. If your prelab involved deriving equations that you will need, repeat them here without the derivation.

In some labs, you will need to go into some detail here by giving logical arguments to convince me that the statements you made in the abstract follow logically from your data. Continuing the debate metaphor, if your abstract said the U.S. should have free trade with China, this is the rest of the debate, where you convince me, based on data and logic, that we should have free trade.

Appendix 2: Basic Error Analysis

No measurement is perfectly exact.

One of the most common misconceptions about science is that science is "exact." It is always a struggle to get beginning science students to believe that no measurement is perfectly correct. They tend to think that if a measurement is a little off from the "true" result, it must be because of a mistake — if a pro had done it, it would have been right on the mark. Not true!

What scientists can do is to estimate just how far off they might be. This type of estimate is called an error bar, and is expressed with the $\pm$ symbol, read "plus or minus." For instance, if I measure my dog's weight to be 52 ± 2 pounds, I am saying that my best estimate of the weight is 52 pounds, and I think I could be off by roughly 2 pounds either way. The term "error bar" comes from the conventional way of representing this range of uncertainty of a measurement on a graph, but the term is also used when no graph is involved.

Some very good scientific work results in measurements that nevertheless have large error bars. For instance, the best measurement of the age of the universe is now 15 ± 5 billion years. That may not seem like wonderful precision, but the people who did the measurement knew what they were doing. It's just that the only available techniques for determining the age of the universe are inherently poor.

Even when the techniques for measurement are very precise, there are still error bars. For instance, electrons act like little magnets, and the strength of a very weak magnet such as an individual electron is customarily measured in units called Bohr magnetons. Even though the magnetic strength of an electron is one of the most precisely measured quantities ever, the best experimental value still has error bars: $1.0011596524 \pm 0.0000000002$ Bohr magnetons.

There are several reasons why it is important in scientific work to come up with a numerical estimate of your error bars. If the point of your experiment is to test whether the result comes out as predicted by a theory, you know there will always be some disagreement, even if the theory is absolutely right. You need to know whether the measurement is reasonably consistent with the theory, or whether the discrepancy is too great to be explained by the limitations of the measuring devices.

Another important reason for stating results with error bars is that other people may use your measurement for purposes you could not have anticipated. If they are to use your result intelligently, they need to have some idea of how accurate it was.

Error bars are not absolute limits.

Error bars are not absolute limits. The true value may lie outside the error bars. If I got a better scale I might find that the dog's weight is 51.3 ± 0.1 pounds, inside my original error bars, but it's also possible that the better result would be 48.7 ± 0.1 pounds. Since there's always some chance of being off by a somewhat more than your error bars, or even a lot more than your error bars, there is no point in being extremely conservative in an effort to make absolutely sure the true value lies within your stated range. When a scientist states a measurement with error bars, she is not saying "If the true value is outside this range, I deserve to be drummed out of the profession." If that was the case, then every scientist would give ridiculously inflated error bars to avoid having her career ended by one fluke out of hundreds of published results. What scientists are communicating to each other with error bars is a typical amount by which they might be off, not an upper limit.

The important thing is therefore to define error bars in a standard way, so that different people's statements can be compared on the same footing. By convention, it is usually assumed that people estimate their error bars so that about two times out of three, their range will include the true value (or the results of a later, more accurate measurement with an improved technique).

Random and systematic errors.

Suppose you measure the length of a sofa with a tape measure as well as you can, reading it off to the nearest millimeter. If you repeat the measurement again, you will get a different answer. (This is assuming that you don't allow yourself to be psychologically biased to repeat your previous answer, and that 1 mm is about the limit of how well you can see.) If you kept on repeating the measurement,

you might get a list of values that looked like this:

203.1 cm	203.4	202.8	203.3	203.2
203.4		203.1	202.9	202.9	203.1

Variations of this type are called random errors, because the result is different every time you do the measurement.

The effects of random errors can be minimized by averaging together many measurements. Some of the measurements included in the average are too high, and some are too low, so the average tends to be better than any individual measurement. The more measurements you average in, the more precise the average is. The average of the above measurements is 203.1 cm. Averaging together many measurements cannot completely eliminate the random errors, but it can reduce them.

On the other hand, what if the tape measure was a little bit stretched out, so that your measurements always tended to come out too low by 0.3 cm? That would be an example of a systematic error. Since the systematic error is the same every time, averaging didn't help us to get rid of it. You probably had no easy way of finding out exactly the amount of stretching, so you just had to suspect that there might a systematic error due to stretching of the tape measure.

Some scientific writers make a distinction between the terms "accuracy" and "precision." A precise measurement is one with small random errors, while an accurate measurement is one that is actually close to the true result, having both small random errors and small systematic errors. Personally, I find the distinction is made more clearly with the more memorable terms "random error" and "systematic error."

The ± sign used with error bars normally implies that random errors are being referred to, since random errors could be either positive or negative, whereas systematic errors would always be in the same direction.

The goal of error analysis

Very seldom does the final result of an experiment come directly off of a clock, ruler, gauge or meter. It is much more common to have raw data consisting of direct measurements, and then calculations based on the raw data that lead to a final result. As an example, if you want to measure your car's gas mileage, your raw data would be the number of gallons of gas consumed and the number of miles you went. You would then do a calculation, dividing

small random errors,
small systematic error

large random errors,
small systematic error

small random errors,
large systematic error

miles by gallons, to get your final result. When you communicate your result to someone else, they are completely uninterested in how accurately you measured the number of miles and how accurately you measured the gallons. They simply want to know how accurate your final result was. Was it 22 ± 2 mi/gal, or 22.137 ± 0.002 mi/gal?

Of course the accuracy of the final result is ultimately based on and limited by the accuracy of your raw data. If you are off by 0.2 gallons in your measurement of the amount of gasoline, then that amount of error will have an effect on your final result. We say that the errors in the raw data "propagate" through the calculations. When you are requested to do "error analysis" in a lab writeup, that means that you

are to use the techniques explained below to determine the error bars on your final result. There are two sets of techniques you'll need to learn:

> techniques for finding the accuracy of your raw data

> techniques for using the error bars on your raw data to infer error bars on your final result

Estimating random errors in raw data

We now examine three possible techniques for estimating random errors in your original measurements, illustrating them with the measurement of the length of the sofa.

Method #1: Guess

If you're measuring the length of the sofa with a metric tape measure, then you can probably make a reasonable guess as to the precision of your measurements. Since the smallest division on the tape measure is one millimeter, and one millimeter is also near the limit of your ability to see, you know you won't be doing better than ± 1 mm, or 0.1 cm. Making allowances for errors in getting tape measure straight and so on, we might estimate our random errors to be a couple of millimeters.

Guessing is fine sometimes, but there are at least two ways that it can get you in trouble. One is that students sometimes have too much faith in a measuring device just because it looks fancy. They think that a digital balance must be perfectly accurate, since unlike a low-tech balance with sliding weights on it, it comes up with its result without any involvement by the user. That is incorrect. No measurement is perfectly accurate, and if the digital balance only displays an answer that goes down to tenths of a gram, then there is no way the random errors are any smaller than about a tenth of a gram.

Another way to mess up is to try to guess the error bars on a piece of raw data when you really don't have enough information to make an intelligent estimate. For instance, if you are measuring the range of a rifle, you might shoot it and measure how far the bullet went to the nearest centimeter, concluding that your random errors were only ± 1 cm. In reality, however, its range might vary randomly by fifty meters, depending on all kinds of random factors you don't know about. In this type of situation, you're better off using some other method of estimating your random errors.

Method #2: Repeated Measurements and the Two-Thirds Rule

If you take repeated measurements of the same thing, then the amount of variation among the numbers can tell you how big the random errors were. This approach has an advantage over guessing your random errors, since it automatically takes into account all the sources of random error, even ones you didn't know were present.

Roughly speaking, the measurements of the length of the sofa were mostly within a few mm of the average, so that's about how big the random errors were. But let's make sure we are stating our error bars according to the convention that the true result will fall within our range of errors about two times out of three. Of course we don't know the "true" result, but if we sort out our list of measurements in order, we can get a pretty reasonable estimate of our error bars by taking half the range covered by the middle two thirds of the list. Sorting out our list of ten measurements of the sofa, we have

| 202.8 cm | 202.9 | 202.9 | 203.1 | 203.1 |
| 203.1 | 203.2 | 203.3 | 203.4 | 203.4 |

Two thirds of ten is about 6, and the range covered by the middle six measurements is 203.3 cm - 202.9 cm, or 0.4 cm. Half that is 0.2 cm, so we'd estimate our error bars as ± 0.2 cm. The average of the measurements is 203.1 cm, so your result would be stated as 203.1 ± 0.2 cm.

One common mistake when estimating random errors by repeated measurements is to round off all your measurements so that they all come out the same, and then conclude that the error bars were zero. For instance, if we'd done some overenthusiastic rounding of our measurements on the sofa, rounding them all off to the nearest cm, every single number on the list would have been 203 cm. That wouldn't mean that our random errors were zero! The same can happen with digital instruments that automatically round off for you. A digital balance might give results rounded off to the nearest tenth of a gram, and you may find that by putting the same object on the balance again and again, you always get the same answer. That doesn't mean it's perfectly precise. Its precision is no better than about ± 0.1 g.

Method #3: Repeated Measurements and the Standard Deviation

The most widely accepted method for measuring error bars is called the standard deviation. Here's how the method works, using the sofa example again.

(1) Take the average of the measurements.

$$\text{average} = 203.1 \text{ cm}$$

(2) Find the difference, or "deviation," of each measurement from the average.

−0.3 cm	−0.2	−0.2	0.0	0.0
0.0	0.1	0.1	0.3	0.3

(3) Take the square of each deviation.

0.09 cm^2	0.04	0.04	0.00	0.00
0.00	0.01	0.01	0.09	0.09

(4) Average together all the squared deviations.

$$\text{average} = 0.04 \text{ cm}^2$$

(5) Take the square root. This is the standard deviation.

$$\text{standard deviation} = 0.2 \text{ cm}$$

If we're using the symbol x for the length of the couch, then the result for the length of the couch would be stated as $x = 203.1 \pm 0.2$ cm, or $x = 203.1$ cm and $\sigma_x = 0.2$ cm. Since the Greek letter sigma (σ) is used as a symbol for the standard deviation, a standard deviation is often referred to as "a sigma."

Step (3) may seem somewhat mysterious. Why not just skip it? Well, if you just went straight from step (2) to step (4), taking a plain old average of the deviations, you would find that the average is zero! The positive and negative deviations always cancel out exactly. Of course, you could just take absolute values instead of squaring the deviations. The main advantage of doing it the way I've outlined above are that it is a standard method, so people will know how you got the answer. (Another advantage is that the standard deviation as I've described it has certain nice mathematical properties.)

A common mistake when using the standard deviation technique is to take too few measurements. For instance, someone might take only two measurements of the length of the sofa, and get 203.4 cm and 203.4 cm. They would then infer a standard deviation of zero, which would be unrealistically small because the two measurements happened to come out the same.

In the following material, I'll use the term "standard deviation" as a synonym for "error bar," but that does not imply that you must always use the standard deviation method rather than the guessing method or the 2/3 rule.

There is a utility on the class's web page for calculating standard deviations.

Probability of deviations

You can see that although 0.2 cm is a good figure for the typical size of the deviations of the measurements of the length of the sofa from the average, some of the deviations are bigger and some are smaller. Experience has shown that the following probability estimates tend to hold true for how frequently deviations of various sizes occur:

> 1 standard deviation about 1 times out of 3

> 2 standard deviations about 1 time out of 20

> 3 standard deviations about 1 in 500

> 4 standard deviations about 1 in 16,000

> 5 standard deviations about 1 in 1,700,000

The probability of various sizes of deviations, shown graphically. Areas under the bell curve correspond to probabilities. For example, the probability that the measurement will deviate from the truth by less than one standard deviation ($\pm 1\sigma$) is about $34 \times 2 = 68\%$, or about 2 out of 3. (J. Kemp, P. Strandmark, Wikipedia.)

Example: How significant?
In 1999, astronomers Webb et al. claimed to have found evidence that the strength of electrical forces in the ancient universe, soon after the big bang, was slightly weaker than it is today. If correct, this would be the first example ever discovered in which the laws of physics changed over time. The difference was very small, 5.7 ± 1.0 parts per million, but still highly statistically significant. Dividing, we get $(5.7 - 0)/1.0 = 5.7$ for the number of standard deviations by which their measurement was different from the expected result of zero. Looking at the table above, we see that if the true value really was zero, the chances of this happening would be less than one in a million. In general, five standard deviations ("five sigma") is considered the gold standard for statistical significance.

This is an example of how we test a hypothesis statistically, find a probability, and interpret the probability. The probability we find is the probability that our results

would differ this much from the hypothesis, if the hypothesis was true. It's not the probability that the hypothesis is true or false, nor is it the probability that our experiment is right or wrong.

However, there is a twist to this story that shows how statistics always have to be taken with a grain of salt. In 2004, Chand et al. redid the measurement by a more precise technique, and found that the change was 0.6 ± 0.6 parts per million. This is only one standard deviation away from the expected value of 0, which should be interpreted as being statistically consistent with zero. If you measure something, and you think you know what the result is supposed to be theoretically, then one standard deviation is the amount you typically *expect* to be off by — that's why it's called the "standard" deviation. Moreover, the Chand result is wildly statistically inconsistent with the Webb result (see the example on page 79), which means that one experiment or the other is a mistake. Most likely Webb at al. underestimated their random errors, or perhaps there were systematic errors in their experiment that they didn't realize were there.

Precision of an average

We decided that the standard deviation of our measurements of the length of the couch was 0.2 cm, i.e., the precision of each individual measurement was about 0.2 cm. But I told you that the average, 203.1 cm, was more precise than any individual measurement. How precise is the average? The answer is that the standard deviation of the average equals

$$\frac{\text{standard deviation of one measurement}}{\sqrt{\text{number of measurements}}} .$$

(An example on page 78 gives the reasoning that leads to the square root.) That means that you can theoretically measure anything to any desired precision, simply by averaging together enough measurements. In reality, no matter how small you make your random error, you can't get rid of systematic errors by averaging, so after a while it becomes pointless to take any more measurements.

Appendix 3: Propagation of Errors

Propagation of the error from a single variable

In the previous appendix we looked at techniques for estimating the random errors of raw data, but now we need to know how to evaluate the effects of those random errors on a final result calculated from the raw data. For instance, suppose you are given a cube made of some unknown material, and you are asked to determine its density. Density is defined as $\rho = m/v$ (ρ is the Greek letter "rho"), and the volume of a cube with edges of length b is $v = b^3$, so the formula

$$\rho = m/b^3$$

will give you the density if you measure the cube's mass and the length of its sides. Suppose you measure the mass very accurately as $m = 1.658 \pm 0.003$ g, but you know $b = 0.85 \pm 0.06$ cm with only two digits of precision. Your best value for ρ is 1.658 g$/(0.85$ cm$)^3 = 2.7$ g/cm^3.

How can you figure out how precise this value for ρ is? We've already made sure not to keep more than twosignificant figures for ρ, since the less accurate piece of raw data had only two significant figures. We expect the last significant figure to be somewhat uncertain, but we don't yet know how uncertain. A simple method for this type of situation is simply to change the raw data by one sigma, recalculate the result, and see how much of a change occurred. In this example, we add 0.06 cm to b for comparison.

$b = 0.85$ cm	gave	$\rho = 2.7$ g/cm^3
$b = 0.91$ cm	gives	$\rho = 2.2$ g/cm^3

The resulting change in the density was 0.5 g/cm^3, so that is our estimate for how much it could have been off by:

$$\rho = 2.7 \pm 0.5 \text{ g/cm}^3 \qquad .$$

Propagation of the error from several variables

What about the more general case in which no one piece of raw data is clearly the main source of error? For instance, suppose we get a more accurate measurement of the edge of the cube, $b = 0.851 \pm 0.001$ cm. In percentage terms, the accuracies of m and

b are roughly comparable, so both can cause significant errors in the density. The following more general method can be applied in such cases:

(1) Change one of the raw measurements, say m, by one standard deviation, and see by how much the final result, ρ, changes. Use the symbol Q_m for the absolute value of that change.

$m = 1.658$ g	gave	$\rho = 2.690$ g/cm^3
$m = 1.661$ g	gives	$\rho = 2.695$ g/cm^3

$Q_m =$ change in $\rho = 0.005$ g/cm^3

(2) Repeat step (1) for the other raw measurements.

$b = 0.851$ cm	gave	$\rho = 2.690$ g/cm^3
$b = 0.852$ cm	gives	$\rho = 2.681$ g/cm^3

$Q_b =$ change in $\rho = 0.009$ g/cm^3

(3) The error bars on ρ are given by the formula

$$\sigma_\rho = \sqrt{Q_m^2 + Q_b^2} \qquad ,$$

yielding $\sigma_\rho = 0.01$ g/cm^3. Intuitively, the idea here is that if our result could be off by an amount Q_m because of an error in m, and by Q_b because of b, then if the two errors were in the same direction, we might by off by roughly $|Q_m| + |Q_b|$. However, it's equally likely that the two errors would be in opposite directions, and at least partially cancel. The expression $\sqrt{Q_m^2 + Q_b^2}$ gives an answer that's smaller than $Q_m + Q_b$, representing the fact that the cancellation might happen.

The final result is $\rho = 2.69 \pm 0.01$ g/cm^3.

Example: An average
On page 76 I claimed that averaging a bunch of measurements reduces the error bars by the square root of the number of measurements. We can now see that this is a special case of propagation of errors.

For example, suppose Alice measures the circumference c of a guinea pig's waist to be 10 cm, Using the guess method, she estimates that her error bars are about ± 1 cm (worse than the normal normal ~ 1 mm error bars for a tape measure, because the guinea pig was squirming). Bob then measures the same thing, and gets 12 cm. The average is computed as

$$c = \frac{A + B}{2} \qquad ,$$

where A is Alice's measurement, and B is Bob's, giving 11 cm. If Alice had been off by one standard deviation (1 cm), it would have changed the average by 0.5

cm, so we have $Q_A = 0.5$ cm, and likewise $Q_B = 0.5$ cm. Combining these, we find $\sigma_c = \sqrt{Q_A^2 + Q_B^2} = 0.7$ cm, which is simply $(1.0 \text{ cm})/\sqrt{2}$. The final result is $c = (11.0 \pm 0.7)$ cm. (This violates the usual rule for significant figures, which is that the final result should have no more sig figs than the least precise piece of data that went into the calculation. That's okay, because the sig fig rules are just a quick and dirty way of doing propagation of errors. We've done real propagation of errors in this example, and it turns out that the error is in the first decimal place, so the 0 in that place is entitled to hold its head high as a real sig fig, albeit a relatively imprecise one with an uncertainty of ± 7.)

Example: The difference between two measurements
In the example on page 75, we saw that two groups of scientists measured the same thing, and the results were $W = 5.7 \pm 1.0$ for Webb et al. and $C = 0.6 \pm 0.6$ for Chand et al. It's of interest to know whether the difference between their two results is small enough to be explained by random errors, or so big that it couldn't possibly have happened by chance, indicating that someone messed up. The figure shows each group's results, with error bars, on the number line. We see that the two sets of error bars don't overlap with one another, but error bars are not absolute limits, so it's perfectly possible to have non-overlapping error bars by chance, but the gap between the error bars is very large compared to the error bars themselves, so it looks implausible that the results could be statistically consistent with one another. I've tried to suggest this visually with the shading underneath the data-points.

To get a sharper statistical test, we can calculate the difference d between the two results,

$$d = W - C \quad ,$$

which is 5.1. Since the operation is simply the subtraction of the two numbers, an error in either input just causes an error in the output that is of the same size. Therefore we have $Q_W = 1.0$ and $Q_C = 0.6$, resulting in $\sigma_d = \sqrt{Q_W^2 + Q_C^2} = 1.2$. We find that the difference between the two results is $d = 5.1 \pm 1.2$, which differs from zero by $5.1/1.2 \approx 4$ standard deviations. Looking at the table on page 75, we see that the chances that d would be this big by chance are extremely small, less than about one in ten thousand. We can conclude to a high level of statistical confidence that the two groups' measurements are inconsistent with one another, and that one group is simply wrong.

Appendix 4: Graphing

Review of Graphing

Many of your analyses will involve making graphs. A graph can be an efficient way of presenting data visually, assuming you include all the information needed by the reader to interpret it. That means labeling the axes and indicating the units in parentheses, as in the example. A title is also helpful. Make sure that distances along the axes correctly represent the differences in the quantity being plotted. In the example, it would not have been correct to space the points evenly in the horizontal direction, because they were not actually measured at equally spaced points in time.

motion of a falling object

Graphing on a Computer

Making graphs by hand in your lab notebook is fine, but in some cases you may find it saves you time to do graphs on a computer. For computer graphing, I recommend LibreOffice, which is free, open-source software. It's installed on the computers in rooms 416 and 418. Because LibreOffice is free, you can download it and put it on your own computer at home without paying money. If you already know Excel, it's very similar — you almost can't tell it's a different program.

Here's a brief rundown on using LibreOffice:

> On Windows, go to the Start menu and choose All Programs, LibreOffice, and LibreOffice Calc. On Linux, do Applications, Office, OpenOffice, Spreadsheet.

Type in your x values in the first column, and your y values in the second column. For scientific notation, do, e.g., 5.2e-7 to represent 5.2×10^{-7}.

Select those two columns using the mouse.

From the Insert menu, do Object:Chart.

When it offers you various styles of graphs to choose from, choose the icon that shows a scatter plot, with dots on it (XY Chart).

Adjust the scales so the actual data on the plot is as big as possible, eliminating wasted space. To do this, double-click on the graph so that it's surrounded by a gray border. Then do Format, Axis, X Axis or Y Axis, Scale.

If you want error bars on your graph you can either draw them in by hand or put them in a separate column of your spreadsheet and doing Insert, Y Error Bars, Cell Range. Under Parameters, check "Same value for both." Click on the icon, and then use the mouse in the spreadsheet to select the cells containing the error bars.

Fitting a Straight Line to a Graph by Hand

Often in this course you will end up graphing some data points, fitting a straight line through them with a ruler, and extracting the slope.
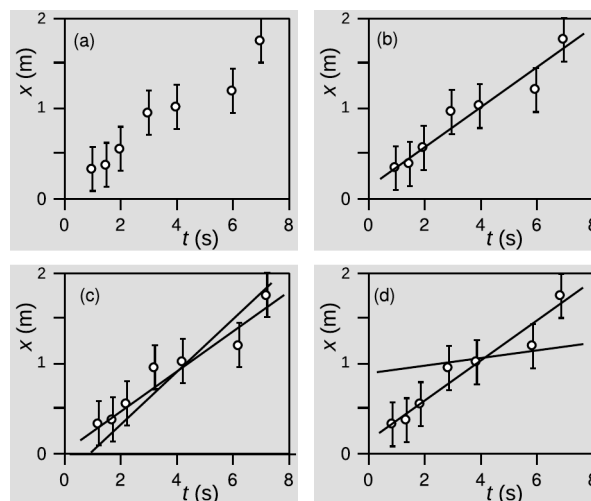
In this example, panel (a) shows the data, with error bars on each data point. Panel (b) shows a best fit, drawn by eye with a ruler. The slope of this best fit line is 100 cm/s. Note that the slope should be extracted from the line itself, not from two data points. The line is more reliable than any pair of individual data points.

In panel (c), a "worst believable fit" line has been drawn, which is as different in slope as possible from the best fit, while still pretty much staying consistent the data (going through or close to most of the error bars). Its slope is 60 cm/s. We can therefore estimate that the precision of our slope is +40 cm/s.

There is a tendency when drawing a "worst believable fit" line to draw instead an "unbelievably crazy fit" line, as in panel (d). The line in panel (d), with a very small slope, is just not believable compared to the data — it is several standard deviations away from most of the data points.

Fitting a Straight Line to a Graph on a Computer

It's also possible to fit a straight line to a graph using computer software such as LibreOffice.

To do this, first double-click on the graph so that a gray border shows up around it. Then right-click on a data-point, and a menu pops up. Choose Insert Trend Line.[1] choose Linear, and check the box for Show equation.

How accurate is your slope? A method for getting error bars on the slope is to artificially change one of your data points to reflect your estimate of how much it could have been off, and then redo the fit and find the new slope. The change in the slope tells you the error in the slope that results from the error in this data-point. You can then repeat this for the other points and proceed as in appendix 3.

An alternative method is to use the LINEST function that is available in many spreadsheet programs. For a description, see `tinyurl.com/ya7wmdft`. Create the following formula in one cell of your spreadsheet: =Linest(y-values,x-value, True.True). Then, in excel, you need to press alt+ctrl+enter. In google sheets, press enter. A table with two columns and five rows will appear. The first number in the first column is the slope of the graph, and the second

---

[1] "Trend line" is scientifically illiterate terminology that originates from Microsoft Office, which LibreOffice slavishly copies. If you don't want to come off as an ignoramus, call it a "fit" or "line of best fit."

number in the first column is the error in the slope.

In some cases, such as the absolute zero lab and the photoelectric effect lab, it's very hard to tell how accurate your raw data are *a priori*; in these labs, you can use the typical amount of deviation of the points from the line as an estimate of their accuracy.

Comparing Theory and Experiment

Figures (e) through (h) are examples of how we would compare theory and experiment on a graph. The convention is that theory is a line and experiment is points; this is because the theory is usually a prediction in the form of an equation, which can in principle be evaluated at infinitely many points, filling in all the gaps. One way to accomplish this with computer software is to graph both theory and experiment as points, but then print out the graph and draw a smooth curve through the theoretical points by hand.
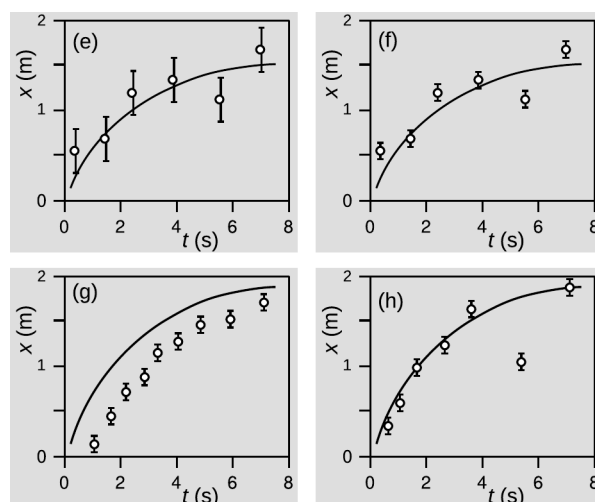
The point here is usually to compare theory and experiment, and arrive at a yes/no answer as to whether they agree. In (e), the theoretical curve goes through the error bars on four out of six of the data points. This is about what we expect statistically, since the probability of being within one standard deviation of the truth is about 2/3 for a standard bell curve. Given these data, we would conclude that theory and experiment agreed.

In graph (f), the points are exactly the same as in (e), but the conclusion is the opposite. The error bars are smaller, too small to explain the observed discrepancies between theory and experiment. The theoretical curve only goes through the error bars on two of the six points, and this is quite a bit less than

we would expect statistically.

Graph (g) also shows disagreement between theory and experiment, but now we have a clear systematic error. In (h), the fifth data point looks like a mistake. Ideally you would notice during lab that something had gone wrong, and go back and check whether you could reproduce the result.

Appendix 5: Finding Power Laws from Data

For many people, it is hard to imagine how scientists originally came up with all the equations that can now be found in textbooks. This appendix explains one method for finding equations to describe data from an experiment.

Linear and nonlinear relationships

When two variables x and y are related by an equation of the form

$$y = cx \qquad ,$$

where c is a constant (does not depend on x or y), we say that a linear relationship exists between x and y. As an example, a harp has many strings of different lengths which are all of the same thickness and made of the same material. If the mass of a string is m and its length is L, then the equation

$$m = cL$$

will hold, where c is the mass per unit length, with units of kg/m. Many quantities in the physical world are instead related in a nonlinear fashion, i.e., the relationship does not fit the above definition of linearity. For instance, the mass of a steel ball bearing is related to its diameter by an equation of the form

$$m = cd^3 \qquad ,$$

where c is the mass per unit volume, or density, of steel. Doubling the diameter does not double the mass, it increases it by a factor of eight.

Power laws

Both examples above are of the general mathematical form

$$y = cx^p \qquad ,$$

which is known as a power law. In the case of a linear relationship, $p = 1$. Consider the (made-up) experimental data shown in the table.

	h=height of rodent at the shoulder (cm)	f=food eaten per day (g)
shrew	1	3
rat	10	300
capybara	100	30,000

It's fairly easy to figure out what's going on just by staring at the numbers a little. Every time you increase the height of the animal by a factor of 10, its food consumption goes up by a factor of 100. This implies that f must be proportional to the square of h, or, displaying the proportionality constant $k = 3$ explicitly,

$$f = 3h^2 \qquad .$$

Use of logarithms

Now we have found $c = 3$ and $p = 2$ by inspection, but that would be much more difficult to do if these weren't all round numbers. A more generally applicable method to use when you suspect a power-law relationship is to take logarithms of both variables. It doesn't matter at all what base you use, as long as you use the same base for both variables. Since the data above were increasing by powers of 10, we'll use logarithms to the base 10, but personally I usually just use natural logs for this kind of thing.

	$\log_{10} h$	$\log_{10} f$
shrew	0	0.48
rat	1	2.48
capybara	2	4.48

This is a big improvement, because differences are so much simpler to work mentally with than ratios. The difference between each successive value of h is 1, while f increases by 2 units each time. The fact that the logs of the $f's$ increase twice as quickly is the same as saying that f is proportional to the square of h.

Log-log plots

Even better, the logarithms can be interpreted visually using a graph, as shown on the next page. The slope of this type of log-log graph gives the power p. Although it is also possible to extract the proportionality constant, c, from such a graph, the proportionality constant is usually much less interesting than p. For instance, we would suspect that if $p = 2$ for rodents, then it might also equal 2 for frogs or ants. Also, p would be the same regardless of what units we used to measure the variables. The constant c, however, would be different if we used different units, and would also probably be different for other types of animals.

slope=2.0

log f

log h

Appendix 6: High Voltage Safety Checklist

Name: _____

_____ Never work with high voltages by yourself.

_____ Do not leave HV wires exposed - make sure there is insulation.

_____ Turn the high-voltage supply off while working on the circuit.

_____ When the voltage is on, avoid using both hands at once to touch the apparatus. Keep one hand in your pocket while using the other to touch the apparatus. That way, it is unlikely that you will get a shock across your chest.

_____ It is possible for an electric current to cause your hand to clench involuntarily. If you observe this happening to your partner, do not try to pry their hand away, because you could become incapacitated as well — simply turn off the switch or pull the plug out of the wall.

Appendix ??: Comment Codes for Lab Writeups

A. General

a1. Don't write numbers without units. (25% off)

a2. If something is wrong, cross it out. Don't make me guess which version to grade.

a3. Your writeup is too long. The length limit is 3 pages, not including raw data.

a4. If your writeup includes printouts, staple them in sideways with a single staple.

a5. See appendix 1 for the format of lab writeups.

a6. Don't state speculation as a firm conclusion.

a7. Leave more space for me to write comments.

a8. Cut unnecessary words. Use active voice. Write in a simple, direct style.

a9. Don't write walls of text. Use paragraph breaks.

a10. Cut any sentence that doesn't carry information.

a11. This paragraph needs a topic sentence.

a12. Express this as an equation.

a13. Don't present details unless you've already made it clear why we would care. Don't write slavishly in chronological order.

a14. The first sentence of any piece of writing must make an implicit promise that the remainder will interest the reader.

B. Raw data

b1. Don't mix raw data with calculations. (25% off)

b2. Write raw data in pen, directly in the notebook.

b3. This isn't raw data. This is a summary or copy.

C. Procedure

c1. Don't repeat the lab manual.

c2. Don't write anything about your procedure unless it's something truly original that you think I would be interested in knowing about, or I wouldn't be able to understand your writeup without it.

D. Abstract – see appendix 1

d1. Your abstract is too long.

d2. Don't recap raw data in your abstract.

d3. Don't describe calculations in your abstract.

d4. The only numbers that should be in your abstract are important final results that support your conclusion or that constitute the purpose of the lab.

d5. Your abstract needs to include numerical results that support your conclusions.

d6. Give error bars in your abstract.

d7. Where is your abstract?

d8. Your abstract is for results. This isn't a result of your experiment.

d9. This isn't important enough to go in your abstract.

d10. What was the point of the lab, and why would anyone care?

d11. Don't just give results. Interpret them.

d12. We knew this before you did the lab.

d13. This lab was a quantitative test. Restating it qualitatively isn't interesting.

d14. This lab is a comparison of theory and experiment. Did they agree, or not?

d15. Your results don't support your conclusions. Write about what really happened, not what you wanted to happen.

d16. One observation can never prove a general rule.

E. Error analysis – see appendices 2 and 3

e1. A standard deviation only measures error if it comes from numbers that were supposed to be the same, e.g., repeated measurements of the same thing.

e2. In propagation of errors, don't do both high and low. See appendix 3.

e3. In propagation of errors, only change one variable at a time. See appendix 3.

e4. Don't round severely when calculating Q's. Your Q's are just measuring your rounding errors.

e5. A Q is the amount by which the output of the calculation changes, not its inputs.

e6. A Q is a change in the result, not the

result itself.

e7. Use your error bars in forming your conclusions. Otherwise what was the point of calculating them?

e8. Give a probabilistic interpretation, as in the examples at the end of appendix 2.

e9. You're interpreting this probability incorrectly. It's the probability that your results would have differed this much from the hypothesis, if the hypothesis were true.

e10. % errors are useless. Teachers have you do them if you don't know about real error analysis.

e11. If random errors are included in your propagation of errors, listing them here verbally is pointless.

e12. Don't speculate about systematic errors without investigating them. Estimate their possible size. Would they produce an effect in the right direction?

G. Graphing – see appendix 4

g1. Label the axes to show what variables are being graphed and what their units are, e.g., x (km).

g2. Your graph should be bigger.

g3. If graphing by hand, do it on graph paper.

g4. Choose an appropriate scale for your graph, so that the data are not squished down. Don't just accept the default from the software if it's wrong. See app. 4 for how to do this using Libre Office.

g5. "Dot to dot" style is wrong in a scientific graph.

g6. The independent variable (the one you control directly) goes on the x axis, and the dependent variable on the y. Or: cause on x, effect on y.

g7. On a scientific graph, use dots to show data, a line or curve for theory or a fit to the data.

g8. "Trend line" is scientifically illiterate. It's called a line of best fit.

S. Calculations and sig figs

s1. *Think* about the sizes of numbers and whether they make sense. This number doesn't make sense.

s2. Where did this number come from?

s3. This number has too many sig figs (e.g., more than the number of sig figs in the raw data).

s4. Don't round off severely for sig figs at intermediate steps. Rounding errors can accumulate.

s5. You're wasting your time by writing down many non-significant figures.

s6. Your result has too many sig figs. The error bars show that you don't have this much precision.

s7. The Calculations and Reasoning section usually just consists of the calculations you've already written.You don't need to write a separate narrative.

s8. Put your calculator in scientific notation mode.