

Conductivity Lab Overview

Ben Nachumi

December 8, 2012

1 Welcome to the solid state electrical conductivity lab!

The electrical conductivity (σ) defines the response of a system to an applied electric field. That is, the conductivity measures the tendency of electrical charges to move because of the field. This of course is the empirical property that makes metals metallic and insulators insulate, depending on the values that it takes. It would be hard to overstate its technological significance. In this lab, you'll be simultaneously learning about the conductivities of various materials, while using some of those facts to measure temperature—and enable more measurements.

In practice, one extracts information about the DC electrical conductivity from the resistivity (ρ), which is the inverse of the electrical conductivity ($\rho = 1/\sigma$). For the purposes of this lab, you can begin thinking of the resistivity as a shape-independent, bulk property belonging to a kind of material. The separation of bulk properties from geometric specifics, which is a good approximation for macroscopic samples, is what allows people to make differently rated resistors from the same kind of wire.

There are two required measurements. The first and simplest measurement is the transition temperature of a piece of superconducting wire. For the second, you will find and fit for the Debye temperature of platinum. Together, these should take up two to three meetings. Elective experiments include: actually making and testing a sample of a particular high- T_C superconductor; measuring the band gap in a semiconductor from the conductivity of a diode; and looking for resistive evidence of the ferromagnetic transition in a piece of nickel.¹

¹You might also be interested in developing a new measurement for this lab. If there is a particular material or technique you would like to try, let us know!

2 Overview of the Main Techniques

This is the rig block diagram for the low-temperature measurement. It's pretty straightforward: the sample sits in a refrigerator, with leads to measure its temperature and electrical resistance. The vacuum insulates the very cold inside of the refrigerator from its room-temperature outside. But there are some subtleties, which we shall now describe.




Figure 1: Rig Block Diagram.

2.1 The 4-Wire Technique and Alternation

A voltmeter doesn't know what it's connected to. If you were to measure the voltage across the sample using the same leads with which you applied the probing current, you would also measure the voltage drop through the current leads. Since the leads can be long and the sample resistance may be small, this can be a very significant source of measurement error. The simplest solution to this problem is to apply the current with one pair of wires, and measure the voltage with a second pair.

Another source of extraneous voltage is the *Seebeck effect* due to the thermal gradient down the voltage leads. Warm electrons diffuse faster down the wire, leading to a voltage proportional to the temperature difference. This material-dependent effect is the basis for the thermocouple thermometer, but for conductors (with small voltage signals) it's a nuisance because inevitable asymmetries in the lead wires will lead to a voltage offset. The solution to this problem is to alternate the direction of the applied current. In the average




Figure 2: Four-Wire Schematic.

of the two measured voltages, this and other asymmetries cancel out.

The Keithley current source and voltmeter can be set up to alternate in what Keithley calls “Delta mode” (more information is in the **Keithley manual**). You can initiate these measurements either from the Keithley front panels, or from the scripts we have set up for you in Matlab (see below). However for measurements of diodes and other intrinsically asymmetric devices, you would use the Keithley in DC mode.

At the start of the experiment, the Keithley should already be set up in Delta mode. For directions on how to switch over to DC mode, see section 3.6 below.

2.2 Refrigeration

To get samples down to low temperatures, you have the use of a closed-cycle helium refrigerator. Compared to the cryostats used in the other low temperature labs, this system is more like a refrigerator in the sense of a kitchen refrigerator, with a compression/expansion cycle of a refrigerant fluid. The important details of how to operate this rather complicated bit of machinery are discussed in the **CTI Refrigerator Manual**. However, if everything is working properly, the instructions in section 3.5 (**Cooling Down**) should suffice.

For now, note that the refrigerator comes in two parts, connected by pressurized helium gas hoses. The big box contains a large pre-compressor; the heavy object to which these leads go is called the *cold head*. It is so clunky because it contains a motor that drives heat-exchanging pistons back and forth. The coldest part of the cold head is at the tip of the *cold finger*. This is where the sample assembly goes. Note that all of the lead wires are thermally anchored all the way down the cold finger to prevent unmitigated thermal conduction from the feedthrough plugs to the sample.



Figure 3: The control boxes for the 4-wire measurement: Top: Keithley 6220 precision current source; Bottom: Keithley 2812A nanovoltmeter; Left: Lakeshore temperature controller.




Figure 4: The rear of the 6220/2812 pair, set up for Delta measurements. For diode measurements, the two serial connections between them are replaced by a GPIB cable.

Next, two heat shields are placed around the sample. These fairly isothermal surfaces both slow radiative heating from the tip of the cold finger and maintain the environment around the sample at a constant temperature. Finally comes the outer tail, which pressure-seals to the feedthrough housing around the shielded cold finger. Once this is done, a vane pump gets most of the atmosphere out of the sealed-off space (otherwise the remnant vapor pressure of the eventually frozen air would leak heat to the sample). At this point, with everything canned up and evacuated, you are ready to start the refrigerator.

Unburdened with a heat load, the refrigerator is meant to cool the tip of the cold finger down to around 10K over 60 – 90 minutes.




Figure 5: The cold head. The sample mounts to the copper backing at the top, and the leads are soldered to the pins.

3 Required Measurements

3.1 Calibrating Thermometers

How do you measure temperatures on an absolute scale? An ideal gas bulb and a blackbody spectrum are “model independent” standards, but these aren’t really convenient. One usually uses secondary devices whose calibrations are ultimately traceable to a standard. One common device is a resistor like a twist of platinum wire. At high temperatures, the resistance versus temperature of many metals is quite linear. But below 30 K, the resistivity of

a platinum sample becomes noticeably less linear and drops away as phonons freeze out. On the other hand, the resistivity of a semiconductor diode increases as thermally excited carriers freeze out. Hence, just where platinum thermometers become useless, the appropriate semiconductor thermometer becomes more sensitive. In principle, one calibrates the temperature/voltage function of the diode using a gas bulb thermometer, or against a previously calibrated Pt resistor, and generates an interpolation table.

Your refrigerator currently has two diode thermometers installed. These are fairly standardized and almost perfectly interchangeable. The production is well-enough controlled that the interpolation curve can be fixed to within a tolerance of about a Kelvin with just two fixed measurements.

The calibration of these thermometers tends to drift over a period of months. If you are the first group doing this experiment in a while, you should recalibrate them using the SOFT CAL procedure detailed in the Lakeshore manual. Read it. For temperature standards, you will can dunk the diodes together into ice water (273.15K) and liquid nitrogen (77.3K). After the calibration is complete, let the thermometers heat up, and make sure that they agree with the mercury thermometer on the room temperature.

3.2 Sample Mounting

For the usual samples (Pt wire, SC wire, diode) the samples should already be soldered to four leads with pins for the plugs on the copper backing. Existing samples are glued between spacers of flat thermally conductive electrical insulators: sapphire disks, or, if those run out, pieces of mica. Each spacer is varnished into contact with a diode thermometer. One of the thermometers sits atop a resistive heater (just some wire wound on a piece of mica) which in turn sits on the coldest part of the refrigerator.

3.3 Sample and Thermometer Wiring

The following tables show how everything should be connected. **The only thing you should have to do is plug the four pins from your sample leads into the appropriate holes (1,4,7, and 10) in the white plug on the copper sample holder in the cold head.**




Figure 6: How the samples are mounted on the copper sample holder. The two thermometers bracket the sample temperature. The sapphire layers provide good thermal contact and electrical insulation.




Figure 7: Pin diagram for the terminals on the copper sample holder.

Sample Connection Table		
Cold head wire color	24-pin Feedthrough pin letter	Wire/Cable/Patch/Wire/Function
Extra Copper (plug 2)	C	White/3/unconnected/uncommitted
Extra Copper (plug 5)	B	Yellow/4/unconnected/uncommitted
Extra Copper (plug 8)	A	Red/5/unconnected/uncommitted
Extra Copper (plug 11)	F	Blue/6/unconnected/uncommitted
Sample Copper (plug 1)	N	Orange-White/7/Keithley/Black copper tab/VS-
Sample Copper (plug 4)	M	Gray/8/Keithley/Black gator/IS-
Sample Copper (plug 7)	P	Purple/9/Keithley/Red copper tab/VS+
Sample Copper (plug 10)	R	Brown/10/Keithley/Red gator/IS+

Thermometry Feedthrough Connection Table		
Cold head wire color	10-pin Feedthrough pin letter	Wire/Patch/Cable/Wire/Function
Heater Copper (backing 9)	A	Red/11/Lakeshore A/White/H1
Heater Copper (backing 10)	H	Purple/12/Lakeshore A/Yellow/H2
Diode A Red (backing 1)	J	White/13/Lakeshore A/Black/IA-
Diode A Green (backing 2)	C	Green/14/Lakeshore A/Blue/VA-
Diode A Black (backing 3)	G	Orange/15/Lakeshore A/Green/VA+
Diode A Gold (backing 4)	B	Black/16/Lakeshore A/Red/IA+
Diode B Red (backing 5)	K	Brown/17/Lakeshore B/Orange/IB-
Diode B Green (backing 6)	D	Yellow/18/Lakeshore B/Red/VB-
Diode B Black (backing 7)	F	Gray/19/Lakeshore B/Black/VB+
Diode B Gold (backing 8)	E	Blue/20/Lakeshore B/Blue/IB+

3.4 Preliminary Tests

After plugging in the sample pins, you should go through the following checklist:

- TEST THE I/V CONNECTIONS! With the Keithley boxes both off, use the hand-held DVM to test the four sample connections. (Across the standard Pt sample, you should see the room temperature resistance of about 108Ω). Turn the Keithley DVM/Current Source on and put them into in DELTA mode (hit DELTA). When the current source says it's ready, hit ENTER. You should be able to read a voltage, and from the voltage and the current applied, calculate a resistance. Is the resistance you read in fair agreement with what you got when you tested your sample leads? To get out of DELTA mode hit EXIT.

- **TEST THE HEATER CONNECTIONS!** With the Lakeshore off, put the DVM across the heater patch points. You should see a resistance on the order of $10 - 200\Omega$ (the precise value will depend on the precise heater).
- **TEST THE THERMOMETER CONNECTIONS!** Turn on the Lakeshore, and make sure your thermometers are reading and are in reasonable agreement with each other. They should respond when, say, you gently touch the sample stack with your fingertip.

3.5 Cooling Down

Now you are ready to refrigerate!

- Can up the sample. Screw down the copper isothermal cap (be very careful not to pinch any wires in the process). After doing this, make sure all of the connections are still behaving sensibly. Then screw down the silvery heat shield.
- Now prepare the outer tail. Making sure the o-ring is well seated in the groove, apply an even layer of silicone grease. With a Kimwipe, gently remove any dust, hairs, or other debris from the mating flange on the feedthrough housing. Also check that there is nothing bridging the thin gap between the central cold head and the feedthrough housing by carefully running the tip of a tweezer around it, wiping errant globs of goo, etc. When everything is clean and appropriately isolated, place the tail over the cold finger so the o-ring sits on the flange.
- **TURN ON THE AIR CONDITIONER IN THE BIG WHITE SOUND-PROOFING BOX.** If you don't do this, the box will get too hot to run the compressor.
- Plug in the vacuum gauge, and plug in the roughing pump. The pump should gurgle, and the pressure in the vacuum jacket should start to drop almost immediately. Let it get down to a steady state.
- Go to the circuit breaker box down the hall, and flip the breaker switch for the compressor. You should hear a noise from the soundproofing box like a refrigerator starting up, and the cold head piston should start to pop and chug. After about 30 seconds, the thermometer readings should start to drop. If you are cooling fast, there will be a significant lag between the thermometers. This is of course because it takes time for the stack of things to come to equilibrium.
- Wait for the refrigerator to get to a stable base temperature. This will take somewhere between one and two hours. Spend this time profitably by familiarizing yourself with the MatLab interface.

3.6 Getting Ready to Take Data

The MatLab file that controls the basic automated measurement is called `supercond_ul.m`. Open it, and *read the comment header*—it will tell you what you need to know about the input parameters, what data goes where, and how to reset the interface between MatLab and the measurement electronics. `supercond_ul.m` runs in two modes: “alone” and “follow”. In “alone” mode, you specify a certain number of points to take, whereas in “follow” mode, you take data until the thermometers hit a certain temperature.

Once you are at base temperature (about 10K), take a few points in “alone” mode: in the MatLab command window, enter `point=supercond_ul(a,5,.1,1,10,"test")`, and follow the prompts. If everything is working as it should, you will see and hear the Keithley boxes take five DELTA points, and these should appear in a graph box on your screen. The data will have been stored in the file `... \data\test`.

To get a sense of your signal-to-noise, take a couple of DELTA measurements using different excitation currents at the same temperature. Because higher currents will give you more significant digits, you will see less scatter in your data as the current increases. On the other hand, you don’t want to use a current that is too large, for this will heat up your sample. For the very small bit of superconducting wire, you can probably get away with .01A. For the Pt sample, which has a room temperature resistance of about 100 Ω , .001A is a safe bet.

Now turn your attention to the Lakeshore box. For manual control, it needs to be in LOCAL mode. You can toggle between LOCAL and REMOTE by hitting the LOCAL/REMOTE button on the front panel: the “R” in the upper right hand corner of the LCD display indicates when the box is in REMOTE mode, i.e. under compute control. To slowly change temperatures when the cold head is near its base temperature, the Lakeshore PID parameters ought to be small—less than or equal to 1.0. In LOCAL mode, you can check the PID settings by pressing the P,I, or D buttons on the front panel.

Now control the sample temperature. In LOCAL mode, press SET POINT, and enter a set point not too far from base temperature—say, 15K. Then, and only then, use the HEATER RANGE button to turn the heater on LOW.

Use `supercond_ul.m` to take some more points while the temperature is changing. Incidentally, if you are unfamiliar with plotting data in MatLab, this is a good time to explore those help menus and/or ask your TA for assistance.

3.7 Taking Data

If you are measuring the T_C of the superconducting wire, your goal is to find the transition temperature as precisely as possible. This means controlling the temperature and moving in small increments. That's about it! Your sample temperature will be bracketed by the temperature on each thermometer, so the average of the two readings is probably the best measure of the sample temperature.

When you measure the Pt sample, you need to scan from base temperature to room temperature. So after you have measured a few times around the base temperature, you should invoke `supercond_ul.m` in “follow” mode with something like `point=supercond_ul(f,290,.1,1,10,"filename")`. Then shut off the compressor; the script will keep taking and plotting measurements until one of the thermometers reaches 290K.

Once you have taken your Pt data, you can fit for the Debye temperature using the predefined fitting routine `BGfit.m` (see below).

3.8 File Management

To help keep a sensible record of the data taken and the students who took it, please set up a personal directory: `...\student's documents\[your_name]`. When you take data with MatLab, you will specify your data file names with that directory. You will also have to add your directory to MatLab's search path. The MatLab scripts that you will be using live in `...\supercond\m.files`. Before using them, copy them into your own directory. Then you can safely modify them if you must.

3.9 Provisions for Fitting Data

All of your data is read into MatLab, which also provides a powerful fitting environment. To fit your data, there are two files that you need. One is a master file that reads in your data, defines the starting parameters for your fit model, and calls a MatLab comparison function, with the name of your second file as an input. The second file defines your fit model—it generates a theoretical function of the temperature. The MatLab comparison function then takes the (R,T) data and compares it in a least-squares sense to your model function, altering the parameters that you fed it to get the best match. Then it returns those parameters to the master file.

The pre-made example is the pair of files `BGfit.m` and `BGcomp.m`. This pair is what you will use to extract the Debye temperature in Pt from your R versus T graph. `BGfit.m` is the master file, which reads the data and invokes the comparison to the model defined

in `BGcomp.m`, then plots the original data, along with the fitted function. `BGfit.m` with `BGcomp.m` should serve adequately as a model for you to follow in defining your own fits, which you will need to do if you measure the gap Δ in a semiconductor.

4 Some Notes on the Elective Measurements

4.1 Non-Alternating Measurements (for Diodes)

For measurements of diodes, the alternating 4-wire measurement doesn't make sense. Fortunately for diodes, the voltages you want to measure will be much larger than the voltage difference due to the asymmetry of the lead wires, so you can get away with a non-alternating 4-wire measurement. Use the script `diode_gpiib.m`. For it to work, you must remove both the regular serial connection and the "fast serial" communication cables between the 6220 and the 2182, reconnect them with a GPIB cable, and set the communication modes of the 6220 and 2182 to GPIB.




Figure 8: The Keithley boxes connected for Diode measurements.

Make sure you know which direction on the diode is “forward.” To avoid damage to the diode and the concomitant bad data set, try not to apply too great a reverse current until and unless you want to see “reverse bias breakdown”. Unless you are trying to break the diode, keep reverse currents to less than $1\mu\text{A}$. For forward current, you probably don’t need to go much higher than a few mA.

The idea is to sit at a given temperature and trace out a representative current vs. voltage curve (of course, you will be setting the current and measuring the voltage). This requires time for the temperature to settle, while your Lakeshore controller applies heat against the refrigerator’s cooling power. So plan on just a few measurements like (100K, 150K, 200K, 250K, 300K). For these elevated temperatures, you will have to change the PID settings—try (10,10,10) as a starting configuration.

4.2 High Temperature Measurements and Furnace Use (for the T_C of Ni)

For high temperature measurements, the sample sits on the end of a high-temperature insert made of alumina, a few millimeters away from a thermometer—in this case is a platinum resistor identical to the one you use as a sample in the Debye temperature measurement. You can safely extrapolate your resistance versus temperature measurement to high temperatures to estimate the temperature of the sample in the oven.

To make the measurements, you will have to patch the thermometer and sample inputs of the Lakeshore and the Keithley using terminals 1& 2 (for the Pt thermometer) and 21-24 (for the sample measurement)—see the table below. Because there is only one thermometer for this measurement, you will be using a different version of the automation script, called `nickel.m`. This reads the resistance of the nickel sample and the resistance of the platinum resistor.

Oven Insert Connection Table		
Feedout Wire	Patch/Cable/Wire(s)	Function
Green (black tape)	1/Lakeshore A/ Black	IA–
	1/Lakeshore A/Blue	VA–
Green	2/Lakeshore A/Green	IA+
	2/Lakeshore A/Red	VA+
Red (white tape)	21/Keithley/Black copper tab	VS–
Red (black tape)	22/Keithley/Black gator	IS–
Red (green tape)	23/Keithley/Red copper tab	VS+
Red	24/Keithley/Red gator	IS+

The “Blue-M” box furnace is programmable—the necessary details are laid out in the marked pages of the **Box Furnace Manual**. The idea is to start from your high set temperature and to ramp down to a few degrees above ambient (say, 30°C) very slowly—the ramp rate should not exceed 0.5°C/minute. Make the set the temperature comfortably above T_C , say 450°C, and wait for the Pt thermometer to equilibrate. Note that the temperature of the Pt thermometer will not necessarily be the same as the oven’s self-read temperature. Let the oven ramp to whatever temperature it thinks it is going to, but trust the Pt thermometer!

5 Making BSCCO

There are facilities for making fresh samples of the high temperature cuprate superconductor, $\text{Bi}_{2+x}\text{Sr}_{3-y}\text{Ca}_y\text{Cu}_2\text{O}_{8+\delta}$ (“BSCCO”) which (under the growth conditions you will use) has a T_C of around 90K.

You can bake up BSCCO as a powder, using the precursor materials BiO_2 (bright yellow), Sr_2O_3 and Ca_2O_3 (both white), and CuO (black). There is a scale, a mortar and pestle, crucibles, an oven, and a press for forming the BSCCO powder into pellets. At the appropriate time in the firing process, you will need liquid nitrogen for quenching.

The BSCCO process will take at least two weeks to complete. For most of this time, the mixed powders are in the kiln, slowly reacting. However you’ll need to schedule when you do the intermediate grinding and pressing. The grinding and pressing will occur in Pupin 1313, where the fume hood and pellet press are located.

5.1 The Process

As you might imagine, the combination of these precursors leads to a very complicated phase diagram. All sorts of intervening materials that can form at elevated temperatures. The phase diagram simplifies if all of the materials are forced to remain solid throughout the reaction. For this solid state reaction process to work, appropriate quantities of the precursor powders are mixed intimately together so that the appropriate ions can diffuse together.

So your first task is to figure out how much of each powder to mix. This of course depends on the precise composition of your material. Read through the provided references and look for the composition that yields the optimal T_{SC} . Then figure out the mass of each precursor you’ll need to make about 10 g of sample. Carefully weigh out the precursors. Use weighing paper and weigh out each species individually. **Wear gloves whenever handling chemicals.**

Once you've weighed out your precursor powders, mix them thoroughly in the mortar. The mixing powders quickly look greenish grey. When thoroughly mixed, the powder should look homogeneous under magnification, without any white particles of uncrushed carbonate. Be prepared to mix for a while.

Once the powder is thoroughly mixed, fill it into an alumina crucible for *calcination*. In this step, the carbonates decompose and the precursors form tiny particles of the layered cuprate. You can find the firing schedule in the Tsukamoto article, but you should also look at the Majewski article to see where you are in the phase diagram. To make phase-pure BSCCO by this method, it's important that you never let the powder melt. To learn how to control the furnace, consult the marked pages in the **Box Furnace Manual**. The resulting powder should look black.

The next step is the annealing step. You'll use the press to press the calcined powder into a pellet. After carefully extracting the pellet, you'll anneal it (following the annealing schedule in Tsukamoto, careful not to bring it over the melting curve). The annealing allows larger particles of the cuprate to grow. Because the moving particles are relatively large, this takes a very long time, on the order of a week; and repeating it after grinding up the annealed disk improves the homogeneity of the final material. To ensure that all of the dopant oxygen stays around, you should quench your samples between firings in liquid nitrogen.

It is more obvious than obvious that you should always wear gloves and a dust mask when working with fine chemical powders. The powders can slosh around in the mortar, so to avoid unnecessary spills, please grind them slowly. It's also obvious that the furnace is hot, and liquid nitrogen is cold. Use safety gloves and tongs to handle both the hot and the cold crucibles. Nitrogen is an asphyxiant: it fills the room and feels like air, but it can displace oxygen. So when you use it, make sure to leave the door open for ventilation.

6 References

6.1 File and Manual Locations

6.2 Manuals

Everything is in `C:\Documents and Settings\Student\My Documents\Conductivity Experiment:`

- This manual: `...\Manual\Conductivity Manual.pdf`
- Tex sources for this manual: `...\Manual\Sources\` directory.
- Matlab scripts: `...\Matlab functions\` directory.

- Keithley manuals: 2182A.pdf and 622x.pdf, in the ... \Hardware manuals\Keithley\ directory.
- CTI (cold head) manual: ... \Hardware manuals\cryostat\CTI
- Oxford Instruments compressor manual: ... \Hardware manuals\cryostat\compressor
- Lakeshore temperature controller manual: paper, somewhere in the room. This one describes the SoftCal procedure.
- Box Furnace manual: ... \Hardware manuals\furnace\TPS and blue binder on the furnace cart.

6.3 Reference Experiments

- BSCCO Recipes: ... \BSCCO\Tsukamoto.pdf, ... \BSCCO\Majewski.pdf
- Reference Experiment for Platinum: ... \Reference Experiments\Platinum.pdf
- Reference Experiment for Diode/band gap: ... \Reference Experiments\Diode.pdf
- Reference Experiment for Nickel: ... \Reference Experiments\Nickel.pdf

<hr/> GPIB mapping <hr/>		
Device	Bus	Address
Lakeshore Temperature Controller	gpib0	12
Keithley 6220 Current Source	gpib0	13
Keithley 2182A Nanovoltmeter	gpib0	14

7 Theoretical Background

7.1 The Drude Model

The standard classical mechanical model for the conductivity/resistivity of solid materials is called the Drude model, and it provides a useful frame of reference for all the systems you'll be looking at. In the DC version of this model, which should be familiar from your introductory physics courses, the electron is treated like a little charged ball, bouncing along through the material under the influence of an accelerating electric field. After every bounce, the velocity is reset to some random value consistent with the temperature of the lattice. Unless there is something to stop the electrons—some sort of boundary or tether—they will tend to rattle on their way forever, with a bias in the direction of the electric field. So in the absence of an explicit binding force, this DC model describes a metallic conductor. The presence of a binding force implies electrical insulating behavior.




Figure 9: The Drude picture of a metal. Electrons are tiny, charged pinballs rattling past obstacles in the general direction of the applied field.

For AC voltages, the distinction between conductors and insulators is muddled, because AC currents in metals go back and forth anyway. Nevertheless, one expects a qualitative difference based on the notion that bound, oscillating electrons should have a resonance frequency (corresponding to the “spring constant” of their tether). So one can generally define a metal in the Drude model as a material for which at least one “resonance” occurs at zero frequency: at least one “spring constant” is set to zero.

It’s pretty simple to derive the full AC-DC Drude model for conductivity. To start, assume that a single charge carrier (mass m and charge q) is accelerated under the influence of an electric field with a form $\vec{E}(t) = \vec{E}_0 \cos \omega t$, while it is also subject to a restoring force $-k\vec{x}$ (\vec{x} is the displacement from equilibrium) and a dissipative force $f\vec{v}$ (\vec{v} is the velocity).

Solve the Newtonian $\vec{F} = m\vec{a}$ equation to find the amplitude and maximal speed of the steady-state solution. Next note that the classical expression for the electrical current density \vec{j} is just $\vec{j} = nq\vec{v}$, where n is the number density of the charges. One slick way to solve the $\vec{F} = m\vec{a}$ equation is to rewrite the electric field as $\vec{E}(t) = \vec{E}_0 e^{i\omega t}$ and take the real part at the end of the calculation. This works because the equation of motion is linear. The answer is

$$\vec{j}(t) = \frac{nq^2 \times i\omega}{k - m\omega^2 + if\omega} \times \vec{E}_0 e^{i\omega t}.$$

The proportionality between \vec{j} and \vec{E} is the frequency-dependent electrical conductivity $\sigma(\omega)$. Its inverse is of course the electrical resistivity, ρ .

If you rationalize the denominator and take the real part of your answer, you will find that it falls into two pieces: one part of the current is in phase with the electric field (goes like $\cos(\omega t)$) and the other is $\pi/2$ out of phase (goes like $\sin(\omega t)$):

$$\text{Re}(\vec{j}) = \frac{nq^2\omega^2 f}{(k - m\omega^2)^2 + f^2\omega^2} \times \vec{E}_0 \cos(\omega t) - \frac{nq^2\omega(k - m\omega^2)}{(k - m\omega^2)^2 + f^2\omega^2} \times \vec{E}_0 \sin(\omega t).$$

What does this mean?

$\vec{j} \cdot \vec{E}$ has the units of a power density. Over a cycle, the out-of-phase product $\sin \cos$ averages to zero, but the in-phase \cos^2 averages to $\frac{1}{2}$. This nonzero part is the cycle-averaged power transmitted to the current. Where does it go? Its proportionality to f is an important hint. Recall that f represents a damping force on a single oscillating charge. What we have described here is a steady state: the kinetic energy of the current is dissipated by the frictional force. By contrast, the out-of-phase part determines the energy flow into and out of the electromagnetic fields with no net loss.

Because we initially assumed $k \neq 0$, we have described a DC ($\omega = 0$) insulator. For $\omega \rightarrow 0$, $\sigma \rightarrow 0$ and $\rho \rightarrow \infty$. However the same analysis goes through with $k = 0$, so the expression for the conductivity for conductors is just the $k = 0$ limit of the general expression. This leads to another familiar interpretation of the dissipation coefficient f in a conductor. Given that the additional velocity of a charge accelerated in a DC field \vec{E} for time τ should be $q\vec{E}\tau/m$, the current density should be

$$\begin{aligned} \vec{j}_{DC} &= \frac{nq^2 \vec{E} \tau}{m} \\ (\sigma_{DC} &= \frac{nq^2 \tau}{m}). \end{aligned}$$

Comparing this to the AC/DC formula we just derived with $\omega = 0$, you'll find that $f = m/\tau$. In other words, m/f is the average time for the uniform acceleration of an electron by the external field, before the electron hits some obstacle and decays out of the current.




Figure 10: The in-phase part of $\sigma(\omega)$ for a metal ($k = 0$). As $f \rightarrow 0$ ($\tau \rightarrow \infty$), the energy absorption per cycle peaks more and more strongly about $\omega = 0$.

7.2 “Perfect DC Conductors”

Things get a bit subtle when we try to imagine what would happen for a perfect DC conductor ($k = 0$ and $f = 0$ simultaneously). Imagine a sequence of graphs of the in-phase part of the conductivity as $\tau \rightarrow \infty$ and $f \rightarrow 0$. The typical Lorentzian narrows, and the peak at $\omega = 0$ simultaneously increases, in such a way that the total area under the curve remains constant. If we rescale the ordinate, we get closer and closer to what seems a perfectly thin spike at $\omega = 0$. This $f \rightarrow 0$ limit is a Dirac δ -function. For a perfect conductor, with no dissipation, the in-phase conductivity happens at zero frequency. Because τ represents the amount of time for which an electron retains the energy that the field gives it, and (for constant m) $f \rightarrow 0$ makes $\tau \rightarrow \infty$, a perfect conductor would be a great place to store kinetic energy, transmit energy without loss, etc. If only life were that easy!

8 What affects the Resistivity

Here is where we start to talk about the effects that you can see with this apparatus. The Drude parameters n and τ turn out to have remarkable material- and temperature-dependences that tell a great deal about the physics of conductivity.

8.1 Impurity Scattering in Metals—a Surprising Temperature-Independence

The idea of perfect metallic conductivity is no relic of the Drude model. The quantum mechanical treatment of the electron states in a lattice also implies that a single electron need not scatter out of its momentum state. So why do metals exhibit finite τ s? There are many mechanisms to consider. The simplest to imagine is impurity scattering. An electric field gives a directed momentum to a conduction electron. Classically, hitting something knocks it off-course. The quantum mechanical process is much the same. So we can get fairly far (though we can't explain everything) by imagining a classical picture.

To be specific, consider the *classical ballistic* model for impurity scattering. Call the density of impurities N_i . Define the *mean free path*, Λ , as the length for which an electron on average hits one impurity. Actually, it's easier to imagine the impurities sweeping toward the charge carrier. Supposing that the radius of one impurity is r_i , then $\Lambda \times \pi r_i^2 N_i \sim 1$. But $\Lambda = v\tau$, where v now stands for the average speed. So $\tau \sim \frac{1}{vN_i}$. This is true: the dirtier the material, the higher the resistance. All else being constant, $N_i \rightarrow 0$ implies $\tau \rightarrow \infty$.

For a fixed N_i , we might expect to be able to understand temperature dependence from the variation of v . If, as you might expect for a classical gas, v were something like $\sqrt{k_B T/m}$ (with $k_B T \gg E\tau$) then you might expect the impurity resistance to increase with temperature. Why it doesn't is due to the peculiarities of quantum mechanics.

Quantum mechanically, the electrons in a crystal are spread throughout the whole volume. Imagine filling this space up to the appropriate electron density n . As you add electrons, the quantum mechanical phenomenon of *Pauli exclusion* limits each state to two electrons. Suppose you put the first two electrons in the state of lowest kinetic energy, that is, zero momentum. Newer electrons are forced into higher and higher momentum states. For a completely isotropic material, you'd end up with a *Fermi ball* of momentum states, filled out to the *Fermi momentum* p_F .

On the other hand, you have a certain number N of electrons in a given crystal, and however you count, the number of occupied states has to be the same. Let's count the momentum states and so relate N to p_F . Using the uncertainty relation $\delta x \delta p_x \approx h$, with $\delta x = L_x$, the x dimension of the crystal, we get $\delta p_x \approx h/L_x$. Repeating for y and z , we




Figure 11: Fermi Ball. The idea is that each momentum owns a particular volume of momentum space (specifically h^3/V , where V is the volume of the sample). So the electrons, of which only two can ever have the same momentum, must stack up.

find that in momentum space, each state must take up a box of size $\approx h^3/V$. Filling states, two electrons per momentum, up to p_F just gives twice the volume of the Fermi ball in units of h^3/V . Therefore

$$N = 2V \times \frac{\frac{4}{3}\pi p_F^3}{h^3}.$$

For macroscopic pieces of metal with N_a atoms each with Z electrons, the number of states per unit volume per increment of momentum is $2 \times h^3/V$, but the total number of electrons $N = ZN_a \propto V$. So you'll always fill, say, a sample of nickel up to the same top momentum. p_F doesn't change with the amount and can be treated as a property of the material.

In the absence of an electric field, the state of minimum energy requires that the Fermi ball be centered around $\vec{p} = 0$. With no electric field, if \vec{p} is occupied, then so is $-\vec{p}$: everything is symmetric, and the total current = 0. But notice that on the electrons in the ground state have an average kinetic energy of

$$\langle KE \rangle_0 = \frac{4\pi \int_0^{p_F} \frac{p^2}{2m} p^2 dp}{\frac{4}{3}\pi p_F^3} = \frac{3}{5} \frac{p_F^2}{2m}.$$

Now suppose an electric force $e\vec{E}$ points to the right. The top states can move over to the right, then the states below those, and so on... until the entire ball has been shifted, piece by piece, a bit to the right. There is now an average rightward momentum, the size of which is limited by the effect of scattering: every so often, some of the rightward electrons are taken out of their right-going momentum state, and added onto the left side of the




Figure 12: Current-carrying state of the Fermi ball. The Fermi ball is displaced by the average momentum shift $e\vec{E}\tau$. One electron has been scattered backwards, and another is going to hop into its place. The net current density comes from the momentum states in the crescent on the right.

Fermi ball. Some other electron gets accelerated into the newly vacant state.

In the equilibrium of the current carrying metal, the rightward acceleration and the scattering balance. The Fermi ball is built up on the right side and diminished on the left side, but only to a point, so that it is approximately the same shape, just displaced by the momentum $e\vec{E}\tau$. For reasonable current densities, this is much smaller than p_F , so the largest momentum of any single electron is *still* on the order of \vec{p}_F .

Hence the proper speed to use at all temperatures for which the metal is solid is the Fermi speed, $v_F = p_F/m$, which for classical particles would correspond to a temperature on the order of 10^4 to 10^5 K. This is 10 – 100 times bigger than the melting point of any metal. *The fact that in metals the impurity resistivity is largely independent of temperature is a consequence of the fact that electrons do not obey classical statistics.*

There is another thing to notice here. You get a macroscopic current when the Fermi ball is pushed one way, and the cancellation between \vec{p} and $-\vec{p}$ isn't perfect. Now look at the current density \vec{j} the following way, as $e/m \times$ (the net momentum density). How much net momentum density is there in the situation where the Fermi ball is displaced by $e\vec{E}\tau$? Let A_F stand for the surface area of the Fermi ball. If $e\vec{E}\tau \ll p_F$, then the net momentum density due of all of the occupied states will be the momentum density of the states in the uncanceled crescent, about

$$\frac{\langle \vec{p} \rangle}{V} = 2 \frac{p_F}{3h^3} \times A_F e \vec{E} \tau,$$

(2 for the spins and the factor of 1/3 because we are displacing in the x direction only) and

$$\vec{j} = 2 \frac{p_F}{3h^3} \times A_F e^2 \vec{E} \tau / m.$$

This means that in the quantum mechanical theory of metals, $n = 2p_F/3h^3 \times A_F = \frac{8}{3}\pi p_F^3/h^3$. But because we filled states up to p_F , this is just N/V . So the Fermi ball ideas return the simple Drude n .

8.2 Phonon Scattering in Metals: Varying τ with Temperature

Faced with the multiplicity of scattering mechanisms, you might not be surprised that the general computation of the electrical resistivity of a metal can turn into a very complicated undertaking. The total scattering rate should be something like the sum of all the different rates taken individually. That is, $\frac{1}{\tau_{total}} = \frac{1}{\tau_{impurity}} + \frac{1}{\tau_{other}}$.² But there is some hope of

²In other words, all else being constant, the various resistivities of the material should add like resistors in series.

disentangling the major contributions because each mechanism has a characteristic temperature dependence.

In the *Bloch-Grüneisen* model of metallic conductivity, the two major process is scattering from lattice distortions. In the harmonic theory of crystal structure, such distortions are made up of *phonons*—vibrational modes of the lattice. A lucky confluence of phonons will cause local variations in the lattice density that act like scattering centers. Bloch and Grüneisen focused on two features of phonons:

1. Phonons carry momentum. Electrons “scatter” from phonons by absorbing or emitting phonons. So a given phonon scatters a given current-carrying momentum state to a definite destination. So different phonons will be more or less effective at reducing the current, depending on their momenta.
2. The number density of phonons of a given energy (and hence momentum) changes with temperature. So unlike impurity scattering, phonon scattering will be temperature dependent.

The first feature implies a geometry problem. Because phonons energies are relatively low compared to $p_F^2/2m$, we only have to consider scattering of electrons near the surface of the displaced Fermi ball (the *Fermi surface*)—we haven’t got the necessary energy to promote an electron from deep inside. The elastic scattering of an electron at the Fermi Surface from a phonon of momentum \vec{q} leads to a loss of momentum (as measured from the current direction)

$$\Delta p = p_F(1 - \cos \theta)$$

and this loss can be averaged over the Fermi surface to get the net effect for that particular \vec{q} :

$$\langle \Delta p \rangle_{\text{Surface}} = 2\pi p_F \int_0^{\theta_{\text{Max}}} n(\theta)(1 - \cos \theta) \sin \theta d\theta,$$

where $n(\theta)$ is the number density of phonons available for scattering electrons through the angle θ . θ_{Max} corresponds to the largest possible scattering angle—obviously belonging to the largest possible \vec{q} .

The second feature brings in statistical mechanics. The scattering is geometrically limited to $\sin(\theta/2) < q/2p_F$. So the question is: what’s the typical q at a temperature T ? This depends on both the number of phonons of a given energy, and the dependence of q on the energy of the phonon. Since phonons are Bose particles, their occupation number as a function of temperature is

$$n_q \sim \frac{1}{e^{E_q/k_B T} - 1} \sim k_B T / E_q$$




Figure 13: Cartoon of a particular phonon on a centered square lattice. The blue atoms are displaced from their equilibrium lattice positions.




Figure 14: Electrons scattering from phonons (point particle picture). The lattice ions in blue are displaced from their equilibrium positions. Electrons (tiny green dots) are drawn to regions of temporarily higher positive charge density.




Figure 15: Phonon scattering cartoon (momentum space version). Here q is the phonon momentum and the initial and final electron states are taken to be on the Fermi surface.

with the last expression as a very crude approximation for phonons of energy $E(q)$ and momentum q .

For the relation between q and E_q , there is the *Debye model*. In the Debye model, the energy of a phonon of momentum q is just cq , where c is the speed of sound in the crystal, and (as with light) q is h/λ . The sound modes have a natural largest momentum q_D (“D” for Debye) corresponding to the inverse dimensions of a cell of the crystal, with a largest energy $E_D = cq_D$. By dividing E_D/k_B , we convert E_D to the *Debye temperature*, Θ_D . So to find the typical $\langle q \rangle_T$ at typical energy $k_B T$, we may estimate its momentum as a fraction of q_D by simple proportionality: $q/q_D = E/E_D \rightarrow \langle q \rangle_T = q_D(T/\Theta_D)$.

As with electron states, each distinct phonon momentum takes up h^3/V of momentum space. There are N_a atoms in the crystal, and 3 ways for each atom to move, which leads to $3N_a$ different possible phonon states (as compared to ZN_a electron states). Hence we can build a *Debye ball* in much the same way we built up a Fermi ball. The difference is that each phonon momentum state can be occupied by many more than two phonons, because phonons are Bose particles. But because the radius of the Fermi ball ($\sim (ZN_a)^{1/3}$) will be comparable to the radius of the Debye ball ($\sim (3N_a)^{1/3}$), q_D and p_F are around the same size.

For temperatures $T \ll \Theta_D$, the shorter wavelength, higher energy phonons “freeze out”: they become much rarer than their low-energy, longer-wavelength, and hence smaller q cousins. Using the fact that the *maximum* phonon momentum, q_D , is about the same as p_F we take the *typical* phonon $\langle q \rangle$ to be much smaller than p_F . So at low temperatures, we can (again, very crudely) approximate $q/p_F \sim \theta$. Retaining the crude approximation that the number of these surviving low-energy phonons is proportional to T , you can calculate

that:

$$\langle \Delta p \rangle_{\text{Surface}} \propto 2\pi p_F T \int_0^{T/\Theta_D \times q_D/p_F} (1 - \cos \theta) \sin \theta d\theta \propto T^5.$$

The connection to a Drude parameter is that the rate of momentum loss from the current is proportional to $1/\tau$. So the above expression says that at low temperatures, as the higher energy phonons freeze out, the resistance vanishes very quickly. At high temperatures, the upper limit of the integral saturates at q_D/p_F , so the temperature dependence for $T > \Theta_D$ will be linear.

A more careful calculation along the same lines yields the *Bloch-Grüneisen formula*

$$\frac{1}{\tau} = v_F \sigma_a \frac{\hbar^2 q_D^2 k_B}{M_a k_B^2 \Theta_D} \left(\frac{T}{\Theta_D} \right)^5 \int_0^{\Theta_D/T} \frac{4z^5 dz}{(e^z - 1)(1 - e^{-z})}.$$

Ignore for now the jumble of constants³, and realize that what's left is a *universal function* of T/Θ_D . So this theory says that, for any metal whatsoever, if you plot temperature in units of Θ_D , and resistance in units of the resistance at Θ_D , then the R-T graphs of all metals will look exactly the same. That's a big prediction! Though not perfectly correct, this simple Bloch-Grüneisen model works remarkably well for a variety of metals, and you will clearly see the effects of the phonons freezing out in the R-T curve for your platinum sample.

8.3 Band Gaps: Varying n with Temperature

In the Drude picture of an insulator, to generate a DC conductivity, you have to break the springs that tether some number of electrons to their nuclei. To free a charge requires an activation energy—but once the charge is free, it can wander around under the influence of an electric field, almost as if it were in a metal. As the temperature increases, one expects a population of thermally excited carriers. Recall that the Drude-model conductivity $\sigma = \frac{nq^2\tau}{m}$. We saw above that the number of electrons at energy ϵ follows the Fermi distribution—they stack up into higher energy states when they are made to encounter each other often. But because insulators are, well, insulating at room temperature, we would expect that the number of free electrons should be small, making electron-electron collisions unlikely. So in an insulator the freed electrons approximate a Maxwell-Boltzmann gas. For the proportion of electrons available to conduct at temperature T , with a tethering energy G , we'd calculate

$$n \propto \int_0^\infty e^{-(p^2/2m+G)/k_B T} \times \frac{4\pi p^2 dp}{h^3} \propto T^{3/2} e^{-G/k_B T}.$$

³ σ_a is the cross section for scattering from a single atom, and will be an important thing to know about when you try to understand what might change at the Curie point in nickel. To understand where the other constants come from, consult the references and/or your lab instructor.

A *semiconductor* is just an insulator with a smallish G , so that at room temperature it will conduct electricity, but at low temperature, it insulates. The difference between crystalline insulators and semiconductors is one of degree, rather than kind.

Of course quantum mechanics modifies the Drude picture. To see how, consider the case of a strongly bonded insulator like diamond. In the ideal diamond lattice, each carbon atom shares its outer four $n = 2$ electrons with four other carbon atoms, in four covalent bonds. There is no saying which carbon atom any of these bonding electrons belong to—in fact, these electrons must live in states that extend throughout the crystal.

As you saw for the case of a metal with extended states, the net electric current is really the effect of a disproportionation in the occupancy of momentum states—e.g., more rightward than leftward states, and the displacement of the Fermi ball. So in an insulator, something must happen to stop the disproportionation. The cause is called a *band gap*.

The band gap is a purely quantum mechanical effect. First, imagine graphing the energy as a function of momentum for an electron. For a free electron, you'd expect the smooth parabola $E = p^2/2m$ —and this is what I used in describing a metal. But this can only have been a great approximation. As you know from optics, a wave moving over a lattice generates reflections which can interfere with the original wave. And you have probably heard of the DeBroglie relationship between wavelength and momentum, $\lambda = h/p$. Now when, in the case of electrons moving through a lattice of ions, λ is about twice the separation between adjacent atoms, the reflection leads to something like a standing wave on a string. The probability density of finding a charge piles up in one of two patterns—either in the region between the nuclei, or right atop the ions. Because of the Coulomb interaction, the energy of the electron state in each case is different. But since the momentum for each state is the same, there will be a discontinuous jump in the allowed energy. There is a band of allowed energies below the critical momentum, and another band of allowed energies above the critical momentum. The band gap is just this forbidden range of energies.

Now imagine filling up the states with electrons, just as in a metal. With the right number of electrons, you can fill states right up to the edge of the gap. In a partially empty band, there are empty states of nearly the same energy allowing the electrons in the top levels to gain small increments of momentum. If the band is filled, there are no nearby states—an electron at the top level of a filled energy band must gain a finite amount of energy to change its momentum. So a crystal with filled bands is an insulator. The filled, lower bands are called *valence bands*, while the upper, empty bands are called *conduction bands*.




Figure 16: A schematic of the energy bands in a crystalline solid.

8.4 Conduction Electrons, Valence Holes, Δ , and Doping

When an electron does acquire enough energy to jump the gap and provide current, it leaves behind a vacant state called a *hole*. This valence hole provides another place for an electron in the valence band to go—but when it does, it leaves behind its own vacancy. In effect, when the valence electron filled the valence hole, the hole moves in the opposite direction. So when an electron jumps from the valence band to the conduction band, *two* charge carriers are made: an *electron-hole pair*. Of course, a conduction electron can also fill a valence hole in the reverse process.

In equilibrium, the rate of collisions between conduction electrons and valence holes should match the rate of spontaneous pair creation:

$$n_e n_h \propto e^{-\Delta/k_B T}.$$

In a pure insulator, the only way to get conduction electrons is pair creation, so the densities of electrons and holes are equal: $n_e = n_h$. Hence for a pure semiconductor

$$n_e \propto e^{-\Delta/2k_B T}.$$

Unfortunately very pure semiconductors are quite hard to make. Crystals are easily contaminated with impurities of different atoms or contain lattice imperfections that can either contribute an electron to the conduction band or create a valence hole, and this makes a

direct measurement of the band gap Δ difficult. But there is a workaround, using semiconductors intentionally *doped* with a comparatively large concentration of impurities. To understand how this works, you first need a sketch of how doping systematically changes the electron and hole densities in semiconductors.

A crystal of Si (valence 4) doped with P (valence 5) will have a surplus of conduction electrons. Such *n-type* dopants contribute filled electron states near the top of the energy gap, which are easily excited into the conduction band. But the holes they leave behind are localized near the dopant and so do not conduct. An *n-type semiconductor* has many more conduction electrons than valence holes. The majority carriers in an n-type semiconductor are electrons.

On the other hand Al (valence 3) dopants in Si contribute holes (the lack of an electron) near the bottom of the gap, just above the top of the valence band. Electrons from the top of the valence band easily jump into these vacant states, where they become localized around the impurity, and leave behind a mobile valence hole. Thus *p-type semiconductors* have more valence holes than conduction electrons. The majority carriers in a p-type semiconductor are holes.

At most temperatures, the number of carriers contributed by dopants far exceeds the number of intrinsically excited carriers. Hence in n-type Si, for example,

$$n_e = n_d + n_h,$$

but the same rate equation still applies, so that

$$n_e n_h = (n_d + n_h) n_h \propto e^{-\Delta/k_B T} \rightarrow n_h \propto e^{-\Delta/k_B T}$$

when $n_d \gg n_h$. So to measure Δ , you can measure the minority carrier density (n_h in an n-type material or n_e in a p-type material). This is much easier and much less expensive than growing ultrapure crystals.

8.5 The Diode

You will be measuring the minority carrier current I_0 in a device called a *junction diode*. This consists of a single piece of doped semiconductor, in which the doping changes fairly abruptly from n-type to p-type: for this reason it is also called an “np junction”. Of course a “pn junction” would also be an accurate name for it.

When the junction is formed, loosely bound dopant carriers from each side wander over the boundary, where they tend to become trapped: doped electrons from the n-type side will fall into the extra electron states provided by the p-type dopants on the p side, and





Figure 17: The basic ideas of of (top) n doping and (bottom) p doping.

extra holes from the p-type side will sit in the extra empty states provided by the n-type dopants on the n side. So when the junction is formed, the n side acquires a net positive charge, and the p side acquires a net negative charge. Once equilibrium is established, there is a potential boundary V_0 to the further flow of majority carriers: holes from the p side to the n side, and electrons from the n side to the p side. In terms of the energy bands, one says that across the junction, the energy levels of the electron states are “bent,” or shifted across the junction. Calling the absolute magnitude of the electron or hole charge q , the magnitude of the shift is qV_0 .




Figure 18: Band bending across an np junction.

In equilibrium, there is no net current flow across the np junction. Yet at any nonzero temperature, there will be minority carriers—holes on the n side and electrons on the p side—which would easily slide downhill. Call this minority leakage current I_0 . It is composed of thermally generated charge carriers, so

$$I_0 = C_0 e^{-\Delta/k_B T},$$

where C_0 has to do with the mobility of the minority carriers (it's like $\tau q^2/m$). But no net current flows in equilibrium. I_0 is cancelled by the small current of majority carriers that have enough thermal energy to get over the potential barrier:

$$I_M(0) = C_M e^{-qV_0/k_B T} = I_0.$$

where C_M depends on the majority carrier mobilities.

Now the way a diode works is that, if you apply a voltage V_a which lowers the potential barrier from qV_0 to $qV_0 - qV_a$, you exponentially increase the majority current flow:

$$I_M(V_a)/I_M(0) = I_M(V_a)/I_0 = e^{qV_a/k_B T},$$

so that

$$I_M(V_a) = I_0 e^{qV_a/k_B T}.$$

If you reverse the polarity of V_a to augment the barrier, you exponentially decrease the majority current flow. So the diode acts like a one-way (non linear) conductor for majority carriers. Calling the direction of majority current flow “positive”, the total current is then

$$I_T(V_a) = I_M(V_a) - I_0 = I_0(e^{qV_a/k_B T} - 1) = C_0 e^{-\Delta/k_B T} (e^{qV_a/k_B T} - 1).$$

This model is good as long as V_a is not too big compared to V_0 . For fixed temperature, you can step through different currents and trace out the I versus V_a curve, and fit for I_0 . After doing this for several temperatures, you can extract Δ (and C_0) from a logarithmic plot of I_0 against $1/k_B T$.

8.6 τ for a Semiconductor

You will be trying to measure Δ by ascribing the temperature dependence of the minority carrier current I_0 to the minority carrier concentration. You might worry whether it's safe to ignore the temperature dependence of τ .

The impurity scattering time for a semiconductor does depend on temperature, though not nearly as strongly as the minority carrier concentration. Here is one way to grasp the main result. For a Maxwell-Boltzmann gas, the *thermal wavelength* corresponding to kinetic energy $\frac{1}{2}k_B T$ of each quantum mechanical particle (of mass m) is

$$\lambda_T = \sqrt{\frac{\hbar^2}{2mk_B T}}.$$

This defines a temperature-dependent volume V_T per particle of $\lambda_T^3 \sim (k_B T)^{-3/2}$. Now $N_i V_T$ will be the number of scatterers in the same volume, so you would expect that for the impurity scattering rate, $1/\tau_i \propto (k_B T)^{-3/2}$, and $\tau \propto (k_B T)^{3/2}$.

The reason the analysis of the impurity scattering time for semiconductors differs from that in metals is that for metals, the charge density n is much larger (on the order of one electron per unit cell) and basically independent of T . Maxwell-Boltzmann statistics are no longer a good approximation, and the appropriate temperature to use for the “thermal wavelength” would be the Fermi temperature T_F (typically $T_F \sim 10^4 \text{K}$), which is proportional to $n^{2/3}$. In other words, the average spread of the charge carrier wavepackets is

metals about a lattice constant—much smaller than the typical impurity spacing—and so in metals the ballistic approach works better.

And what of the phonon scattering? If we were to repeat the same sort of analysis that we made for phonon scattering in a metal, the big difference would be that the number of electrons is smaller than for a metal. At room temperature or lower, the Fermi ball would be much smaller—and hence the geometric restriction, which is the source of the strong temperature dependence when the phonons freeze out, would no longer be in force. So phonon scattering in insulators has a comparatively weak temperature dependence, probably more like the Bloch-Grüneisen formula for high temperatures, ie $1/\tau_{ph} \propto T$.

So for the total τ , we'd expect

$$\frac{1}{\tau} = \frac{1}{\tau_i} + \frac{1}{\tau_{ph}} = aT^{-\frac{3}{2}} + bT,$$

and so

$$\rho \propto \frac{1}{n\tau} \propto e^{\Delta/k_B T} \times (aT^{-3} + bT^{-\frac{1}{2}}).$$

To extract Δ from a measurement of σ , you would write this as a function of $\beta = 1/k_B T$, and consider the logarithm:

$$\ln \sigma = -\ln \rho = -\Delta\beta + \dots + \mathbf{constants}.$$

The \dots terms are on the order of $\ln \beta$, and these are the ones you are thinking of ignoring. The rate of change of $\ln \beta = 1/\beta = k_B T$. Compare this to the linear coefficient Δ : at “low” temperatures, $k_B T/\Delta \ll 1$, and only the strong exponential temperature dependence matters when measuring σ .

8.7 Band Gap Temperature Dependence

If you take a more careful look at the theoretical considerations which lead to the band structure of a solid and the gaps between bands, it becomes apparent that the size of the band gaps depend on the separations between adjacent atoms of the lattice. Because the lattice parameters change slowly as a function of temperature, the band gap itself will vary. Experimentally, the band gaps of Si and Ge (the most common semiconductors for diodes) vary by around 10% between room temperature and 100K. The general decrease in band gap as temperature increases is consistent with the idea that the interactions that underly the band structure decrease as the lattice thermally expands.

9 Electronic Phase Transitions

Having seen ice melt and water turn to steam, you are familiar with the notion of phase transitions. In the case of water, the chemical identity of the molecules stays the same, but they occupy space and interact very differently with one another in the distinct physical phases. Often, you hear of these in a hierarchy: “solid, liquid, gas,” or “gas, liquid, solid.” In going one way or the other through the list, you notice that water vapor is rarefied, the molecules weakly interacting to the point that a gas will expand to fill whatever container it’s in. Liquid water is shapeless, but the molecules draw on each other enough that the fluid has a definite volume. And ice has a definite shape—the interactions in ice are surely the strongest.

A similar thing can happen with the electrons in a solid. You might think of the solid as defining the available space, and (for most, but not all, phase transitions) the electrons as the analogue of the water molecules. Some electrons, by interacting with each other, often by an intermediary such as the motions of the lattice, change in phase from an uncorrelated “gas,” to a much more correlated, more “condensed” substance—something more akin to a liquid or solid. Since the electrons are responsible for the electrical conductivity, the phase transition often has a definite signature in the resistivity measurement. Here are two electronic phase changes that you can investigate.

9.1 Itinerant Ferromagnetism in Nickel: Changing σ_a

At room temperature, a piece of nickel metal is magnetic. The magnetization \mathbf{m} is the relative difference between the densities of oppositely polarized spins. Below a Curie-Weiss temperature T_C , the spins of the electronic magnets tend to align with each other. But magnetic nickel is also a metal—in fact, most well-known magnets are also metals. Being a metal requires conduction electrons. These of course bounce around through the metal.

The question is, then, with all of this bouncing around, how the electrons in a magnet maintain their polarization—and whether the polarized ones are even the same as the conducting ones. It turns out that in magnetic nickel, the conduction electrons are indeed polarized. There is a model, due to Stoner, which explains how this can occur, and the same ideas have consequences for the resistivity.

In lone atoms of some d -series metals like nickel, the d electrons tend to align with one another. In a piece of nickel metal, the itinerant electrons may interact when two of them occupy d states on the same nickel atom, one right after the other. In that case there would be a large effective force causing the spins to align—it is energetically favorable for them to do so. Two such electrons, in mutual agreement, may go on their way, and interact with

other nickel atoms. On the other hand, two oppositely polarized electrons will not be able to go through the same nickel at nearly the same time, and the second electron will instead be deflected. In other words, consider the effective scattering cross section for the second electron to pass through a given nickel atom. If the first and second electron are polarized parallel to each other, you expect one cross section, $\sigma_{\uparrow\uparrow} = \sigma_{\downarrow\downarrow}$. If they are oppositely polarized, you expect a different cross section, $\sigma_{\uparrow\downarrow}$. The second scattering process is stronger: $\sigma_{\uparrow\downarrow} > \sigma_{\uparrow\uparrow}$.

But all of the electrons have to hop from nickel to nickel—there’s no other way to get around. In an environment where the majority of the electrons are already polarized in the same direction, it costs energy to point the other way. In Stoner’s model, T_C is the temperature at which this extra energy becomes scarce enough that the polarization interactions dominate. Which polarization gets the upper hand is a matter of chance, but once a sufficiently large local fluctuation happens to polarize enough electrons one way, the whole piece of nickel goes over. There will be the occasional “wrongly polarized” electron, but they don’t last for long in any one place. At lower temperatures, deeper into the magnetic state, the free energy scarcity worsens, and the number of electrons that have the “wrong” polarization become ever rarer.

In the Bloch-Grüneisen picture, the resistivity is proportional to σ_a , which is the average cross section for scattering by a single atom: the size of σ_a determines the high-temperature slope of the resistance versus temperature graph. From the above argument, this will change as the conduction electrons get polarized. As the temperature decreases below T_C , you would expect the resistance to drop below the trend for $T > T_C$, because as the magnetism sets in, polarization mismatches should become ever more rare. Right at T_C , there ought to be a kink in the resistance versus temperature curve as the cooperative effect suddenly kicks in. To see this, consider the rate of encounters between \uparrow and \downarrow electrons on the same nickel atom. The scattering rates will be proportional to

$$\sigma_{\uparrow\uparrow} \frac{n_{\uparrow}^2}{2} + \sigma_{\downarrow\downarrow} \frac{n_{\downarrow}^2}{2} + \sigma_{\uparrow\downarrow} n_{\uparrow} n_{\downarrow},$$

since $n_{\downarrow} + n_{\uparrow} = n$, and $n_{\uparrow} - n_{\downarrow}$ is the magnetization \mathbf{m} , and since $\sigma_{\uparrow\uparrow} = \sigma_{\downarrow\downarrow}$, this expression can be rewritten as

$$\sigma_{\uparrow\uparrow} \frac{n^2 + \mathbf{m}^2}{4} + \sigma_{\uparrow\downarrow} \frac{n^2 - \mathbf{m}^2}{4},$$

so the change in the resistivity due to \mathbf{m} should be proportional to

$$\mathbf{m}^2 \frac{(\sigma_{\uparrow\uparrow} - \sigma_{\uparrow\downarrow})}{4}.$$

Hence by measuring the “missing resistance” (the discrepancy between the resistance predicted by the $T > T_C$ trend and the actual $T < T_C$ resistance), you can estimate $|\mathbf{m}|$ as a function of temperature.

9.2 Superconductivity: More than $\tau \rightarrow \infty$

While the idea of perfect metallic conductivity is understandable within the context of the Drude model for independent electrons, *superconductivity* (which by contrast is an actually attainable state) is not. While normal currents in a metal decay by scattering, for supercurrents there is an energy gap—analogue to the band-gap that makes insulators work—which means that scattering itself requires energetic activation. Hence materials in the superconducting state are qualitatively different from merely excellent regular conductors.

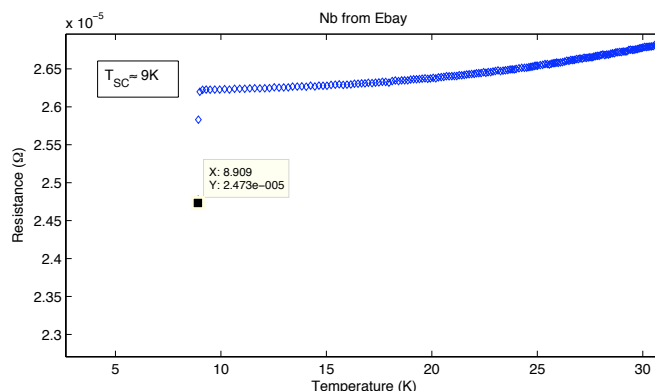


Figure 19: This ledge in the resistance indicates the onset of the SC transition to zero resistance. The sample was cut from Nb wire purchased from a jewelry supplier on an online auction site. The low temperature limit of the apparatus prevents you from seeing the full transition.

How does an energy-gapped state arise in a conductor, where the Drude electrons are “untethered?” In one fairly well-understood class of superconductors (Bardeen-Cooper-Schreiffer or BCS-type), the motions of the electrons become coordinated by sharing the energy benefit of lattice deformations. The overlap of the coordinated electron wavefunctions leads to a bulk coherent state, a quantum mechanical phenomenon, which is energetically favorable below the superconducting transition temperature T_{SC} . The electronic correlations themselves create a new “moving tether” between paired charges. But since all of the electrons are identical, they are *all* tethered into the same (macroscopic) quantum state.

Superconductors are therefore qualitatively different from “perfect” conductors where we picture each Drude electron moving on its own. The difference is especially apparent in a superconductor’s response to an applied magnetic field. First consider the theoretical response of a perfect metal to a changing magnetic field. According to Faraday’s law of in-

duction, the changing magnetic field would induce current loops that oppose the change in the field. Hence the field in a perfect conductor would be stuck at its initial value. On the other hand, a superconductor tends either to expel (type I) or sequester (type II) magnetic field lines so that most of the sample volume is free of magnetic field. A superconductor spontaneously generates electric currents that keep magnetic fields from the bulk of the sample.

Here's a rough sketch of how this large diamagnetic response, known as the *Meissner effect*, comes about. So we don't have to worry about boundary effects, consider an infinitely long, thick slab of superconductor. It is always tempting to think of electrons like little balls of charge, with velocity \vec{v} , such that $\vec{j} = ne \langle \vec{v} \rangle$. But to construct a localized particle out of waves requires a particular sum of momentum states. Superconductivity depends explicitly on the correlated multi-electron wavefunction, and this single-particle picture breaks down: the momentum states are added up in a very different way. Because of the energy gap between this ground state and the single-particle-like excitations, the applied vector potential can't change the shape of the wavefunction by much.

It seems paradoxical that current can both flow and not flow in the “ground” state, depending on whether we are applying a magnetic field. The resolution to this paradox resides in the difference between the $\langle m\vec{v} \rangle$ -momentum of a localized particle and the canonical momentum $-i\hbar\nabla$ in the Hamiltonian. In the presence of a vector potential \vec{A} , the operator for $m\vec{v}$ is modified from $-i\hbar\nabla$ to $-i\hbar(\nabla - e\vec{A}/c)$. Since the ∇ operator describes the shape of the wavefunction in space—which can't change because of the energy gap—the expectation value of current $\langle ne\vec{v} \rangle$ must be proportional to $\langle -ne^2\vec{A}/mc \rangle$. So the vector potential acts as a direct push on the *totality* of “particles” that comprise the superconducting state, resulting in something like a Galilean translation of a rigid solid. By applying the appropriate Maxwell's equations, you will find that this supercurrent exponentially screens out a magnetic field.

Superconductivity was first discovered in solid mercury by Heike Kammerlegh Onnes in 1911, soon after he had succeeded in the liquefaction of helium. Superconducting wires with the ability to carry large currents are used in the production of the very strong cryogenic electromagnets used in diagnostic NMR machines. In the 1980's, Bednorz and Muller discovered superconductivity at the relatively high temperature of 35 K in a layered copper-oxide. As part of this lab, you have the opportunity to make (or, if time doesn't permit, at least measure) a similarly structured compound that becomes superconducting anywhere from 70 – 90 K, depending on the details of how you produce it. Although the recipe for making this material is well known, there is as yet no consensus about the interaction which makes its superconducting transition occur.

10 Theory References

- **The Solid State: From Superconductors to Superalloys**, André Guinier and Rémi Jullien, pp. 74–120 (fundamentals of electrical conductivity), 144–156 (ferromagnetism).
- **Principles of Solid State Physics**, J.M. Ziman, chapter 2 (phonons), chapter 7 (transport properties), chapter 10 (magnetism) and chapter 11 (superconductivity).
- **The Feynman Lectures on Physics, vol. III**, R.P. Feynman, chapter 13 (on the electron states in a crystal lattice), chapter 14 (semiconductors), and chapter 21 (superconductivity).