A Resonant Circuit

I. Equipment

Oscilloscope
Capacitor substitution box
Resistor substitution box
Inductor
Signal generator
Wires and alligator clips

II. Introduction

A circuit consisting of a capacitor, an inductor, and a resistor, called an LRC circuit, is a perfect analog to a damped mechanical oscillator. Although they are fairly simple systems, oscillator have a rich variety of behaviors. Oscillators in one guise or other occur in so many systems in nature, and in so many artificial devices as well, that they are arguably the most important systems in physics. The purpose of this lab is to use the LRC circuit to bring out some of the characteristics of oscillators. Along the way we'll also make concrete some rather arcane features of AC circuits.

We'll be building on the experience gained in your earlier lab, when you explored the discharge of an RC circuit using the oscilloscope. The only new element we're adding is an inductor.

First, a little theory. As you've learned in class, a capacitor stores charge, and develops a voltage in response to that charge. We're just getting to *inductors*; these are little coils of wire which develop a *back-emf* when the current through them *changes*.

Now think for a moment of a mass on a spring, the simplest of mechanical oscillators. When the mass is momentarily at rest at the end of its travel and the spring is compressed, all the energy is in the potential energy of the spring. When the mass is whipping through the middle of its travel and the spring is relaxed, all the energy is in the kinetic energy of the mass. So it goes, with the energy continually changing hands between kinetic and potential. The spring can hold the energy in the stretching or squeezing of the bonds between its atoms, or the inertia of the mass stores energy as kinetic energy. The oscillator doesn't simply stop dead at the end of its travel, because the spring's restoring force starts pushing the mass back toward the middle. And it doesn't stop dead at the middle of its travel, because the mass' inertia carries it through to compress the spring again.

If we have a simple circuit consisting of an inductor and a charged-up capacitor, and then we close the circuit, a very similar situation happens. The electrical energy can be stored on the capacitor (you'll recall that the energy is $\frac{1}{2}CV^2$, or it can be stored in the magnetic field of the inductor. When the switch is closed, the charge starts to drain off the capacitor through the inductor. This builds up the current through the inductor, until the capacitor is completely discharged, but at that point the inductor has a big current in it which can't shut off immediately; the inductor fights abrupt changes of current. This makes the current act as if it had 'inertia'; accordingly, the capacitor charges up in the opposite polarity! As the capacitor is charged, it acts like a compressed spring and starts pushing charge back the other way, so the current decreases; when the current has decreased to zero, the capacitor has all the charge it can get, and its voltage now starts forcing a current back through the inductor. The current builds up in the inductor, until the capacitor is discharged; but again the current in the inductor can't stop immediately, so it charges the capacitor back up just like it was at the start. And so on and on and on, indefinitely. *The capacitor acts like a spring, and the inductor acts like the inertia of the mass.* When the current is maximum, the inductor has all the energy in it, in the form of the energy of the magnetic field. When the current is zero, the capacitor has the maximum charge on it, and all the energy is in the form of the electric field. So the charge and current oscillate back and forth, and energy changes form between electric and magnetic, just as energy changes form between potential and kinetic in the mechanical oscillator.

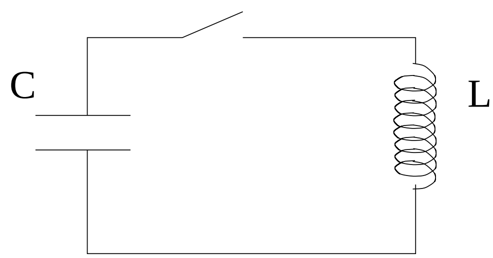

LC circuit.

Fig. 1: Note the circuit symbol for the inductor looks like a spring, but it actually plays the role of the *inertia* in the circuit; the capacitor behaves like the spring.

How long will it take for the charge and current to go through one full cycle? The analysis for this is a little complicated, but not too bad. Let's let $q(t)$ be the charge on the capacitor. A current which tends to *dis*charge the capacitor will then be $i = -dq/dt$. The voltage across the capacitor is $q(t)/C$, and the emf across the inductor is $-L\,di/dt = L\,d^2q/dt^2$. Since the sum of all these voltage drops around the circuit must be zero, we have

$$L\frac{d^2q}{dt^2} + \frac{q}{C} = 0,$$

or

$$\frac{d^2q}{dt^2} + \frac{q}{LC} = 0.$$

If you recall from Physics 3, a similar analysis of a mass on a spring gives

$$\frac{d^2x}{dt^2} + \frac{kx}{m} = 0.$$

Mathematically, these two equations are identical and must have identical solutions. If you'll recall, the angular frequency of the oscillation for a mass on a spring is $\sqrt{k/m}$; so by inspection, we have the *natural frequency* of an LC circuit,

$$\omega_0 = \frac{1}{\sqrt{LC}}.$$

If we start with some charge Q_0 on the capacitor, we expect

$$q(t) = Q_0 \cos \omega_0 t,$$

and for the voltage across the capacitor we'll have

$$V_C(t) = \frac{Q_0}{C} \cos \omega_0 t.$$

The current, defined positive as the capacitor drains, is

$$i(t) = -\frac{dq}{dt} = Q_0 \omega_0 \sin \omega_0 t.$$

Now, a real circuit won't oscillate forever, because the sloshing charge loses energy to the resistance in the wires (and as we'll see, even if there were no resistance, there would be energy lost to electromagnetic radiation.) A more realistic circuit is the LRC circuit, in which a resistor is in series with the capacitor and the inductor:

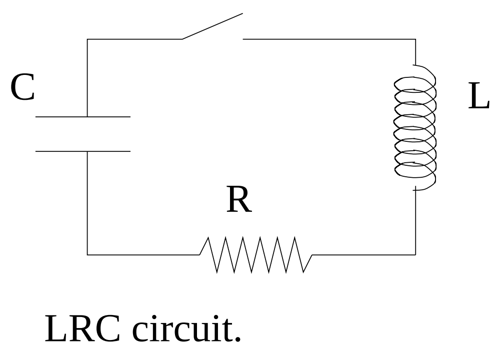

Fig. 2: The LRC circuit, with a switch. Imagine that the capacitor is charged up and the switch is closed; this will provoke a transient ringing response in the circuit, provided R is not too large.

How will this respond when the switch is closed? If R isn't too large, the effect is fairly simple; the circuit still oscillates, but the oscillation dies away exponentially, with a time constant related to the resistance – the more resistance, the shorter the time constant. Specifically, the charge on the capacitor as a function of time looks like

$$q(t) = Q_0 e^{-t/\tau} \cos \omega t,$$

where

$$\tau = \frac{2L}{R} \quad \text{and} \quad \omega = \sqrt{\omega_0^2 - \frac{1}{\tau^2}}.$$

If unless R is pretty big, ω is almost the same as ω_0, so we can just use $\omega \approx \omega_0 = 1/\sqrt{LC}$.

What will this look like? The $e^{-t/\tau}$ part is a falling exponential, just like in the RC circuit – at the end of every time interval τ, this function has died away to $1/e$ (roughly 0.37) of the value it had at the beginning of the interval. Multiplying these together gives a function something like this:

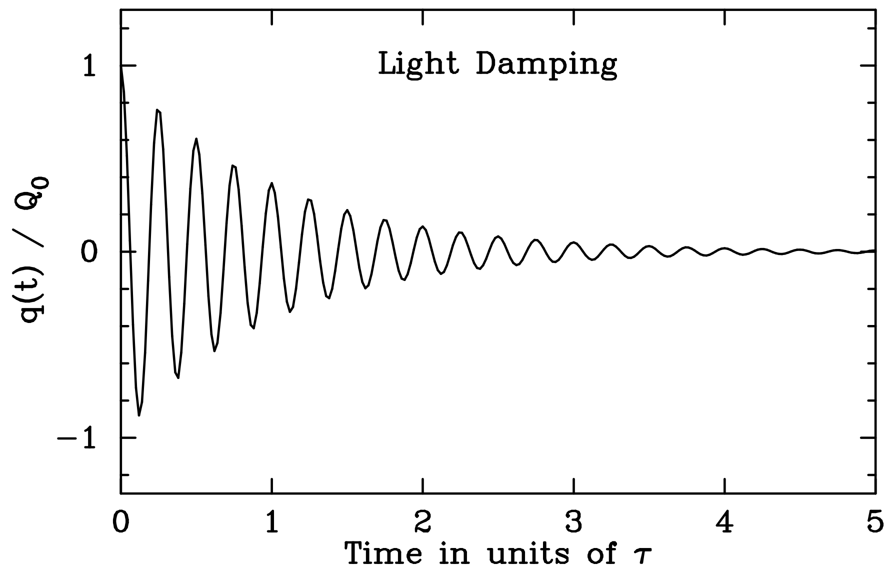

Fig. 3: Decay curve for fairly light damping. I've arbitrarily chosen parameters so that the sinusoid oscillates four times for each τ, i.e., during the time it takes the amplitude to decay by a factor $1/e$. Note how the oscillation dies away after an interval of several times τ has passed.

Curiously enough, if you turn the damping way up – that is, if you make the resistance bigger – you can make the oscillation go away altogether. To make the oscillation go away, it turns out that you need to make

$$R \geq 2\sqrt{\frac{L}{C}}.$$

When these are equal, the oscillation barely vanishes, and the situation is called *critical damping*; when R is greater than $2\sqrt{L/C}$, this is called *overdamping*.

This is easy to understand using a mechanical analogy. Suppose you have a pendulum with an extension on it which swings in some fluid which drags on the pendulum. If you swing it in some low-viscosity fluid like water, the pendulum would swing back and forth for a while, with its amplitude slowly decreasing. But if you use gooier and gooier liquids, the pendulum will swing less and less; finally, when you make the fluid gooey enough, the pendulum will just drift in and come to rest at the middle, without developing enough momentum to overshoot its equilibrium point. The gooiness of the liquid is just like the resistance in the circuit.

So far we've been concentrating on the *transient* response of the LRC circuit – the response of the circuit to a sudden change, like closing the switch with the capacitor all charged up. In the lab, we'll study this by driving the circuit with a relatively low-frequency square wave – the abrupt change in the driving voltage is just like closing a switch, and the circuit 'rings' in response.

But the transient response is only part of the story. There are also interesting behaviors in response to a sinusoidal driving voltage, which turns the circuit into a *damped, driven, harmonic oscillator.* So now let's consider this circuit:

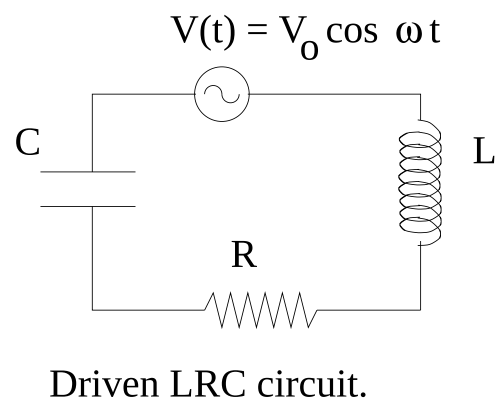

Fig. 4: Driven LRC circuit. The little yin-yang-like thing is a power supply delivering a sinusoidally varying voltage. The sum of all the voltage drops around the circuit must be zero, so the sum of the voltages across the L, R, and C must be the same as the driving voltage.

Here, instead of just closing a switch to drive the circuit, we're putting in a voltage which follows a sinusoidal curve with time. The driving frequency is ω_d – it can be any frequency, and in particular it need not be the same as the circuit's 'natural' frequency which it follows in its transient response.

When we first turn the sinusoidal voltage on, the circuit will respond in a complicated way, vibrating both at the natural frequency of the circuit and at the frequency of the driving voltage. But the vibrations at the natural frequency will soon die away, because the resistance in the circuit eats them up. The initial vibrations are called *transients* because they usually go away pretty quickly. After the transients are gone, the currents in the circuit will be oscillating *only* at the driving frequency – these oscillations persist because the driving voltage keeps feeding energy into them.

The situation we have now is, again, identical to the oscillators you saw in Physics 3 (see pp. 432 – 435 in Hecht's book for a quick qualitative discussion). Basically, if the driving frequency is close to the natural frequency, one sees big currents in the circuit. If the driving frequency is far from the natural frequency, the currents are smaller. If you sweep the driving frequency from way below the natural frequency to way above it, you'll see a substantial increase in current as you pass the natural fequency. This is called a *resonance* of the circuit. The less resistance there is in the circuit, the more dramatic the increase in response at the natural frequency, and the sharper the 'resonance peak' becomes. We can easily see this in the lab.

It's instructive to pick apart this circuit in some detail – the book does this on pp. 844 – 857. Because this material is so dense, it might help to have another exposition of it to bring out the main points, so here it is.

In order to understand what's going on, we need to think about the response of each circuit element to a sinusoidal driving voltage. So let's do each one in turn.

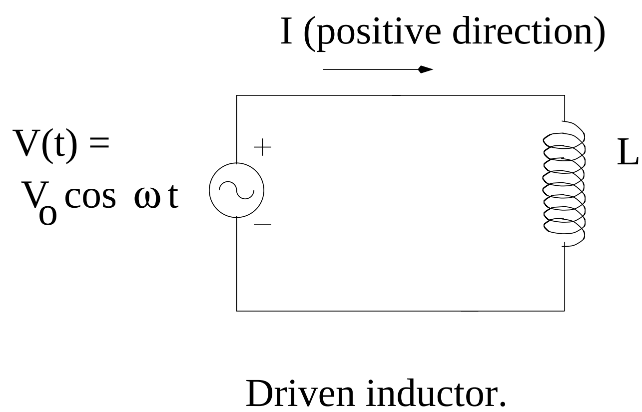

Fig. 5: A driven inductor, all by itself. The arbitrary direction of 'positive' voltage is shown, as is the direction of 'positive' current. The back EMF of the inductor must balance the driving source at each moment.

First, consider the inductor. You'll recall that it develops a back EMF, given by

$$\mathcal{E} = -L\frac{dI}{dt}.$$

Suppose we hook up our sinusoidal driving voltage (from a signal generator) to an inductor only, and assume that the resistance in the circuit is zero. At first glance you might think that the inductor would behave as a short circuit, and infinite currents would flow – but that ignores the essential point that *inductors 'fight' changes in current.* Since the signal generator's voltage is changing all the time, the current is changing all the time, and the inductor is going to fight it.

This leads to a very interesting relationship between the current and the voltage. If we go around the circuit, all the voltage drops have to sum to zero – that's Kirchoff's loop rule. Consider the moment when the voltage is zero. At that moment, the inductor must also have zero EMF, which means that the current through it is *not changing.* But the current is also following a sinusoid, and the only time a sinusoid is not changing is at its maximum or miminum. So the *current is largest when the voltage is zero.* Now let's consider the time when the voltage is *largest* (that is, most positive*) This voltage will be trying to *increase* the current. Since the *voltage* is maximum, the *rate of change* of the current in the inductor must be maximum, and the current must be changing toward more positive values. A sinusoidal curve has its steepest (positive) *slope* as it crosses zero from minus to plus. So *as the voltage goes over its maximum, the current is just crossing zero on its way up.* You can see from this that the *current across an inductor will lag behind the voltage across it, by 1/4 of a cycle.*

How much current will we get? If the inductance is very large, a small rate of change of current will generate a big EMF, so the currents will be small if the inductance is large. Furthermore, if the frequency of the driving signal is high, the voltages change quickly, so the current will be small. Considerations of this kind lead to the idea that the inductor has a sort of 'resistance-like' quality, called *reactance*, which is given by

$$X_L = \omega L.$$

So our old rule that

$$V = IR$$

becomes in this case

$$V_0 = I_{\max} X_L = I_{\max} \omega L.$$

From this you can see that L must have units of ohm seconds, since ωL must have units of ohms.

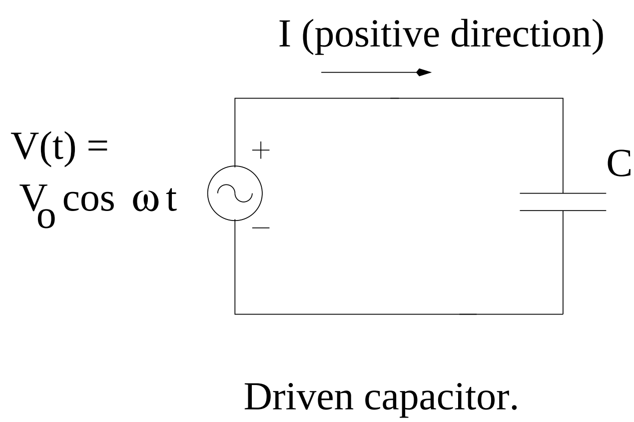

Fig. 6: A driven capacitor, all by itself. The charge on the capacitor at each instant must be enough to give a voltage balancing the power supply. This leads to a different relationship between voltage and current than in the inductor.

* For puroposes of analyzing the circuit, we have to arbitrarily label one side of the AC source '+' and the other one '−'. Then, when a current circulates from + to − we call it a positive current, and when the current is going the other way, we call it negative.

Now let's play the same game with a capacitor hooked up to the oscillator all by itself. Here, the capacitor always holds just the right amount of charge to balance the voltage of the power supply. To find the current, we only need consider how fast the charge is going onto or leaving the capacitor. When the capacitor is fully charged, the voltage is maximum, but the current is zero (since at that moment no charge is flowing in or out of the capacitor). As the voltage comes up through zero, the charge on the capacitor is changing at a maximum rate to track the voltage, so the current is maximum. So now, as the *current* goes through its maximum, the *voltage* is just going through zero on its way up. So now the *current leads the voltage, by 1/4 cycle*, exactly opposite to what happens with the inductor.

How much current will flow? The greater the capacitance, the more charge will need to slosh back and forth to keep up with the voltage, so the more current we'll have. And the higher the frequency, the faster the charge will slosh back and forth, so the greater the current. Reactance is like resistance, so the more current flows, the smaller the reactance. That means that the capacitive reactance will be *inversely* related to both ω and C. A more detailed analysis shows that

$$X_C = \frac{1}{\omega C}.$$

Once again, inductance and capacitance are like opposite sides of the coin – inductors choke current, and capacitors like current. The voltage across an inductor leads the current, and the voltage across the capacitor lags the current.

Finally, let's look at the easy case – a sinsuoidal voltage source connected across a simple resistor. Now Ohm's law holds instant-by-instant, and the current and the voltage across the resistor stay right in step.

We're finally in a position to understand what happens when we drive this circuit. The voltage across the inductor will lead the current, and the voltage across the capacitor will lag the current, while the voltage across the resistor keeps pace with the current. Hecht's book explains a very nice way of keeping track of all this on pages 852 – 853, the idea of *phasor addition*. We think of each voltage as a little arrow with its tail at the origin, and we let these little arrows spin round and round at the driving frequency ω_d. We represent the maximum voltage across the resistor with one arrow – since the current tracks the voltage across the resistor, this arrow also serves to represent the current. As with any 'phasor', the real voltage which this phasor represents moment by moment is just the projection of the phasor on one of the axes as the little arrow spins round. Another arrow, which we draw at 90 degrees to the first arrow, represents the maximum voltage across the inductor. We draw this 90 degrees *ahead* of the voltage across the resistor, since the voltage across the inductor *leads* the current in phase. Then we draw another arrow opposite this, to represent the maximum voltage across the capacitor. This is all shown in Fig. 21.18a, on page 853. *The (vector) sum of these three arrows must represent the driving voltage*, by Kirchoff's loop rule. So we can see at a glance how whether the net current leads or lags the driving voltage.

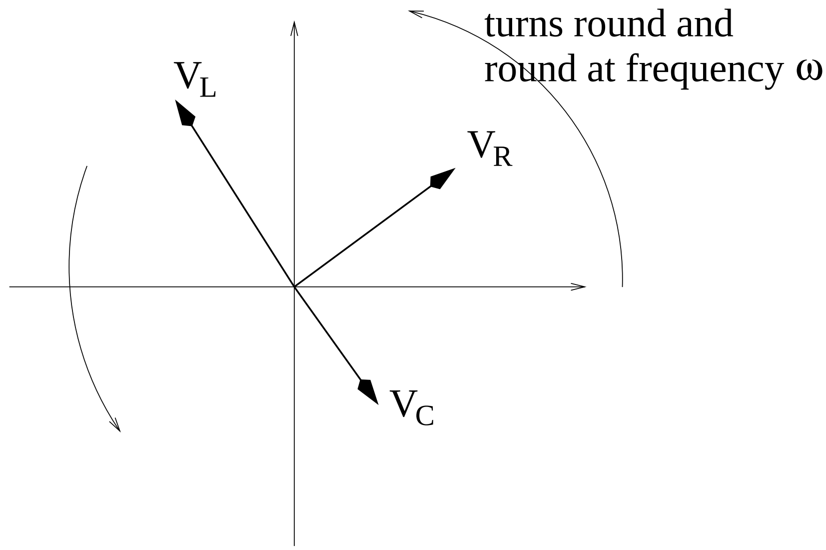

Fig. 7: Voltage phasors in an LRC circuit. Think of these as all whirling about at frequency ω; the actual voltage across any component at any instant is the projection of its phasor on an axis (maybe the horizontal axis, it doesn't matter). The length of each arrow is proportional to the resistance or reactance of that component. The sum of all the arrows (not shown) is the total voltage of the driving source. Because the current stays in step with the voltage across the resistor, the resistor's phasor also indicates the phase of the current.

This vector addition will of course depend on how long each of these vectors is. Remember that

$$V = IR,$$

and the reactances are like resistances, so that the maximum voltage across the inductor is just $V_{\max\ \mathrm{L}} = I_{\max} X_{\mathrm{L}}$, and the maximum voltage across the capacitor is $V_{\max\ \mathrm{C}} = I_{\max} X_{\mathrm{C}}$. So the relative lengths of these voltage phasors are proportional to the relative sizes of the reactances and the resistance. If the inductive reactance is large (big L, big ω), the inductor will have a big voltage across it, and the inductor will dominate; the voltage will lead the current. If the capacitave reactance is large (small C, low ω), the capacitor will dominate, and the voltage will lag the current (See Figs. 23.19 and 23.20 on p. 879). If the two just balance, then the voltage will be right in phase with the current. For this to happen,

$$X_L = X_C \quad \text{or} \quad \omega L = \frac{1}{\omega C}.$$

When this happens, we have

$$\omega = \sqrt{\frac{1}{LC}}$$

... the natural oscillation frequency, which is also the 'resonant frequency'. At this special frequency, the current and voltage are just in step, you get the largest current, and the maximum power is transferred into the circuit.

III. Procedure

With this very large helping of theory, we can proceed to actually do the lab. Luckily this is a little simpler.

Transient effects. If we hooked up the circuit in Fig. 2, we'd only get one chance to look at the transients each time we closed the switch, then we'd have to charge up the capacitor again, and so on. That would be a drag. So, what we do instead is just connect the circuit to a signal generator, and set it up to generate a square wave. Each time the voltage jumps, the square wave 'kicks' our circuit, leading to a transient effect. We want the square wave to repeat slowly enough that the transients in the circuit have time to die away over each cycle of the square wave, but other than that, the exact frequency of the square wave doesn't matter.

The first order of business, as usual, is to be sure your oscilloscope works. Attach the signal generator to the oscilloscope and fiddle until you can see the square wave nicely. You'll want a square wave with about a 1 kHz repetition rate.

Now connect the inductor and the capacitor substitution box in series with the power supply. Set the capacitor box to a lowish value of C, so your resonant frequency is high. Leave the resistor substitution box out of the circuit for the moment. Attach the oscilloscope across the capacitor. Fiddle until you can see a transient curve something like Fig. 3. Adjust the capacitor box to higher and lower values of C, and notice the effect on the frequency of the fast oscillation.

Select a value of C and measure the frequency of the fast oscillation. Compare your result to the theoretical value of ω. It's most accurate if you count off several cycles of the oscillation, figure out how many horizontal divisions this corresponds to, and find the period for one cycle P by dividing the time by the number of cycles. Don't forget that $\omega = 2\pi/T$.

Now, measure how long it takes for the amplitude of the oscillations to decay by a factor $1/e$, just as you did in the RC circuit in a previous lab. From this observation, and the value of L, determine the R in your circuit. There is resistance in your circuit, even though there's no resistor box – it's worth keeping this in mind. It might be interesting to measure the value of R for your inductor by hooking it to a digital multimeter, and comparing to what you get. Why can't you measure the resistance of a capacitor the same way?

Now, put the resistor box into the circuit, set to its lowest value. Find the trace again. Crank up the resistance and watch the trace. When the resistance you've added is equal to the resistance you inferred for the rest of the circuit, you should have cut τ in half; did you? Keep cranking the resistance up and see if you can see critical damping. Compute using L and C how much resistance you should need for critical damping, and compare to the total resistance in your circuit.

The Driven LRC Circuit. Before doing this, note the resonant frequency of your circuit for the setting of C you've been using. Set your resistance to a moderate value, well below critical damping. (Don't forget to convert the ω to the number of cycles per second f, by dividing by 2π.)

Switch your signal generator over to generate sine waves. Set up the oscilloscope to look at the voltage across the resistor. Set your signal generator to the range at which you expect the resonant frequency, and use the dial on the signal generator to tune up and down. Can you see the resonance? How does the signal generator reading at the maximum amplitude compare to the theoretical value of the resonance?

Let's see if we can get some curves similar to Fig. 23.22 (p. 881). Set your resistor box on its lowest value (what's the resistance in the circuit, now?). Tune the signal generator through the resonance and measure the amplitude at about a dozen frequencies spaced to cover the resonance. Plot these amplitudes as a function of frequency in your lab book. Now double the effective resistance of the circuit, and do this again. Do your curves match expectations?

Now try this. Put your oscilloscope across the inductor, and sweep through the frequencies. Note the amplitude as a function of frequency. You may have to use two of the factor-of-10 ranges on the signal generator to see the limiting behavior of the amplitude at low and high frequencies. Do the same for the capacitor. Which would you expect to show a large voltage at high frequency? Which at low frequency? Do your observations support this?

Finally, let's explore the *phase relationships* we worked out earlier. Here we'll use a new feature of the oscilloscope, the ability to look at two signals at once. The oscilloscope has an 'A vs. B' setting. This uses the input from one channel to drive the horizontal axis of the oscilloscope, and the input from the other to drive the vertical axis. So now, instead of having the little spot sweep across the screen at a constant rate, it vibrates about over the screen in response to the two inputs. The relation between the horizontal and vertical scales is arbitrary – you can set the gains independently, and you'll have to – but the phase relationship is set by the input signals.

Think about this setup for a bit. You're putting a sinusoidal signal into each axis. When the signals driving the two axes are in phase, the little spot is driven back and forth in a straight line at some angle to the axis (ideally 45 degrees if the two signals are the same amplitude and the gain settings on the two axes are the same). But when the two signals are 90 degrees out of phase, the little spot moves around in an ellipse with its long axis horizontal or vertical. When the two signals are some other angle out of phase, you get an ellipse with its long axis at some angle to the x and y axes.

Try plotting the voltage across the resistor on the horizontal axis, and the voltage across the power supply on the vertical. The resistor voltage keeps track with the current in the circuit, while the power supply voltage is of course the driving force. What do you expect to happen as you tune the power supply through the resonance? Where do you expect a straight line? Do you see the expected behavior? Try the same with the capacitor and the inductor.