MIT OpenCourseWare
http://ocw.mit.edu

8.02 Electricity and Magnetism, Spring 2002
Transcript – Lecture 20

Today, I will quantify the ability of a circuit to fight a magnetic flux that is produced by the circuits themselves.

If you have a circuit and you run a current through the circuit, then you create some magnetic fields, and if the currents are changing then the magnetic fields are changing.

And so there will be an induced EMF in that circuit that fights the change, and we express that in terms of a self-inductance: L, self-inductance, and the word self speaks for itself.

It's doing it to itself.

Magnetic flux that is produced by a circuit is always proportional to the current.

You double the current, the magnetic flux doubles.

And so it is the proportionality constant that we call L, that is the self-inductance, and so therefore the induced EMF equals minus d phi/dt, that is Faraday's Law.

And so that becomes minus L dI/dt.

L is only a matter of geometry.

L is not a function of the current itself.

I will calculate for you a very simple case of the self-inductance of a solenoid.

Let this be a solenoid and this is a closed circuit, and we run a current I through the solenoid, and the radius of these windings is little r.

Let's say there're N windings and the length of the solenoid is little l.

Perhaps you'll remember that we earlier derived, using Ampere's Law, that the magnetic field inside the solenoid is mu 0 times I times capital N divided by l.

This is the number of windings per meter.

If we attach an open surface to this closed loop, very difficult to imagine what that open surface looks like -- we discussed it many times -- inside this solenoid you have sort of a staircase-like of surface.

That magnetic field penetrates that surface N times because you have N loops.

And so the magnetic flux, phi of B, is simply the area by little r squared, which is the surface area of one loop, because I assume that the magnetic field is constant inside the solenoid, and I assume that it is 0 outside, which is a very good approximation.

So we get pi little r squared surface area of one loop, but we have N loops and then we have to multiply that by that constant magnetic field.

So we get an N squared because we have an N here, mu 0 I divided by L.

And this we call L times I.

That's our definition for self-inductance.

And so the self-inductance L is purely geometry.

It's pi little R squared, capital N squared divided by L times mu 0.

Let me check this.

Pi little r squared, I have a capital N squared, mu 0, that's correct, divided by little l.

And so we can calculate, for instance, what this self-inductance is for a solenoid that we have used in class several times.

We had one whereby we had 2800 windings.

R I think was something like 5 centimeters -- you have to work SI of course, be careful -- and we had a length, was 0.6 meters.

We had it several times out here, and if you substitute those numbers in there, you will find that the self-inductance of that solenoid is 0.1 in SI units and we call the SI units Henry, capital H.

It would be the same as volt-seconds per ampere, but no one would ever use that.

We call that Henry.

Every circuit has a finite value for the self-inductance, however small that may be.

Sometimes it's so small that we ignore it, but if you take a simple loop, a simple current, just one wire that goes around -- whether it is a rectangle or whether it is a circle it doesn't make any difference -- it always produces a magnetic field.

It always produces a magnetic flux through the surface, and so it always has a finite self-inductance.

Maybe only 9- nano Henrys, maybe only micro Henrys, but it's never 0.

And so now what I want to do is to show you the remarkable consequences of the presence of a self-inductor in a circuit, and I start very simple.

I have here a battery which has EMF V.

I have here a switch, and here are the self-inductor.

We always draw a self-inductor in a circuit with these coils, and we also have in series a resistor, which we always indicate with this, these teeth.

And I close this switch when there is no current running.

In other words, at time T equals 0 when I close the switch, there is no current.

When I close this switch the current wants to increase, but the self-inductance says uh-huh, uh-huh, take it easy, Lenz law, I don't like the change of such a current.

So the self-inductance is fighting the current that wants to go through it.

There comes a time that the self-inductance loses the fight, if you wait long enough, and then of course the current has reached a maximum volume, which you can find with Ohm's Law, because the self-inductance itself has no resistor.

Think of the self-inductance as made of super-conducting material.

There's no resistance.

And so without knowing much about physics, you can make a plot about the current that is going to flow as a function of time.

You start out with 0 and then ultimately, if you wait long enough, you reach a maximum current which is given by Ohm's Law, which is simply V divided by R.

And you slowly approach that value.

And how slowly depends on the value of the self-inductance.

If the self-inductance is very high, it might climb up like this, so this is a high value for L.

If the self-inductance is very low, that is a low value for L.

If the self-inductance were 0, it would come up instantaneously, but I just convinced you that there I no such thing as 0 self-inductance.

There's always something finite, no matter how small.

And so this is qualitatively what you would expect if you use your stomach and if you don't use your brains yet.

There's nothing wrong with using your stomach occasionally, but now I want to do this in a move civilized way, and I want to use my brains, and when I use my brains I have to set up an equation for this circuit.

And if you read your book, you will find that Mr. Giancoli tells you to use Kirchhoff’s Loop Rule.

But Mr. Giancoli doesn't understand Faraday's Law, and he's not the only one.

Almost every college book that you read on physics do this wrong.

They advise you to use Kirchhoff’s Loop Rule, which says that the closed loop integral around the circuit is 0.

That, of course, is utter nonsense.

How can it be 0?

Because there is a change in magnetic flux, and so it can only be minus d phi/dt.

I advise you to go to the 8.02 website and download a lecture supplement that you will find in which I address this issue and hit it very hard.

So the closed loop integral of E dot dL, if you go around the circuit, is not 0, is minus d phi/dt -- Faraday's Law -- so it's minus L dI/dt.

So we have to go around to circuit and we have to apply Faraday's law and not Kirchhoff’s Loop Rule, which doesn't apply here.

This is the plus side of the battery and this is the minus side, so the electric field in the battery is in this direction.

The electric field in the self-inductance is 0 because the self-inductance has no resistance, it's super-conducting material, and so the electric field in the resistor -- if the current is in this direction, which it will be -- then the electric field in the resistor will be in this direction.

So now I am equipped to write down the closed loop integral of E dot dL.

I start here and I always go in the direction of the current, and I advise you to do the same.

I don't care that you guess the wrong direction for the current.

That's fine.

Later, minus signs will correct that, they will tell you that you really guessed the wrong direction, but always go around the loop in the direction of the current, because then the EMF is always minus L dI/dt.

If you go in the direction opposed to the current, then it is plus L dI/dt and that could become confusing.

So I always go in the direction of the current, and so I first go through the self-inductance.

There is no electric field in the self-inductance, so the integral E dot dL -- in going from here to here -- is 0.

This is where the books are wrong.

It is 0.

Now, I go through the resistor, and so now I get plus IR, E and dL are in the same direction.

Ohm's Law tells me it's IR.

In the battery, I go against the electric field, and so I get minus V.

That now equals minus L dI/dt, and this is the only thing and the only correct way to apply Faraday's Law in this circuit.

You can write it a little differently, which may give you some insight.

For instance, you could write that I can bring V and the L, and L, to one side -- so I can write down that V minus L dI/dt equals IR.

It's the same equation when you look, and the nice thing about writing it this way is that since dI/dt is positive here -- it's growing in time -- the induced EMF, which is this value -- notice it's always opposing the voltage of my battery -- and that's what Lenz's Law is all about.

It's not until dI/dt has become 0 that V equals IR, and that happens, of course, if you wait long enough.

And so we have to solve that differential equation, and what is often done that you bring all the terms to the left side and that you get an, a 0 on the right side.

And so what you often see is that L dI/dt plus IR minus V equals 0.

And because we have a 0 here, some physicist thinks that this is an application of Kirchhoff’s Rule.

This is nonsense.

You can always make it 0 here by bringing all the terms to this side.

The closed loop integral of E dot dL is not 0, the closed loop integral of E dot dL is minus L dI/dt, but when I shift minus L dI/dt to this side I get 0 here.

And of course the people who write these books know that this is the right answer and so they manipulate it so they get this equation and they call that Kirchhoff’s Rule.

Sad, and also embarrassing.

So, this is the equation that you have to solve.

Some of you may have solved this equation in 8.01 already.

Surely you didn't have an I here.

You may have had an X here for the position, but you probably solved it when you had friction.

Maybe you didn't.

I will give you the solution to that differential equation.

It's a very easy solution.

The current as a function of time is a maximum value times 1-e to the minus R divided by L times T, and I max -- that is the maximum current -- is V divided by R.

And let's look at this in a little bit more detail.

First, notice that when t equals 0 that indeed you find I equals 0.

Substitute in here t equals 0, you get 1-1.

So you find, indeed, that I equals 0.

Substituting there t goes to infinity, then you find that I indeed becomes V divided by R, which is exactly what you expect.

If t becomes infinity, then clearly the self-inductance has lost all its power, so to speak, and the current is simply V divided by R, the maximum current that you can have.

And so that's a must, that's a requirement.

If you wait L over R seconds -- and believe it or not, if you have some time, convince yourself that L over R indeed, as a unit, is seconds -- then the current I is about 63% of I max, because if t is L over R, then you get 1-1 divided by E, and that is about 0.63.

And if you wait double this time, then you have about 86% of the maximum current.

In other words, right here -- if I wait L over R seconds -- this value here is about 0.63 times the maximum value possible, and it's very, it's climbing up and asymptotically approaches then, ultimately, the maximum current which is V divided by R.

Make sure you download that lecture supplement that you'll find on the Web.

Now, what I'm going to do is all of a sudden I'm going to make this voltage 0.

The way I could do that is by simply shorting it out.

Of course on the blackboard, I can simply remove it.

So it's not there.

The current is still running and all of a sudden at a new time -- t 0, I define the time t 0 again -- the voltage is 0.

And now comes the question what is now going to happen?

Well, the self-inductance doesn't like the fact that the current is going down, so it's going to fight that change, and so you expect that the current is not going to die off right away, but you expect that the current is going to go down sort of like so.

And you want, when you wait long enough, you want that, at t goes to infinity, where t equals 0 the current is still maximum, so it's still V over R, but when you go to infinity -- if you wait long enough -- then, of course, the current has to become 0.

And as the current dies out, heat is being produced in that resistor at a rate of I square R joules per second, and then there comes a time that the current becomes almost 0 and then the whole show is over.

And so we can also calculate the exact time behavior by going back to our differential equation and make V 0.

Where is that differential equation?

Is it hiding?

Oh, there it is.

So I solved this differential equation, but this is now 0.

And the solution to that differential equation is that I as a function of time is I max times e to the minus R over L times t, and that exactly has all the quantities that you want it to have, because notice that at t equals 0 -- when you put in t equals 0 -- the current is indeed maximum, and that's what you require.

That's the moment that you make that capital V 0, the current was still running.

But notice that when t goes to infinity -- if you wait long enough -- that indeed the current goes to 0.

And if you wait L over R seconds, then you are down to about 37% of your maximum current.

So if you now go to I, I re-define t equals 0 here, so if now I wait L over R seconds then this value here is about 37% of that value.

So you've lost 63%.

And so you see, this is the consequence of the fact that the circuit is capable of fighting its own magnetic flux that it is creating.

When the current was running happily here, with the battery in place, the current was, let's say, all the time I max -- at least very close to that value -- V over R.

And so all the time, there was heat produced in the resistor.

I square R joules per second.

Who was providing that, uh, energy?

Well, of course, the battery.

But now, when I take the battery out, there is still current running, and that means while the current is dying there is still heat produced in that resistor, and that heat slowly comes out until the current ultimately becomes 0.

Now where does that energy come from?

Well that energy must come from the magnetic field that is present in the solenoid, and this idea -- that we have energy that comes out in the form of heat, which really was there ear-, earlier in the form of a magnetic field -- allows us to evaluate what we call the magnetic energy field density.

Let me first calculate how much heat is produced as the current goes from a maximum value down to 0.

Hmmmm, I'll have to erase something.

I'll erase this part here.

So at any moment in time, the current is producing heat in the resistor, and so if I, if my voltage becomes 0 at time t equals 0, then this is the amount of heat, uh-uh, no square here.

I square R dt, integrated from 0 to infinity, is the total heat that is produced as the current dies out, but I know what this current was -- I just erased it, if I still remember it -- so I can bring I max outside and

I can bring the resistance outside and then I get the integral from 0 to infinity of e to the minus R over L times t dt.

And this is a trivial integral.

This integral is, uhm, L divided by 2 R.

Oh, by the way, it is I squared, so I have a 2 here.

It's very important.

Don't forget the 2.

And so that integral is L divided by 2R, and so if now I look at the product of I square maximum R L divided by 2 R, I get one half L times I maximum squared.

So this comes out in the form of heat, and I max is then the maximum current that we had when the current was flowing after a long time.

I can now by, by manipulating numbers, I can now calculate how much energy there was in that field per cubic meter, because the magnetic field was exclusively inside that solenoid.

And if I know that the energy that is ultimately coming out is one half LI squared, then I have here I, and so I can replace that I there by B divided by mu 0 times L divided by N and here I have L.

And so if I substitute in here the value for L that we have on the blackboard there, and we substitute for I the value that we have there -- you can drop the maximum now -- this simply tells you, then, that any moment in time that I have a current I running through a solenoid, that the energy that is available in the solenoid in the forms of magnetic energy is one half LI squared.

And so when you do that, you substitute capital L and capital I, you will find that one half LI squared then becomes B squared over 2 mu 0 times pi little r squared times l.

You check that at home.

It's simply a substitution.

But this is the volume of the solenoid where the magnetic field exists, and we have assumed that the magnetic field is 0 everywhere outside.

And if you accept that, then you see that we now have a result for the magnetic field energy density -- that is how much energy there is per cubic meter -- that is, of course, this value.

Because this is the total energy of the magnetic field, if we know the current, and this is the volume of the magnetic field.

So the magne-, magnetic field energy density is then B squared divided by 2 mu 0, and this is in joules per cubic meter.

So in principle, if you knew the magnetic field everywhere in space, then you can integrate over all space and you can then calculate how much energy is, uhm, present in the magnetic field.

And earlier in this course, we did something similar for electric fields.

We calculated the electric field energy density.

Perhaps you remember whatit was.

It was one half epsilon 0 kappa times E squared.

It was also in joules per cubic meter.

Now in the case of an electric field, this represents the work that I had to do to arrange the charges in a certain configuration.

In the case of a magnetic field, it represents the work that I have to do to get a current going inside a pure self-inductor.

That means the resistance of the self-inductor is 0, and it takes work because the solenoid will oppose the building up of the current, and so I have to do work.

So there's a parallel between the two.

I can make you see, in a quite dramatic way, how strong self-inductances can fight their own current, and the way I'm going to do that is with the setup there, whereby I have a twelve volt car battery and I have two light bulbs.

I have here an enormous self-inductance L, 30 Henry.

We will learn later in the course how you make such a high self-inductance, and then here is a light bulb.

The light bulb has a resistance of 6 ohm.

This self-inductance, there is nothing we can do about it.

It happens to have 4 ohm resistance.

We don't have a self, we don't have s-, superconducting wires here, so it also has a resistance of 4 ohms.

Forgive me for that but there is nothing I can do about it.

I have here another resistance of 4 ohms and a light bulb, which is the same one as that one -- also 6 ohms -- and then here is my car battery plus a switch.

And the car battery is twelve volts, and I'm going to throw the switch, turn it on.

You will see that this light bulb comes on almost instantaneously.

There is no self-inductance in this loop -- well, maybe a few microhenry or even less -- but in this loop here there is this huge self-inductance, and so the self-inductance says take it easy to the current, take it easy, just wait, and you will very slowly see that light bulb come on.

And we can calculate how long it takes, because we have here 10 ohms -- 6 ohms and 4 ohms, so L divided by R, we have 30 divided by 10 -- so that is 3 seconds.

So what that means is that even after 6 seconds, which is twice this time, even then the current through this light bulb is only 86% of its maximum.

But that means since the light, of course, is proportional to I square R in the light bulb, that the light is only 75% of maximum, and even if you wait 9 seconds, then the light that comes out here is still only 90% of it's maximum, whereas this one comes on immediately.

The reason why we put the 4 ohm here, we want this part to be -- from an ohmic resistance point of view -- to be identical as this part.

So that's why you artificially added here the 4 ohm, because the 4 ohm is always there in the self-inductance, there is nothing we can do about it.

And so you are going to see a remarkable example.

This is one light bulb.

That is the one that is here, and this is the light bulb that is there, and the self-inductance is in this incredible monster, here.

I'll make the lights, change the light setting a little bit so that we can see the lights of the bulb.

The bulbs are only 8-watt bulbs, they're not very, very strong, not very bright bulbs, but the effect will be very clear.

So if you're ready for this -- so this is the one that I think comes on immediately, and this is the one that takes its time -- three, two, one, zero.

One, two, three, four, five, six, you see how slow?

It's still not very bright, it's still getting brighter.

It's still not as bright as this one.

It's getting there, you can actually do some timing by counting.

By the way, the values of the resistance that I gave you are when the light bulbs are hot.

Three, two, one, zero, one, two, three, four, five, six, seven, eight, nine, ten, eleven, twelve, I'm still seeing it getting brighter.

What you're seeing here is a remarkable example that the self-inductance is fighting itself.

That's why the name self-inductance is so nice.

Now I want to go one step further and I want to power the LR circuit with a AC power supply.

If you have an AC power supply -- so it's changing all the time, the voltage -- now, of course, the self-inductance is fighting back all the time, not just only in the beginning as you saw in this circuit, but now, of course, it is active almost all the time.

So we can do away with this, and so now we replace the battery by a AC power supply, which we normally put just a wiggle there, and here is the self-inductance L, and here is the resistor R.

And let the voltage provided by this power supply be V 0 times the cosine of omega T, omega being the angular frequency.

And now, I have to apply Faraday's Law, not Kirchhoff’s Rule -- Faraday's Law -- when I go around this circuit and I set up the differential equation.

And of course the differential equation is going to be exactly like this, except that V, now, is V 0 times the cosine of omega T.

And now I have to solve for this differential equation.

That's the only difference, so I don't have to start from ground zero.

And the solution to this differential equation is quite remarkable and not so intuitive.

The current -- as a function of time, now -- is V 0 divided by the square root of R squared plus omega L squared times the cosine of omega t minus phi, and the tangent of phi, that angle phi, is omega L divided by R.

And of course we need some time to digest this.

The first thing that you notice is that there is a phase lag between the current and the driving voltage.

If phi, as you're going to see, is, uhm, 90 degrees, then the current is delayed by one quarter of a cycle, though the fact that the current comes later than the driving voltage perhaps is intuitive, because the self-inductance is fighting the change in the current.

So it's perhaps not so surprising that there's going to be a delay, that the current comes a little later.

If you look at this equation here, then what you have in front of the cosine term is obviously the maximum possible current, because the cosine term is just oscillating between plus 1 and minus 1, and so this here is the maximum current that you can ever get.

In one full cycle you get positive and you get negative net value.

And notice here the role of omega L really plays the role of a resistance, and in fact the dimension of omega L is ome-, is, is, is ohms.

It really plays the role of a resistance, and if omega is very high then the resistance here becomes very high and so your current becomes very low.

Well, that's intuitively pleasing because if omega is high, then the changes -- the dI/dt's -- are very, very high, and therefore if there are very fast changes the induced EMF is going to be high, and so the current will be low.

Also, if L is very high, then the system also is capable of fighting back very hard, and so it puts up a large resistance.

So it's also pleasing that you see the L there, downstairs.

If omega is very low, in your mind you can make omega 0.

You don't even have alternating current when omega is 0, then you have DC which is direct current.

So when you make omega 0, you simply get I is V 0 divided by R.

That's Ohm's Law, that's obvious that you get that.

Let's now look at the phase angle.

The tangent of phi is omega L divided by R.

If the self-inductor is very large, then the system has a strong ability to fight back, so it can delay that current by a large amount, and the same is true is omega is high.

If omega is high then the time changes are very, changes occur on a very small time scale, and so the system can fight back.

Because remember, you always have the EMF proportion to dI/dt.

And so it's also pleasing to see that omega and L are upstairs here.

Either one of being large it can fight back and it can hold back the, the current.

I have worked out a situation whereby we have an LR circuit -- this is on the Web, you can download that so you don't have to copy it -- and the reason why I have these values is because it is directly coupled to a demonstration that I will do shortly.

You see, an L in series with an R, the L is 10 millihenry and the R is 10 ohms, and let V 0 be 10 volts.

And here you see three frequencies, 100 hertz, 1000, and 10000 hertz, and here you see the values for omega.

You have to multiply hertz with 2 pi.

And look now at omega L.

At low frequency -- a hundred hertz -- omega L is 6.3 ohms.

Compare that with the 10 ohms.

They'd be comparable, but now look for instance at 10000 hertz.

The omega L is huge.

It's 630 ohms.

So it entirely determines -- so to speak -- the resistance of that circuit.

And so the current that is going to run -- at least this is the maximum current is this value, which we also saw here on the blackboard -- that current at high frequently is enormously reduced.

It's 50 times lower than this current at low frequency, even though they have the same value for V 0.

And then you see here the phase angles.

And the reason why you have these values is that I can make you listen to this, I can make you hear this, because your hearing is very good at 100 hertz, and since all of you are young you can probably hear even 10000 hertz, maybe some of you can even hear 20 kilohertz.

When you get older you lose your high frequencies.

In fact, my frequency cut off is somewhere near 4000 hertz.

I'm going to make you listen to music, and there will be violins which produce probably 3, 4, 5000 hertz, and then I will turn on, all of a sudden, the 10 millihenry.

So first I will make you listen to music whereby there is no 10 millihenry in there, and then I will turn on the 10 millihenry, and what you will hear, that the violins disappear because the current reduction is now huge on the high frequency but is very little on the low frequency, and that's the idea of what a self-inductor can do for you.

So if you listen to this [playback of classical music with violin] there's no self-inductor in now.

There's no music either.

[laughter] OK.

[Playback of classical music continues, without violin notes.] OK.

It's different music.

No self-inductor.

[Click] Self-inductor.

The high frequencies are gone.

[Playback of classical music continues, again with violin notes] No self-inductor.

[Click] You can turn a violin concerto into a cello concerto.

[laughter].

You just cut the violins out.

OK.

I cannot make you listen to the phase shift, not even in the case of the 90 degree phase shift, and that is quite obvious because what does it mean that there is a 90 degree phase shift?

It means that during one cycle of 10000 hertz -- which takes only 1/10000 of a second -- that the high frequencies are shifted by only 25 microseconds, and there's no way that your ears, your ears can hear that.

The fact that the composer wanted those violins to come in 25 microseconds earlier than, than they do of course is something you cannot hear.

So I cannot make you listen to the phase shift, but for the phase shift I have something else.

And for that something else, I'm going to return to the, to my last lecture, in which we levitated a woman -- magnetic levitation.

And so I'm going to return to that idea and grind a little deeper than we did when I gave that lecture, just before spring break.

I had a coil and I was driving that coal, coil with 60 hertz AC.

And let's assume that looking from above that the current was running in clockwise direction, which is exactly what I assumed when I discussed this with you, and so the magnetic field is coming down like this -- magnetic dipole field -- produced by this coil.

And then we had here, we had a conducting plate under there, and I said to you when this magnetic field is increasing in strength, then there's going to be an induced EMF here, which tries to imp-, oppose that change, and so the induced current that is going to run, which we call eddy current, is going to run in this direction.

If this is clockwise, this current is going to be counter clockwise, so it's going to produce a magnetic field in this direction.

It opposes the change of the increasing magnetic field: Lenz's Law.

And since the two currents are in opposite direction, the two repel each other.

And you bought off on that, and we levitated a woman.

However, no one asked me the question, what happens a little later in time when the magnetic field is still in the downward direction but it is decreasing, since it is an AC current, there comes a time that the magnetic field will be decreasing in time.

Now the EMF here must flip over because Lenz says sorry, we don't like the decrease.

The moment that the EMF flips over, this current will flip over.

The two currents are now in the same direction, and they will attract each other, and so there goes your magnetic levitation.

Half the time attraction, half the time they repel each other.

But yet we did elev- levitate a woman, and the secret lies in the self-inductance.

This current that runs here runs over a path -- which is very difficult for me to anticipate -- which has a certain resistance, R; and it has a certain self-inductance, L.

We know what omega is; that's about 360.

And so we do get in this conductor, we get the current, the induced current here, is delayed by a phase angle given by this equation, is delayed over the induced EMF.

The EMF immediately coupled to what this coil is doing, but the induced current is delayed.

And I have something that will allow that, allow you to see that, perhaps in even more detail.

This red curve is the current through the coil, the coil that you see there above.

And when the current is above the black line it's clockwise and when it is below the black line it's counter clockwise.

The vertical scales are arbitrary.

The green curve is the EMF, which is induced in the conductor.

Notice when the magnetic field increases, when the current goes up in the coil, that the EMF in the conductor is in such a direction that it opposes the change of that magnetic field.

But now when the magnetic field goes down, when the current in the coil goes down, immediately the EMF flips over, which is what I just mentioned to you, and therefore if the induced current and the induced EMF were in phase with each other, half the time you would have attraction and half the time you would have that the two repel each other, and that won't give you magnetic levitation.

Here, what you see is a blue curve which represents the induced current.

I call it the eddy current.

If there is no phase shift between the induced EMF and the induced current, notice that half the time the blue curve and the red curve are in opposite direction.

When they are in opposite direction they repel each other.

When they are in the same direction they attract each other.

But now, if I have a phase delay so that the induced current comes later than the EMF -- and I'm going to do something dramatic, I'm going to shift it by 90 degrees so the current is now 90 degrees delayed relative to the induced EMF -- look now that the red curve and the blue curve are always in opposite direction.

And so now there is hundred percent of the time a repelling force.

The coil repels the conductor and the conductor repels the coil.

Now in the case when we levitated the woman, I am sure that the phase delay was not 90 degrees but maybe it was only 30 or 40 degrees, but the net result is here the shift is not 90 degrees, the net result is that you get, on average, a repelling force.

And so, the secret of the repelling force, in the case of the levitation of this coil and therefore of the levitation of the woman, lies in the fact that there is a finite self-inductance [tapping noise] in here.

If R is 0, then of course we have a super conductor, then phi is always 90 degrees.

When R is 0, this is infinitely high, so now we get a 90 degree phase shift, and I did a demonstration whereby I had a little magnet floating above a superconductor.

That was an ideal case.

Phi was then 90 degrees, so they always repel each other.

Today I want to do a more controlled demonstration, whereby I can actually calculate the self-inductance and I also can calculate the resistance.

And what I will do today is I will have a coil which is stationary, and I will have a conductor which is not stationary.

Here is my coil.

AC, 60 hertz.

There's the coil.

And I have a ring, and the ring is made of aluminum, and I know exactly the dimensions of this ring.

I know the radius -- it's about 5 centimeters -- I know the thickness, I know everything.

It's an aluminum ring and has a radius of about 5 centimeters.

Since I know all the dimensions, I can calculate the resistance of that ring.

You should be able to do that, too, if I gave you the dimensions.

And so the resistance of that ring very roughly is about 7 times 10 to the minus 5 ohms.

It's very small.

I can all-, this is at room temperature, by the way, the lower temperature the resistance is lower.

I can also calculate very roughly what the self-inductance is of that ring.

Now that's not so easy, because here when I calculated the self-inductance, the magnetic field was constant.

I assumed it was constant, uniform inside.

That's not the case when you have a ring.

You have a dipole field.

However, I just assumed that the magnetic field was the same everywhere at the surface of the ring -- and with that assumption admittedly I could be off maybe by 20 or 30% -- with that assumption I find that the self-inductance is 10 to the minus 7 Henry.

I know what the omega is.

It's 360, roughly, and so I find that omega L over R for this ring is about one half, run at that frequency omega.

And that gives me a phase angle phi of 25 degrees, and therefore the ring is going to be repelled by the coil and of course the coil is going to be repelled by the ring.

I'll put the ring here, and the ring is supported by this, so this ring cannot fall over.

The only difference between this experiment and that one is first of all I can be very quantitative there, I can actually calculate the phase angle, where here that's almost impossible.

Here, it is the conductor that I'm going to make levitate and the coil is stationary.

In here it was the conductor that was stationary, and the coil is floating, but of course the idea is exactly the same.

And so what I want to do now is make you see it there, actually, hmmmm, and we'll have to, it should come up there.

Yes, there you see this, this ring.

Maybe I should first show you the whole setup.

So this ring goes over here -- this is an aluminum ring -- and I'm going to make it levitate by simply running 60 hertz, 110 volts through this coil and I hope I do nothing wrong.

Oh no, I have to turn on my AC.

Oh my God, that was not, gee, that was not my intention.

A good thing we don't have a woman sitting on the ring now.

My idea was to have it levitate.

By the way, you did see that it was repelled, that was quite clear.

I had the current too high.

[laughter] I had the current too high so we'll have it a little lower, and I will make the current come up very slowly and then I want you to see that it levitates.

There it is, levitating.

Oh, oh, off the screen.

There it is.

And I can turn it over and it's still levitating, of course.

And the secret is this phase delay introduced by the self-inductance.

I have another ring here which has a, a slot -- also aluminum -- same ring, but it has a slot.

Well, the EMF in this ring is going to be identical.

There's no difference.

In fact, the self-inductance of this ring is identical, but the resistance of this ring is huge because there's a slot in there and the resistance is almost infinitely high.

And so if the resistance is infinitely high, no matter what L is and what omega is, phi is going to be 0.

It won't repel.

Half the time it attracts, half the time it repels.

[metallic clanking sound] That means nothing happens, no magnetic levitation.

Same EMF, same self-inductance, but an infinite resistance and here you see magnetic levitation.

Since the induced current in the ring is extremely small because of its very high resistance, the force on the ring -- whether it's repelling or attracting -- in any case is practically 0.

So that alone is enough reason for the ring not to move at all.

All right, I hope to see all of you tomorrow during our exciting testing of the motors.