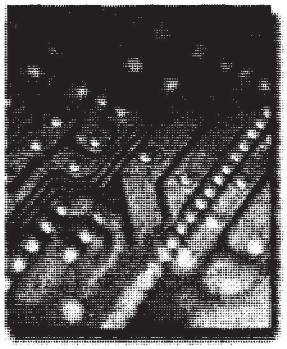
- Don't allow nonrelated plane areas to overlap.
- Connect reference planes together at only a single point.
- Don't route signal traces across a split or across an unrelated plane.



17 LOSSY LINES AND EYE DIAGRAMS

In days gone by we considered traces on PCBs to be (almost) perfect conductors. On rare occasions we considered their resistance and maybe propagation time, but usually not. Traces almost never contributed to signal integrity problems.

Then, as rise times shortened, we began to worry about signal reflections. In Chapter 10 we learned that signal integrity problems caused by reflections could be managed if we designed our traces to look like transmission lines and if we addressed questions like trace termination. Instead of traces resembling perfect conductors, traces resembled distributed networks of LC components. But if we could measure an effective input impedance whose value is given by Equation 10-1

$$Z=\sqrt{L/C}$$

we could terminate traces in components with the same impedance and still control the reflections.

However, rise times continue to shorten and now new issues are coming into play. Their solution will not be the purview of PCB designers, but designers will need to understand the issues if they are going to continue to grow in their profession.

LOSSY LINES

A transmission line that we consider to be infinitely long and whose input impedance is given by Equation 10-1 is considered to be lossless. That is, whatever signal goes into the front of it arrives at the far end without any distortion. This is not a bad model for well-designed lines at lower frequencies. Although nothing can

INTEL-1010
U.S. Patent 7,151,430

be truly "lossless," lines can be designed that are generally close enough for most purposes.

But there *are* losses, and as frequencies (or frequency harmonics) get high enough, these losses can become significant and can begin having negative effects on our boards. There are two general areas of losses we need to be concerned about. Both are beyond the scope of this book, so they are simply summarized here so we can illustrate some of their effects. That way, if your circuit design engineers begin setting board design parameters or designing complex termination strategies you will have some understanding of what they are doing and why they are doing it.

Skin Effect Perhaps you have heard about something called the *skin effect.* This describes the tendency of the current to travel along the outside skin of the trace or wire instead of being evenly distributed throughout the entire cross-section of the wire.

The explanation for why this is so is complicated and subtle. If you have read and understood Appendix B you learned that when a current changes rapidly in a wire it (a) generates a changing magnetic field around the wire, which (b) generates a current in the opposite direction in the wire.

This opposing current is what causes the "inertia" we talk about when we describe inductance. At the very first instant, the reverse-generated current is the same as the forward-generating current, and no current flows. But soon the forward current overcomes the reverse-generated current and it does begin to flow. The resistance to the first increments of current creates what we call inductance.

Now consider a wire with a circular cross-section, as shown in Figure 17-1. Assume a current flows evenly through that cross-section and that the current is increasing. The increasing current generates an increasing magnetic field that in turn generates a current back in the opposite direction. Where is that current concentrated?

The strength of the magnetic field is inversely proportional to the square of the distance. Consider a point along the centerline of the wire. All other points that make up the cross-section of the wire are within one radius of the centerline. Now consider another point along the circumference of the wire. More than half the other points that make up the cross-section of the wire are further than one radius away from this point. Therefore the effect of the magnetic field must be *weaker at the circumference* than it is along the centerline.

INTEL-1010
U.S. Patent 7,151,430

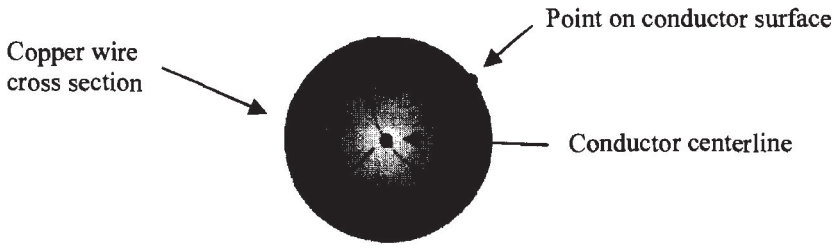


Figure 17-1 Very high-frequency currents will travel on the outside skin of a conductor where the inductive coupling is less (shading suggests current density).

This then means that the *very first* place a sharply increasing current will begin to flow is along the circumference (surface) of the wire, rather than at the centerline (or uniformly across the entire cross-section). At lower frequencies this effect is very transient. Current will quickly spread to the entire cross-section. But consider what happens when the frequency is *very* high and current changes direction *very* quickly. Just when the magnetic field (self-inductance or inertia) is beginning to be overcome, the driving current changes direction. Now we start over again and current begins first to flow on the outer circumference in the opposite direction. Before that current distribution can spread throughout the entire cross-section of the wire, the driving current changes direction again. In this way, the current distribution is always strongest at the outer circumference. In the extreme, when the frequency is high enough, the current distribution never spreads uniformly to the center of the conductor. The higher the frequency, the greater this effect is, and current only flows along a thin portion of the circumference, commonly called the skin of the wire.

Many different applications capitalize on this effect. Transoceanic communication cables, for example, are often made of strong but not necessarily highly conductive steel for strength. The steel cables have a very thin copper (or sometimes even gold) foil wrapped around the circumference. The thin foil is where the current flows. Overall conductivity is not affected because the high-frequency signals are only flowing along the skin anyway.

The skin effect manifests itself as a higher resistance associated with the wire or trace, and the effect increases with frequency. The effective skin depth is

inversely related to the square of the frequency. Thus, this is one place where the effective resistance becomes a function of frequency. Since only a fraction of the total cross-section of the wire or trace carries current, and since resistance is inversely related to cross-sectional area, the effective resistance the signal sees increases with skin effect and therefore with frequency.

Dielectric Absorption We think of dielectric materials as insulators. FR4, for example, is fiberglass. We think of it as passing no current at all on our circuit boards. As very high-frequency currents flow on our traces, however, the very fast-changing currents cause motion at the molecular level of the material. That movement requires energy, which is absorbed from the signal. That, in turn, attenuates the signal. The signal level is reduced, therefore, as energy is absorbed in the dielectric material.

The net effect is as if there were a resistor across the signal leads loading down the signal. The effect is a function of frequency. So again, we have a situation where losses appear as if we have a frequency-dependent resistor across the signal leads.

LOSSY LINE MODEL

In Chapter 10 we illustrated the model of a (lossless) transmission line, shown again here as Figure 17-2:

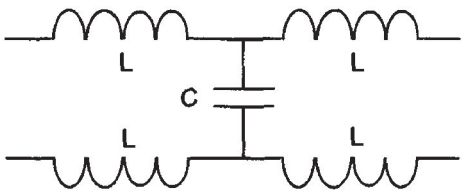


Figure 17-2 A lossless transmission line model.

When we add the "lossy" components to the model it becomes Figure 17-3.

INTEL-1010
U.S. Patent 7,151,430

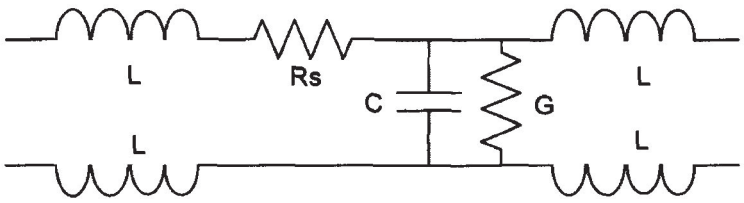


Figure 17-3 A lossy transmission line model.

In Figure 17-3, the Rs represents the (frequency-dependent) skin effect losses and G represents the (frequency-dependent) conductance associated with the dielectric absorption.[1] We use conductance rather than resistance in the model because it simplifies some of the arithmetic. As frequency increases both Rs and G increase. Thus, the higher frequency harmonics of the signal are attenuated.

One significant consequence of losses in the line is that the characteristic impedance of the line changes with frequency. We noted in Chapter 10 that the characteristic impedance of a lossless transmission line is given by the formula in Equation 10-1:

$$Zo = \sqrt{\frac{L}{C}}$$

The characteristic impedance of a lossy line is given by the formula shown in Equation 17-1:

$$Zo = \sqrt{\frac{Rs + j\omega L}{G + j\omega C}} \qquad [17\text{-}1]$$

The frequency range at which lossy lines become an issue is, of course, dependent on many factors. In general, loss effects can become measurable in the 100 MHz range and can become design issues in the 1.0 GHz range.

[1]Recall from Chapter 5 that conductance, G, is the inverse of resistance, R.

EYE DIAGRAMS

Eye diagrams are an effective way to visualize the effects caused by lossy lines. But first we need to know what eye diagrams are and how they are created. Start with Figure 17-4.



Figure 17-4 Typical data string.

Figure 17-4 illustrates a typical data string of bits with value 1 (high) or 0 (low). If we were looking at this on an oscilloscope screen we would see the right-hand bit first. Then it would shift to the right (in time) and we would see the next bit. So if we were going to read the data string as it passed by a point in our circuit we would read (from right to left) 0-0-1-1-0-1-0-0, and so on.

More generally, if we looked at a specific point in a circuit and watched a large number of data streams passing by that point, and looked only at the three most recent bits, there are only eight possible combinations of bits that would pass by. Figure 17-5 identifies them. In Figure 17-5 the numerical codes represent the data bits as they move past a point in the direction shown. Thus we "see" them in order from right to left.

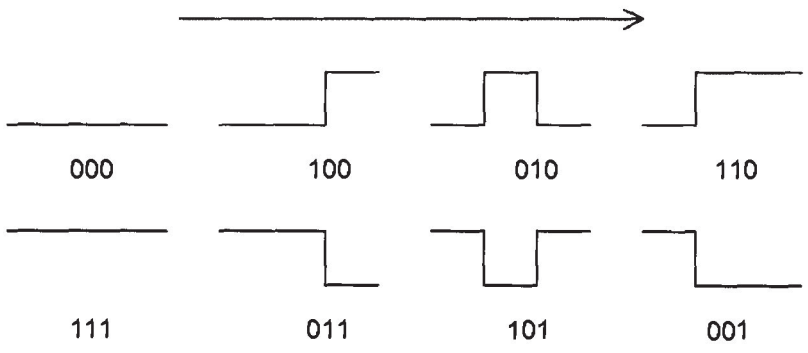


Figure 17-5 Eight possible combinations of a three-bit stream as they pass by a specific point in our circuit. (Read data stream from right to left.)

INTEL-1010
U.S. Patent 7,151,430

Now if we took a time exposure of a large number of data bits passing a point, and superimposed them on a single graph, they would look like Figure 17-6.

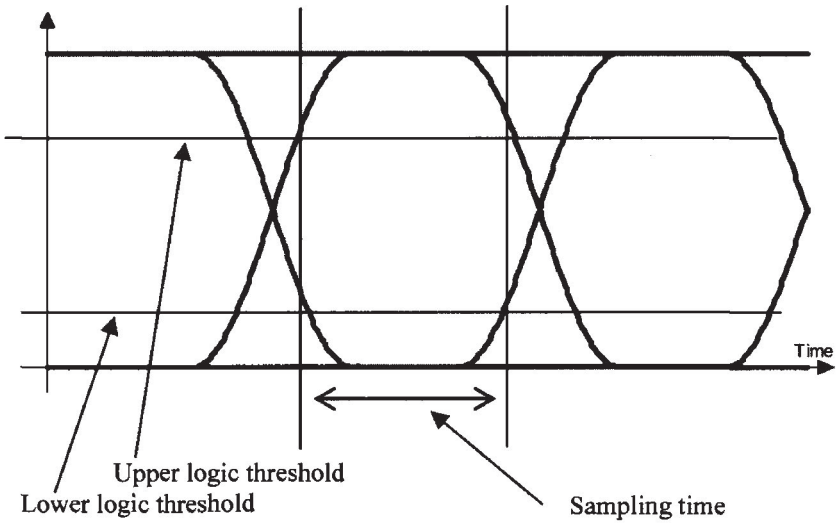


Figure 17-6 All possible combinations, superimposed, of three data bits passing by a point in a circuit.

Figure 17-6 represents an ideal eye diagram. It shows all possible combinations of three data bits passing by a point in a circuit. The transition between logic states is very clean and predictable. The upper and lower logic threshold levels are indicated in the figure. The vertical lines represent the sampling time during which there will be no logic errors. That is, if the data stream is sampled between these times, the data will always be in a correct and valid state. This is referred to as an eye diagram because of the obvious resemblance of the open area to an eye.

When there are losses along a line, the net effect is to attenuate (reduce in amplitude) the higher frequency harmonics. This can cause distortion in the signal. Depending on circumstances, the distortion can range from relatively minor to crippling, causing the system to simply not function. HyperLynx introduced a new version of its LineSim software package in 2003 that allows simulation of lossy lines and eye diagrams. Figures 17-7 and 17-8 illustrate a simulation of a signal at the receiving end of a 20-in. transmission line under ideal and lossy assumptions. This

entire area is so complex that I offer these simulations only as indications of how lossy lines may behave, not as suggestions regarding what to do about them.

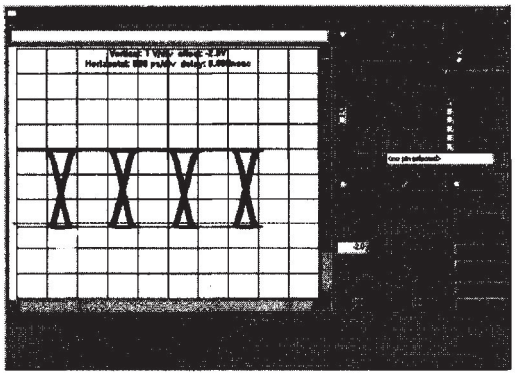


Figure 17-7 Eye diagram simulation for a properly terminated 20-in. transmission line.

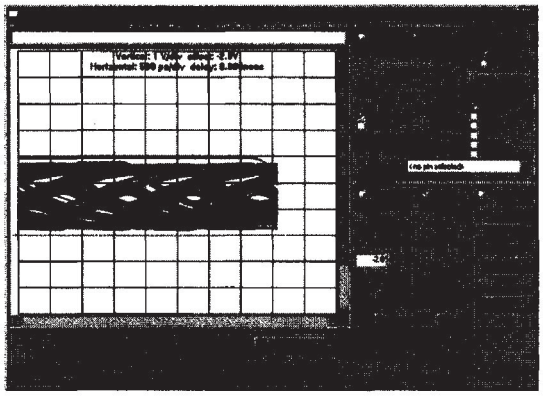


Figure 17-8 Eye diagram for the simulation in Figure 17-7 with significant lossy assumptions.

EQUALIZATION

If the lossy characteristic of a transmission line means that higher-frequency harmonics are attenuated, then we may be able to compensate for that with high-frequency equalization. By that we mean adding active circuits that pre-

amplify higher frequency harmonics or passive circuits that act as high-pass filters, attenuating the lower frequency harmonics. Much research is under way at the present time regarding how to do that.

Passive Equalization Passive high-pass filters are a relatively simple way to try to compensate for losses along a line. They consist of simple parallel RC circuits. As frequencies increase, the current passes through an RC equalizing filter with little attenuation, but lower frequency harmonics are attenuated somewhat. Since this is exactly the opposite of what happens along the lossy line, a properly constructed equalizing filter can (at least in theory) exactly compensate for the losses in the line.

Such RC filters can be placed at the beginning of the line (much like a series terminating resistor is), at the end of the line, or can even be built into the receiving device. Figure 17-9 shows what happens with the simulation of Figure 17-8 when a simple parallel RC filter circuit is placed at the beginning of the line. The compensation is not perfect, but the improvement in the signal is evident.

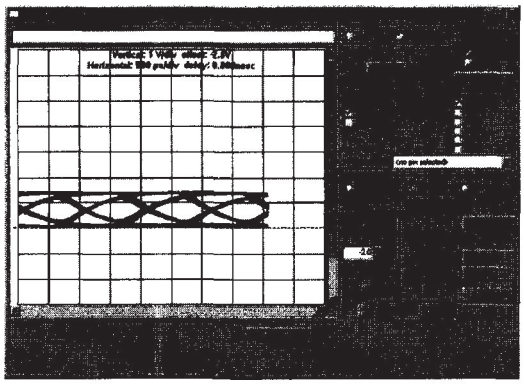


Figure 17-9 RC equalization at the driving end of the line can improve the eye diagram for the simulation in Figure 17-8.

Active Equalization Suppose we have a bit string with a pattern that looks like 0-1-0-1. Here there is a rapid transition from a logical zero to a logical one and back to a logical zero. This represents a very high-frequency sequence traveling down the line, and the signal transitions may be attenuated if the line is lossy. Compare that to a string like 0-1-1-1. There is a rapid transition from the first logical zero to the first logical one, but then the signal stays at the logical one level. This represents a higher frequency transition followed by a lower frequency stabilization.

Let's define the following strategy:

1. Assume the signal transitions between zero and one.
2. Every time a signal transitions from one state to another, we boost its magnitude by .5.
3. If the next clock pulse involves no transition, we boost the amplitude of the signal by .25.
4. After that we will not alter the signal until there is another transition.

This kind of equalization (compensation) can be done in the circuit at the chip level. If done that way, it is called *active precomp*. Figure 17-10a illustrates a bit string without compensation and Figure 17-10b illustrates the same bit string with active precomp. Notice how the active precomp acts to increase the higher frequency parts of the bit string, compensating for the later attenuation that will occur as the signal passes down the line.

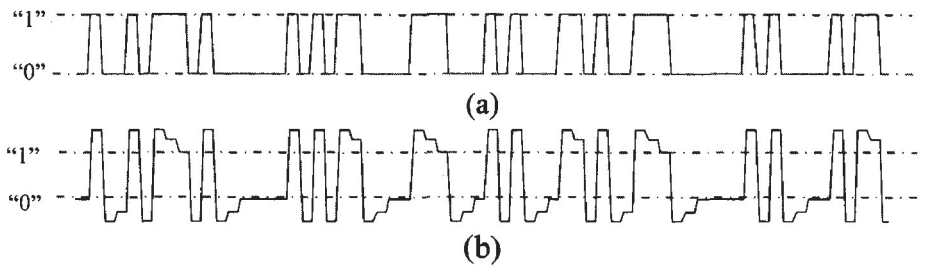


Figure 17-10 Bit streams (a) without any compensation and (b) with high-frequency precompensation. Note that rapid transitions in the signal are amplified to compensate for subsequent losses in the transmission line.

SUMMARY

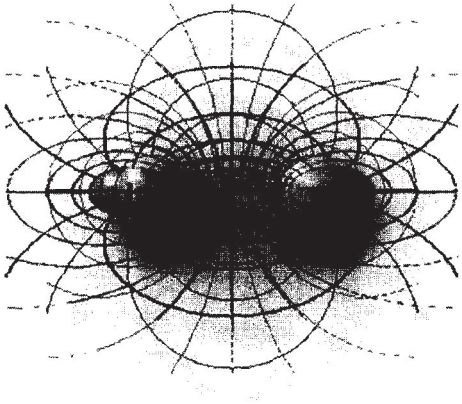
Lossy line analysis and equalization is going to be a rapidly growing area in the next several years. It will impact device manufacturers, many of whom are already experimenting with or actually designing in and supplying equalization cir-

cuits in their products. Passive circuit modules are appearing for both precomp and postcomp placement along lossy lines. Simulation tools will be adding lossy line simulation capabilities to their arsenals.

All of this means that board designers will have to become more conversant with the topic. Board designers will not be responsible themselves for designing equalization circuits for the applications, but they will have to understand them. As we learn more, there are certain to be placement and design rules for equalization circuits just as there now are placement and design rules for transmission line termination circuits.

INTEL-1010
U.S. Patent 7,151,430

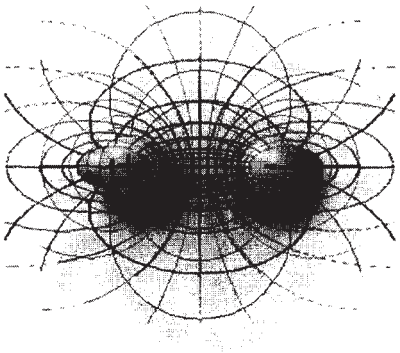
PART 3



APPENDICES, GLOSSARY, and INDEX

If I covered everything that interested me about the subject of signal integrity and circuit boards, I would never have been able to finish this book. The appendices included in Part 3 assemble several different types of information helpful to the reader. Included here is information regarding how to obtain demonstration files and freeware calculators; tutorials on why inductors induct, on logarithms, and on complex algebra; and a description on how TDRs and VNAs work and why they are useful to us. The last appendix, J, discusses one of my favorite topics—right-angle corners.

INTEL-1010
U.S. Patent 7,151,430



A UltraCAD'S SQUARE WAVE SIMULATOR

This simulation (see Figure A-1) is referenced in Chapter 1. Download the file *www.ultracad.com/prenticehall/square.zip* to a convenient place on your hard drive and unzip it. Execute the resulting square.exe file.

A square wave does not exist naturally in nature. We can think of a square wave as actually being formed by, or made up of, a series of natural cosine waves following the relationship:

$$\cos(\omega t) - \frac{\cos(3\omega t)}{3} + \frac{\cos(5\omega t)}{5} - \frac{\cos(7\omega t)}{7} + \ldots \text{ etc.}$$

where ω is the fundamental frequency of the wave and t is time. The subsequent terms in the relationship contain harmonics. That is, if the fundamental frequency, ω, equals 1.0 kHz (1,000 Hz), then 3.0 kHz (3,000 Hz) is the third harmonic of the fundamental. This is important because if our circuit is going to faithfully deal with the shape of the signal (say a square clock pulse going down a trace) it must be able to faithfully deal with all the different harmonics that are contained within it. This can actually be a very large number.

INTEL-1010
U.S. Patent 7,151,430

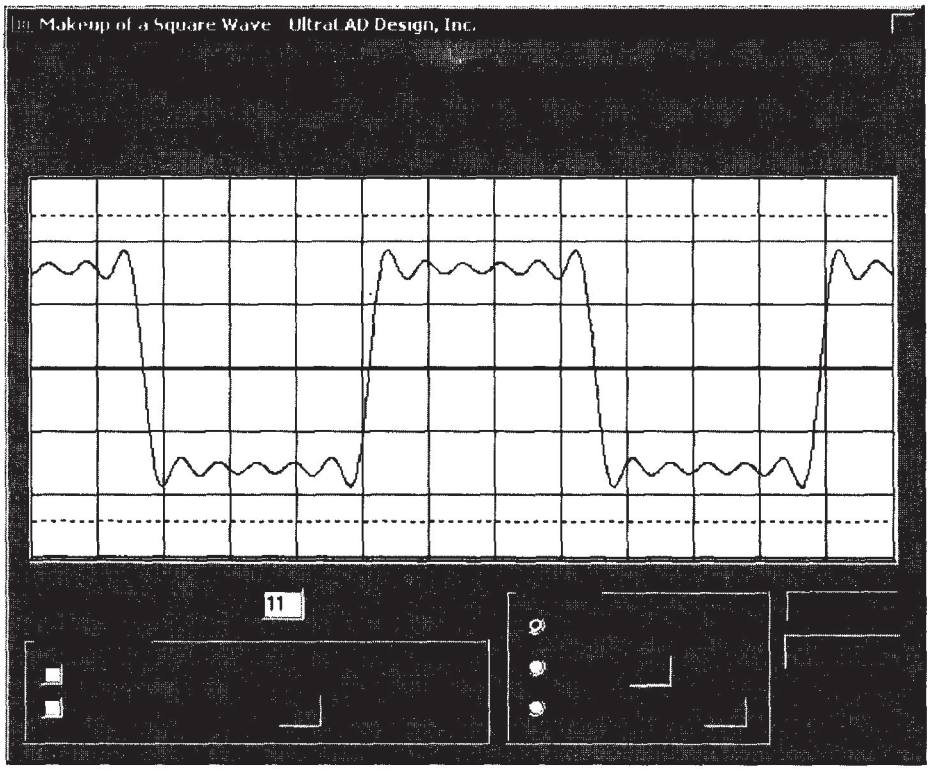


Figure A-1 UltraCAD's Square Wave Simulator.

This simulation can suggest what happens if the circuit does not pass all of the harmonics equally well, and helps you visually understand the makeup of a square wave. Start by entering a number for the "Include all harmonics through:" box. The number 11 is a reasonable place to start. Then press Display. The graph shows the fundamental frequency and the first eleven harmonics following that. (Note that only the odd harmonics enter into this relationship, so the first eleven harmonics only add five harmonics above the fundamental to the equation.) The curve looks like a square wave, but it is not perfect.

Change the number of harmonics to see the effect on the shape of the curve. Determine what you feel is an acceptable approximation to a square wave and see how many harmonics are included. It is that many harmonics the system must be able to deal with, without distortion, to produce the curve you think is acceptable.

This simulation allows you to experiment in some other ways, too. For example, select some number of harmonics and look at the waveform. Then select the

INTEL-1010
U.S. Patent 7,151,430

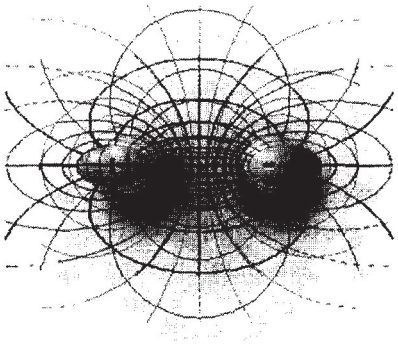
Exclude button and enter the number of a harmonic to exclude from the equation. Then select Display. Note how the waveform distorts. This type of thing could happen, for example, if somehow a filter existed in the circuit at the frequency you just selected that acted to "kill" that particular harmonic. That filter could happen inadvertently, or it could exist as a result of parasitics that no one knew about or expected. This is an insidious problem that can be very difficult to diagnose. (Note that if you exclude an even harmonic nothing happens because even harmonics are not contained in the series relationship.)

If you select the Exclude all except option, you can look at the relative magnitude of any individual harmonic.

The Expand Scale option allows you to see a close-up view of the rise time of the waveform. It expands the display in the approximate area of the fifth vertical grid line. This allows you to see a close-up view of any distortion or ringing associated with the rising edge.

The Shift higher frequencies by option will add an increasing phase shift to each harmonic. Suppose the circuit has some capacitive loading. The capacitive loading can cause a negative voltage phase shift (voltage lags current). If you enter a negative number in the box, each harmonic will be shifted by a multiple of this phase shift. Try these numbers: Enter 99 for the number of harmonics to include and –4 for the shift. Then select Display. The display looks almost exactly like the ringing that is common in many circuits.

Caution: The Shift higher frequency by option does not really reflect real-world conditions. I have included it to *suggest* some effects, but a real-world phase shift would be much more complicated than this. Treat this particular feature as interesting, but not necessarily representative.



B WHY INDUCTORS INDUCT

ELECTROMAGNETIC CONCEPTS

There were four people who were particularly important in the development of electromagnetic theory. Their groundwork was crucial to the ultimate understanding of this complex topic.

First, Charles Augustin de Coulomb has been credited with Coulomb's Law (1785) in which he stated that there are two kinds of charge, positive and negative. Like charges repel, and unlike charges attract, with forces proportional to the product of their charge and inversely proportional to the square of their distance.

Carl Friedrich Gauss is generally credited with stating that every magnetic pole is actually part of a dipole that includes an equal and opposite pole. One pole cannot exist without its opposite. Magnetic force is applied along a vector with a direction that is a line along which the force acts. Magnetic force is inversely proportional to the square of the distance.

André Marie Ampere is credited with Ampere's Law: An electrical current is accompanied by a magnetic field with a direction at right angles to the direction of the current flow.

Finally, in 1831, Michael Faraday published Faraday's Law of Magnetic Induction: A *changing* magnetic field is accompanied by an electric field that is at right angles to the change of the magnetic field.

These people, over a 50-year period, established some very important groundwork, but they were all (particularly Faraday) experimental scientists, not

INTEL-1010
U.S. Patent 7,151,430

great mathematicians. It was James Clerk Maxwell who combined these important observations and laws into a single set of equations that comprised a closed system. By that, we mean that every phenomenon and every relationship is self-contained in the set of equations, and any change in any one variable will flow through to every other variable in the system. There are no outside variables or forces needed or allowed.

This set of equations has become known as Maxwell's Equations, and they were published in his famous Treatise on Electricity and Magnetism in 1873. Maxwell is often treated as though he developed all the theory behind the equations, which is not true. For almost 100 years before him, people were developing significant parts of what would become known as electromagnetic theory. Maxwell's considerable contribution was seeing how to consolidate the work that several others had done earlier into a single, closed set of equations. That was no small feat, and Maxwell's Equations are still, today, crucial to our work and understanding in this area.

If you are wondering why I am going to spend a little time on electromagnetic theory, let me give you two reasons. First, it is at the heart of inductance. Inductance (and related issues like crosstalk) is an electromagnetic effect. Second, the first two letters of EMI stand for "electromagnetic." We need an understanding of this theory to help us understand destructive EMI effects.

Right-Hand Rule

Let's begin our discussion by looking at Ampere's Law. If a current flows through a wire, there will be a magnetic field around the wire (Figure B-1). The strength of the magnetic field is directly related to the amount of current that flows. Small currents have weaker magnetic fields; larger currents have stronger magnetic fields. We can find the direction of this magnetic field by the right-hand rule. Point the thumb of your right hand in the direction of the current flow. The magnetic field force curls around the wire in the direction in which your fingers curl.

INTEL-1010
U.S. Patent 7,151,430



Figure B-1 Right-hand rule.

This effect can be easily demonstrated at home using the famous compass experiment (Figure B-2). Connect a wire to a battery and place it under a compass so that the current flows toward the North Pole. Point your thumb in the direction in which the current flows (north). Notice how your fingers curl around the wire. The magnetic force around the wire, and above the wire (closest to the compass), will be in an easterly direction. When the wire is connected to the battery, the compass needle will swing toward the east, aligning with the magnetic force around the wire caused by the current flow. Disconnect the wire from the battery and the compass needle will point back toward the north.

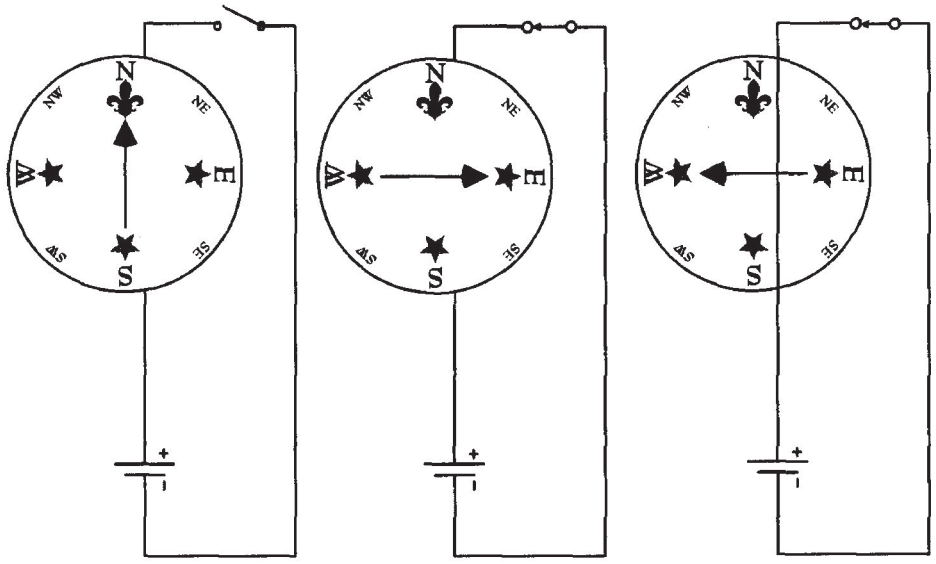


Figure B-2 The classic compass experiment.

If you now place the wire above the compass, notice how your fingers below the wire (closest to the compass) curl toward the west. When the wire is connected to the battery, the compass needle will swing to the west, aligning now with the magnetic force under the wire. Disconnect the wire from the battery again and the compass needle will swing back toward the north.

This simple experiment confirms the absolute relationship that exists between electric current and magnetic fields. Remember that current is the flow of electrons, and it is the magnetic property associated with electrons that is at the heart of this relationship.

Coil

We can strengthen and focus the magnetic lines of force if we coil the wire (Figure B-3). Note that the current through the wire in Figure B-3 is flowing from bottom to top. If, at any point along the wire, we point the thumb of our right hand in the direction of the current flow, our fingers curl upward inside the coil and downward outside the coil. The magnetic lines of force tend to be focused inside the coil, especially if the coils are close together. Since our fingers curl upward inside the coil, the north pole of the magnet will be at the top of the coil.

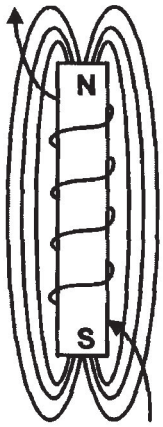


Figure B-3 Wire coils help "focus" the magnetic field.

INTEL-1010
U.S. Patent 7,151,430

The strength of the magnet will be increased if we (a) increase the current, (b) add more turns to the coil, or (c) add a ferrite core inside the coils to further focus the magnetic force.

SIMPLE MOTOR

In Figure B-3 we have taken a coil and passed a current through it to create an electromagnet. If we now place that coil between two other magnets, as shown in Figure B-4, we have the makings of a simple motor. The north pole of our coil will be repelled by the magnet with the north pole facing it, and it will be attracted to the magnet with the south pole facing it. Thus, there will be (in Figure B-4) a clockwise torque acting in the coil.

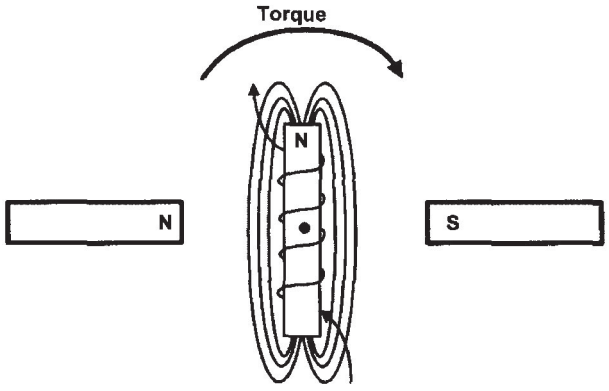


Figure B-4 Simple motor.

This is, in fact, the theory and design behind a basic electric motor. Real motors are more complicated in their design, of course. There are more complex mechanisms for getting current to the coil and for creating the internal magnetic fields, but the basic principle behind any electrical motor is as shown in Figure B-4.

If you look closely at the heart of the meter movement for a basic analog electrical meter, it will look much like Figure B-5. Current through the coil creates a torque that tries to rotate the coil in a clockwise direction. A spring pulls the coil back. The degree of rotation will be the point where the force created by the magnetic field just balances the restraining force of the spring. Look closely at the meter

INTEL-1010
U.S. Patent 7,151,430

movement of a typical analog Volt-Ohm meter (sometimes called a volt-ohm-milliammeter, or simply VOM) and you can see the magnet assembly, the coil, and the restraining spring.

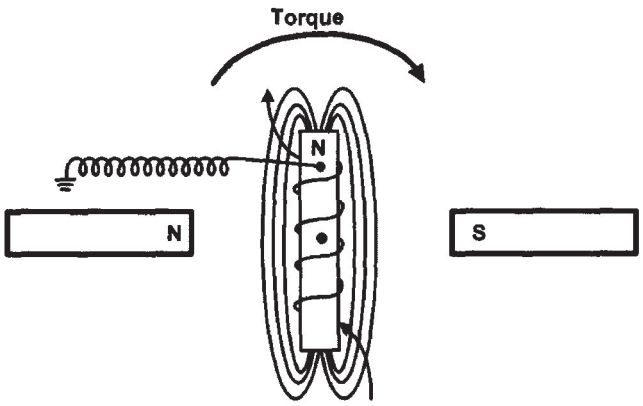


Figure B-5 Fundamental concept of an electric analog meter.

SIMPLE GENERATOR

Suppose we now take our coil and place another magnet beside it, as shown in Figure B-6. If the magnet is not moving, nothing much happens. Faraday's Law states that a changing magnetic field is accompanied by an electrical field, but says nothing about a stationary magnetic field. If we move the magnet toward the coil, however, the magnetic lines of force around the magnet cut across the coil's wires as the magnet moves. Thus, the wires see a changing magnetic field, and an electrical current is induced in the coil. The magnitude of the induced current is related to (a) the strength of the magnetic field, and (b) how fast the field is changing. Faster changing, stronger fields induce stronger electrical currents.

There are some important aspects here. First and foremost, the magnetic field must be changing. This concept is often difficult for people to grasp when they are studying this for the first time. If the magnet is a bar (iron) magnet, changing magnetic fields are created through motion (as suggested in Figure B-6). A changing magnet field can also be created by passing a changing current though a coil. Either way, something must be *changing* to induce an electrical current in a nearby coil. If everything is stationary, no current will be generated.

INTEL-1010
U.S. Patent 7,151,430

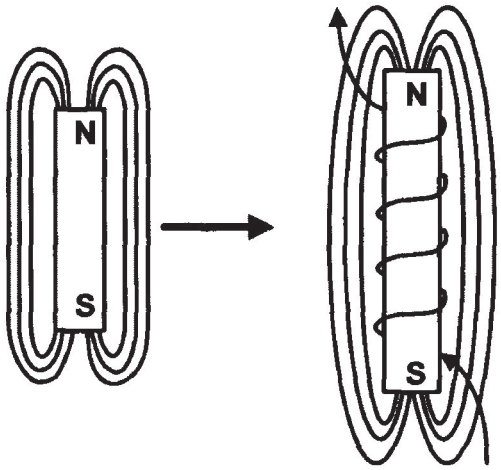


Figure B-6 A moving magnet can generate a current in an adjacent coil.

Looking ahead for a moment, one way to generate a changing magnetic field is with a changing current. The *strength* of the magnetic field is directly related to the magnitude of the current. The *rate of change* of the magnetic field is directly related to how fast the current changes. One of the biggest culprits in high-speed design signal integrity problems is di/dt—how fast current changes per unit time. Here is di/dt in the context of a changing magnet field. A changing current, di/dt, causes a changing magnetic field that can generate an unwanted noise signal in an adjacent wire.

Referring back to Figure B-6, it can be instructive to ask which way the induced current will flow. Is it going to flow upward through the coil, or downward? The answer can be tricky. The induced current flowing through the coil generates its own magnetic field. According to Lenz's Law (Heinrich Friedrich Emil Lenz), what is determined is the induced magnetic field. The induced magnetic field will be generated so that it will oppose the changing magnetic field generating the electrical current in the first place. We can then determine the direction of the induced current through the right-hand rule. Curl your fingers in the direction in which the induced magnetic field must be oriented; your thumb will then point in the direction of the induced current.

In Figure B-6, the moving magnet has a north pole at the top and it is moving toward the coil. The induced magnetic field in the coil will be generated with a

INTEL-1010
U.S. Patent 7,151,430

north pole at the top of the coil to repel the approaching magnet. Thus the electrical current will flow in whatever direction is necessary to generate a magnetic north pole at the top of the coil. In Figure B-6 this direction is from bottom to top in the coil.

In Figure B-7 the coil is wound in the opposite direction. The induced north pole still must be at the top of the coil, so in this case the induced current must flow from top to bottom in the coil to generate a north pole at the top.

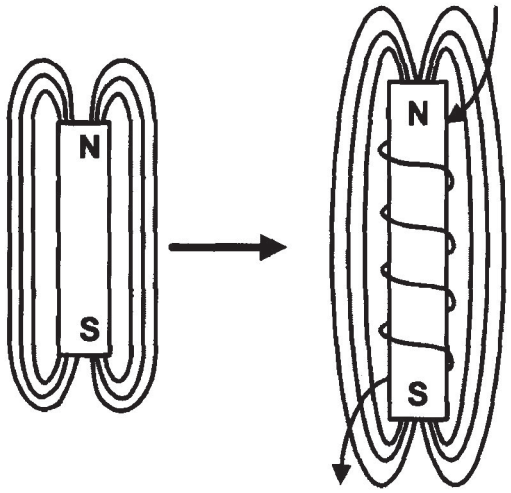


Figure B-7 The direction in which the current flows depends on how the wire is coiled.

If the magnet moves away from the coil, the situation is exactly reversed. The changing magnet field induces a current in the coil, but this time the induced current flows in the opposite direction. That is so the top of the coil becomes a south pole and attracts the moving magnet, trying to stop its motion.

Thus, there is no easy rule to help define in which direction an induced current will flow. The basic rule is that the induced current will flow in a direction so that the magnetic field it generates will counteract the changing magnetic field that induces it in the first place. The actual direction depends on how the coil is wound.

INTEL-1010
U.S. Patent 7,151,430

When you think about it, the fact that the induced magnetic field counteracts the inducing field must be so. If it were the other way around, everything would be reinforcing. A moving magnet would generate a current to attract the magnet faster, which would induce a larger current, and so on. A runaway situation would result. We know intuitively that this doesn't happen. Induced currents always act to counteract the forces inducing them.

Figures B-6 and B-7 illustrate the basic principles behind a generator. As with motors, the real-world designs are more complicated for reasons of efficiency and manufacturability, but all electrical generators are based on this simple basic concept.

Figure B-4 illustrates a basic motor. But what if we mechanically turned the coil, instead? It would then see a changing magnetic field caused by its motion within and relative to the magnetic field from the stationary magnets. Thus, if we mechanically rotated the coil in Figure B-4, a current would be induced (generated) in its coil. In that case, Figure B-4 would illustrate a generator.

If Figure B-4 can illustrate both a motor and a generator, what is the difference between them? The difference between a motor and a generator depends on where and how energy is applied. If electrical energy is applied and motion is generated, we have a motor. If motion is applied and electrical current is generated, we have a generator. The same device can be both a motor and a generator, depending on how it is used.

You can demonstrate this for yourself by going to an electronics store and buying a simple 3-volt DC motor. Hooking it up to a battery will confirm that it is a motor. Then connect a sensitive VOM (ammeter) to the terminals of the motor and turn the shaft with your fingers. (An analog meter works best, because the generated signal only exists for the short moment that you turn the shaft. Remember, only *changing* magnetic fields generate an electrical force.) You will see the ammeter needle jump momentarily as your fingers turn the shaft, confirming that you are generating a signal.

A car has an electrical starter motor and also a separate electrical generator (alternator). These are different devices (not a single device used for both purposes) because the demands on them are extreme and they each need to be optimized for their specific use. A gasoline-powered golf cart, however, usually has a single device that acts as both a motor and a generator. When you push the gas pedal down, the circuit closes and the battery turns the device (motor) to start the engine.

INTEL-1010
U.S. Patent 7,151,430

When the engine starts, it turns the device (generator) faster and generates enough current to recharge the battery connected to it. This is a practical example of a single device that operates as both a motor and a generator in the same application.

INDUCTANCE

Now, let's look at the situation in Figure B-7 but skip the magnet; we'll use another coil, instead (Figure B-8). If we apply a changing current through the coil on the left, it will generate a changing magnetic field that will induce a current in the coil on the right.

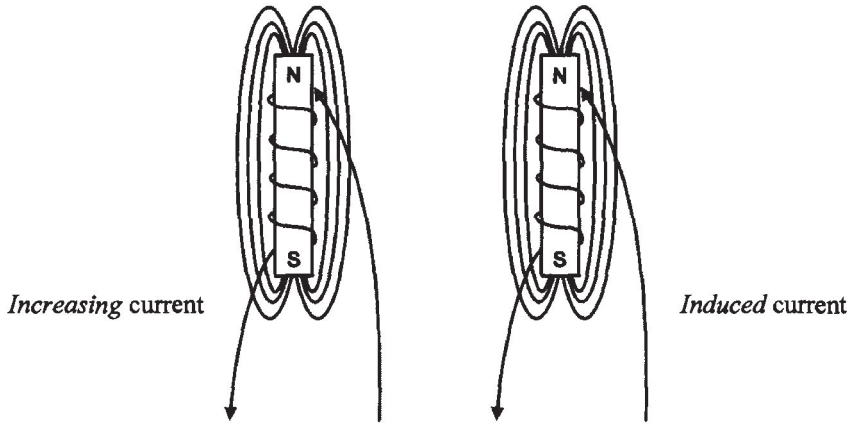


Figure B-8 Very basic transformer action.

This is the principle behind a basic transformer. Current applied to one winding induces a current in another winding.

Figure B-8 illustrates a pretty poor way, actually, to make a transformer. The most obvious improvement would be to wind the two coils (the driven, or primary, and the induced, or secondary) on the same axis, close to each other. There are many other tricks that are beyond the scope of this discussion that people use to make very efficient transformers. Nevertheless, Figure B-8 illustrates a case in

INTEL-1010
U.S. Patent 7,151,430

which one winding can (and will) induce a current in an adjacent winding, which is the definition of a transformer.

Now if Figure B-8 illustrates how one coil can induce a current in an adjacent coil, so does Figure B-9. The only difference between Figures B-8 and B-9 is that we have straightened out the windings into straight wires. As we have seen before (the compass experiment) a current in the wire on the left generates a magnetic field around the wire. If we apply a changing current to the wire on the left, the magnetic field it generates will also be changing. That changing magnetic field will cut across the wire on the right, and will induce a current in the wire on the right. The direction of the induced current will be such as to (a) generate its own magnetic field which will (b) counteract the magnetic field from the wire on the left. Thus, the induced current will be in the opposite direction to the inducing current.

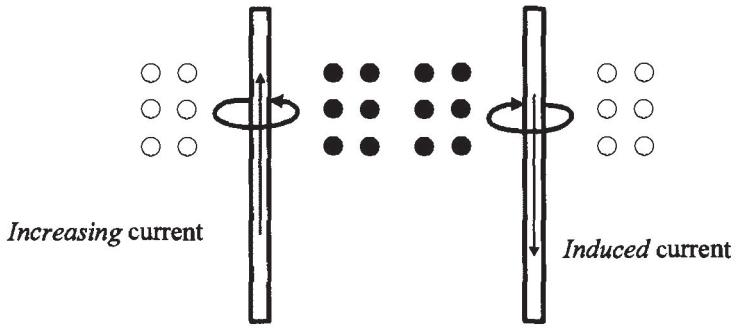


Figure B-9 Two adjacent wires will couple together.

The coupling illustrated in Figure B-9 is pretty weak. It would be stronger for the coils illustrated in Figure B-8 and, of course, much stronger if we had a transformer whose windings were on a common axis. Nevertheless, there will be coupling. The strength of the coupling will be related to (a) how close the wires are to each other, and (b) how strong the changing magnetic field is. The changing magnet field will be stronger for (a) higher currents and (b) higher frequencies. Here is our problem di/dt again. Higher values of di/dt generate stronger, faster changing magnetic fields, which in turn can cause greater induced currents in adjacent wires.

INTEL-1010
U.S. Patent 7,151,430

On a PCB we might have traces only 6 mils apart carrying signals with rise times (dt) of a nanosecond. This is close enough, with a high enough di/dt, that coupling can be a problem. We know this type of coupling as crosstalk.

This type of coupling is called *mutual inductance*. A current in one wire is affecting the current in an adjacent wire. Now let's carry this discussion one final step (Figure B-10), by removing the adjacent wire.

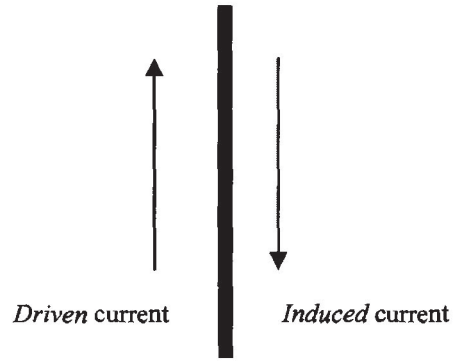


Figure B-10 A single wire will couple into itself.

A changing current in a wire will cause a changing magnetic field that will induce a current *in itself* in a direction that will counteract the changing magnetic field that generates it.

This is called *self-inductance*, or simply inductance. The induced current will be in the opposite direction of the driven current, thereby opposing or "fighting" the driven current. This is the source of the inertia we talked about in Chapter 3. The initial flow of current causes an induced current in the opposite direction, which *for the first instant* results in a net flow of no current at all (until the initial inertia can be overcome). Then, if the change in the magnetic field slows (i.e., the current doesn't change), the induced current begins to reduce and more net current flows. In the steady state case, there is no change in the current, no change in the magnet field, no induced reverse current, and therefore no more inductive effect.

If the driven current is constantly changing (as with an AC sine wave), the magnetic field is constantly changing, and there is a continuing reverse current be-

INTEL-1010
U.S. Patent 7,151,430

ing generated that impedes the primary current flow. This is the inductive impedance that is equivalent to resistance in the DC case.

It is probably intuitive that the inductance here (Figure B-10) is pretty small. It takes either a lot of current or a rapidly changing magnetic field for this little amount of inductance to have much effect. In high-speed circuits where we are dealing with rise times (dt) in nanoseconds, this small amount of inductance can begin to pose a problem. This is why we in the industry stress rise time, not frequency, as the issue. For the same magnitude current, a square wave will have a significantly faster rise time than will a sine wave, so a square wave will generate faster changing magnetic fields, which in turn will be more susceptible to smaller values of inductance.

You have probably picked up on the idea that the magnitude of the inductance is a function of the configuration. Figure B-11 illustrates some general guidelines as to how inductance might change with configuration (increasing from left to right).

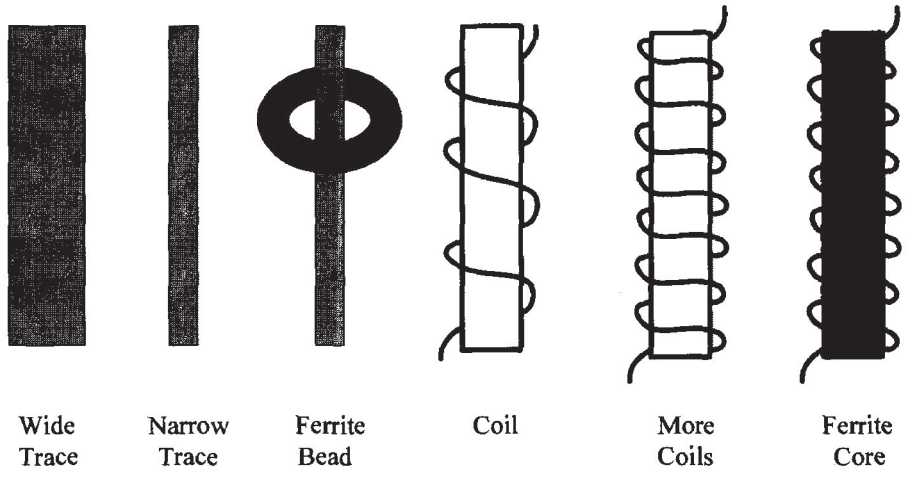


Figure B-11 Methods (left to right) for increasing inductance.

Larger wires have smaller inductance than smaller wires. Think of this as being the result of smaller current densities and smaller magnetic field densities the more wire there is to "spread out" over. On a PCB that is why you sometimes see recommendations that traces between bypass capacitors, for example, and IC pins be as "fat" as possible. That is also why power supply planes generally have lower (smaller) inductance than traces.

INTEL-1010
U.S. Patent 7,151,430

Since inductance is a function of length (i.e., we can think of inductance per unit length) shorter wires have less inductance than longer wires. That is why we tend to keep lead lengths and trace lengths as short as possible on PCs.

We can increase inductance at a point by placing a ferrite bead over the wire. Ferrite (iron) has special magnetic properties and helps "focus" magnetic fields. Therefore, the magnetic fields will be stronger around the ferrite bead. Thus, a ferrite bead will increase inductance and its effects (resisting rapidly changing currents) at that point.

A wire shaped as a coil will have higher inductance than a straight wire. The inductance will increase as (a) we increase the number of coils and/or (b) "tighten" the coils (so there are more coils per unit length). Just as a ferrite bead will increase the inductance of a wire, a ferrite core within a coil will increase the inductance of the coil.

Although there are many ways to create capacitance, there are only a few ways to make inductors. For the most part, all fabricated inductors are coils of wire, probably with a ferrite core. When used in power supplies, inductors tend to be pretty big (at least relatively speaking) compared to other types of components. Wirewound inductors tend to have some parasitic capacitance between the windings. Thus, it is much harder to fabricate a "pure" inductor than it is to fabricate a "pure" resistor or capacitor.

For these reasons, circuit designers don't design very many inductors into their circuits. When an inductor might be useful in a design, we often try to find another way to achieve the same result.

There are lots of little inductors on our boards caused simply by lead lengths. These might be traces or they might be component leads. Even the wirebond leads inside IC packages create little inductors that can be of concern in circuits with fast rise times. These are discussed particularly in Chapter 15 when we talk about ground bounce and the role of bypass capacitors in circuits.

INTEL-1010
U.S. Patent 7,151,430

RETURN CURRENTS

Many of us have heard that when a trace is routed close to a plane, the return current "wants" to be on the plane directly under the trace. The reason is because that is where the impedance is lowest. And the reason for *that* is mutual inductance.

Consider the two traces shown in Figure B-12. Current i_1 is coupled into Trace 2 with a coupling coefficient k. The coupling coefficient, k, is determined primarily by how close Trace 2 is to Trace 1 (distance D). The coupled current goes in the opposite direction. If Trace 2 carries the return current, then the coupling helps "boost" the return current. The return current, in turn, couples into Trace 1, helping to "boost" current i_1. If the currents are "boosted" (i.e., the traces are closer together), then it takes less energy to move the current along than if the currents are not boosted (i.e., the traces are further apart). If it takes less energy to move the same current, that means the impedance to the current flow must be less. Therefore, the impedance is lower if the traces are closer together than if the traces are further apart.

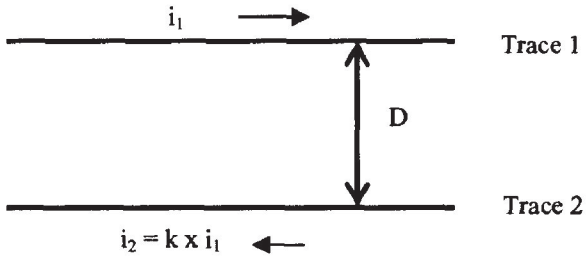


Figure B-12 A changing current in Trace 1 will couple into an adjacent Trace 2.

This is exactly the same phenomenon we were talking about with Figure B-9. Current in one wire can electromagnetically couple to the other through mutual inductance. In the special case where the other wire carries the return current, this coupling actually helps the flow of current by effectively reducing the impedance to the flow.

Now consider Figure B-13. Assume current is flowing down the trace and the return current is flowing back along the plane at point a. The return current will boost the primary current. But if the return current flows back at point b, the boost

will be stronger because the distance between the current and its return path will be shorter. If the boost is stronger, the impedance is lower. The point of lowest impedance will be where the boost is strongest (i.e., point c when return current is directly under the trace). And this tendency will be stronger as our friend, di/dt, gets larger—that is for higher speed circuits.

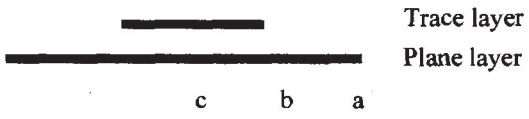
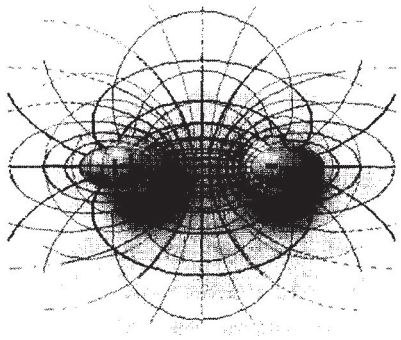


Figure B-13 Signal return currents will be positioned on the plane directly underneath the signal trace.

The fact that the return current wants to flow directly under the trace is also very important in other contexts. It has implications for such things as impedance control (control of reflections), loop areas (EMI radiation), and even crosstalk. As designers, we sometimes do things that prevent the return current from flowing where it wants to. The implications of that are almost always bad.

INTEL-1010
U.S. Patent 7,151,430



C LOGARITHMS

This appendix covers the definition of logarithms (logs), how we convert from logs of one base to logs of another base, and why we care about using logs in the first place.

DEFINITION

A logarithm is an exponent of a number we call the base of the log. In general, if we have an expression of the form $y = a^x$, we say we take the log (to the base a) of y to obtain x, or $\log_a(y) = x$. This looks kind of abstract, so lets look at the most common form of logarithm called, appropriately enough, a common log:

$$\text{If } y = 10^{.5}, \text{ then}$$
$$\log(y) = \log(10^{.5}) = .5$$

(We omit the subscript 10 in common logs, assuming it is understood.)

In this illustration, the base is 10 and the log is .5. Similarly, the (common) log of 100 (10^2) is 2 and the log of one million (10^6) is 6. A more interesting case is the log of something like 486,397:

$$\log(486397) = 5.687, \text{ since}$$
$$10^{5.687} = 486,397$$

INTEL-1010
U.S. Patent 7,151,430

(Note: In all following illustrations there might be slight differences due to rounding. The use of logs with real numbers rarely results in exact results.)

Logs work in reverse, too. For example, if we say the log of a number is something like 3.875 (i.e., log(y) = 3.875), what is the number? This is called an inverse log operation and is written as $\log^{-1}(3.875) = y$. The answer is 7,499, since $10^{3.875} = 7{,}499.0$.

Before computers and calculators we had to look logs up in tables. Now, virtually every handheld calculator will take a log of a number simply by pushing a button, and an inverse (or anti-) log by pushing two buttons.

Base 10 is the most common log, hence the name. There is another base that we routinely use in electronics, the base e, where e is the natural constant equal to 2.7182818..... To differentiate between the common log and the natural log, we use usually use the name ln() to mean the natural log. So, if we have an expression $y = e^x$, then:

$$\ln(y) = x$$

For example, if $y = e^{7.25}$, then $\ln(y) = 7.25$. Similarly, if we have y = 1,000, then $\ln(y) = \text{Ln}(1{,}000) = 6.91$ (since $e^{6.91} = 1{,}000$). And if $\ln(y) = 6.25$, then $y = \ln^{-1}(6.25) = 518$, since $e^{6.25} = 518$.

There are three special log values. Log(1.0) always equals zero, no matter what base we are working in. Any value raised to the zero (0) power is one. The log of zero (log(0)) does not exist in any base. There is no power a number can be raised to whose result is zero. Anytime we take the log to the base n of the value n, the result is one. That is, the log(10) = 1.0, ln(e) = 1.0; even $\log_x(x) = 1.0$.

OTHER BASES

Although base 10 and base e are the most common bases for logarithms, conceptually any base can be used. For any expression $x = a^y$, then $y = \log_a(x)$. For example, we might routinely want to work in base 2 (binary system). Thus, if $y = 2^5 = 32$, then $\log_2(y) = \log_2(32) = 5$. We might also want to work in base 8 (octal system). So $\log_8(247) = 2.65$, since $8^{2.65} = 247$. Or, we might want to work in the

INTEL-1010
U.S. Patent 7,151,430

hexadecimal (hex) system, base 16. We could even work in a system that makes little intuitive sense, such as 7/8. For example, what is $\log_{7/8}(125)$?

CONVERTING BETWEEN BASES

Although we have tables, and now calculators, for common and natural logs, we generally don't readily have them for any other base. It is not too difficult, however, to convert from any one base to any other base.

To start with, the very first pair of equations we started this appendix with present a special case. The general case is:

$$\text{If } y = 10^{.5}, \text{ then}$$
$$\log(y) = (.5)\log(10)$$

We reduce this to $\log(y) = .5$ since $\log(10) = 1$. But if we were in another base, like base 2, then the correct result would be:

$$\log_2(y) = (.5)\log_2(10)$$

Now, how do we determine log(y) for any base? For example, what is $\log_{7/8}(125)$?

(1) Let $y = a^{x1}$
(2) Therefore, $\log_a(y) = x1$
(3) Also, $\log_{10}(y) = (x1)\log10(a)$

In this example, $y = 125$, $a = 7/8$, and x1 is the answer we are looking for.

Note that we can now solve for x1:

$$x1 = \log_a(y) = \log_{10}(y)/\log_{10}(a)$$

Therefore, to find the log of y for any base, take the log of y to base 10 and divide by the log (base 10) of the base of interest. In our earlier example,

$$\begin{aligned}\log_{7/8}(125) &= \log_{10}(125)/\log_{10}(7/8)\\ &= 2.0969/(-.05799)\\ &= -36.1586\end{aligned}$$

We can check this by seeing if $(7/8)^{-36.1586} = 125$, and it does, as we can confirm from most calculators or a spreadsheet.

USEFULNESS OF LOGS

So what? Why do we care about logs? That's a fair question. I'll give you three answers in this appendix.

First, before the days of electronic calculators (about 1970 or so) a complex problem could be very difficult to work out. For example, suppose you had to solve for

$$y = (25.875)(36.54^2)(74.36^{1.75})/.8278^{.9}$$

Today, we'd put that in a spreadsheet or grind through it on a calculator. Before that we would simplify it using logs:

$$\log(y) = \log(25.875) + 2\log(36.54) + 1.75\log(74.36) - .9\log(.8278)$$

That may not look too simple, either, but it has been reduced to four log look-ups (in a table), three multiplications, and some addition. When we had a numerical answer on the right side, we'd do an inverse log and solve for y.

So, at least during the days before electronic calculators and computers (i.e., when problems were solved with paper, pencils, and slide rules) logs simplified complex calculations.

A second purpose for logs is to linearize an expression. Suppose you wanted to do a curve fitting analysis and find the best curve that fit the equation:

$$y = Bo(x1^{B1})(x2^{B2})$$

INTEL-1010
U.S. Patent 7,151,430

That is, for a set of data x1,x2,y, what is the best guess for the coefficients Bo, B1, and B2? Setting up a least squares curve fitting model for this is not too difficult, but it is nonlinear. It can be converted to an easier linear form using logs:

$$\log(y) = Bo + (B1)\log(x1) + (B2)\log(x2)$$

which looks more like the familiar

$$y = bo + b1(x1) + b2(x2).$$

This form may be easier for doing a least squares fit for the coefficients Bo, B1, and B2.

We often use logs for scaling graphs. Suppose we want to graph the reactance of a capacitor as a function of frequency (recall $Xc = -1/j\omega C$). If the horizontal scale is linear, it is impossible to look at the relationship in the region of, say, 1 kHz, 10kHz, 100kHz, and 1MHz on the same sheet. If we can see details in the kHz range, then 1 Mhz will be far away (in the next room or block). If we can see detail in the MHz range, then the kHz range will be but a spec.

By converting the graphical scales to logarithmic rather than linear, the graph becomes much easier to read. See Figure C-1. The linear graph has a shorter range than the log one, and still offers much less detail than the log one does.

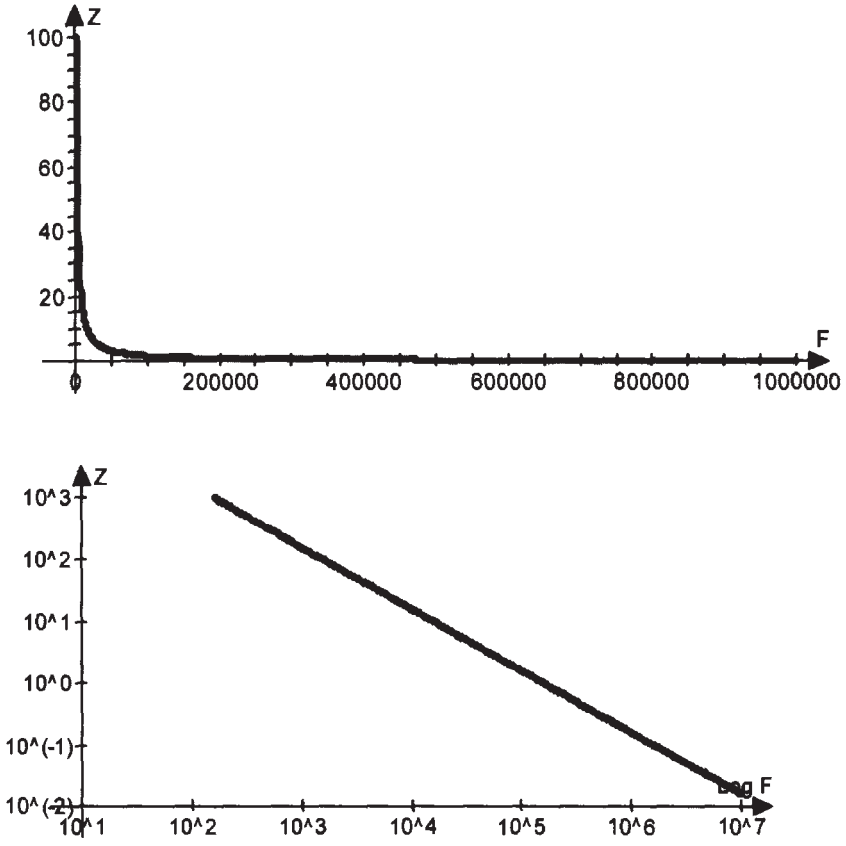
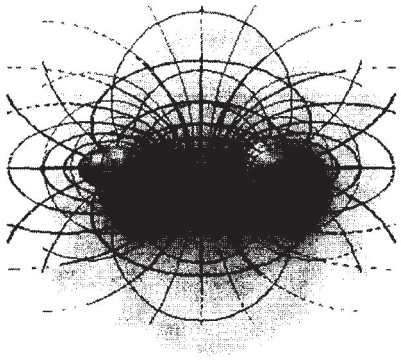


Figure C-1 Reactance of a 1 uF capacitor graphed on linear (top) and log (bottom) scales.



D PHASE SHIFT SIMULATION

BACKGROUND

This simulation is intended to help you visualize the types of phase shift that occur in circuits as the relationship between circuit components changes. Download rlc.zip from *www.ultracad.com/prenticehall* and copy it to a convenient place on your hard drive. Unzip it to obtain rlc.exe. When you execute this file, the program will start (Figure D-1).

All component values are initially set to 1.0 (Ω, uH, and uF, respectively). The angular frequency, ω, is fixed at 10^6. You will note that this is the resonant frequency under these conditions. Component values can be changed by the sliders across the top and the values are shown in the text boxes at the right. The total impedance and the phase angle are also shown at the right, and the system is at resonance when the components are selected such that the phase angle is zero.

The schematic of the basic series RLC circuit is shown at the left. The circuit simulation can be set up for either constant current input or constant voltage input by the selection buttons at the lower right of the window. The check boxes across the bottom select what is displayed. The options include the voltage across the resistor (V(R), green), the inductor (V(L), yellow), and/or the capacitor (V(C), blue). The input current ((I(in) in mA) and the total voltage (V(in), which is the sum of the other three voltages) can also be displayed. The white curve is always the constant input variable (voltage in or current in) and the red curve is the other input.

INTEL-1010
U.S. Patent 7,151,430

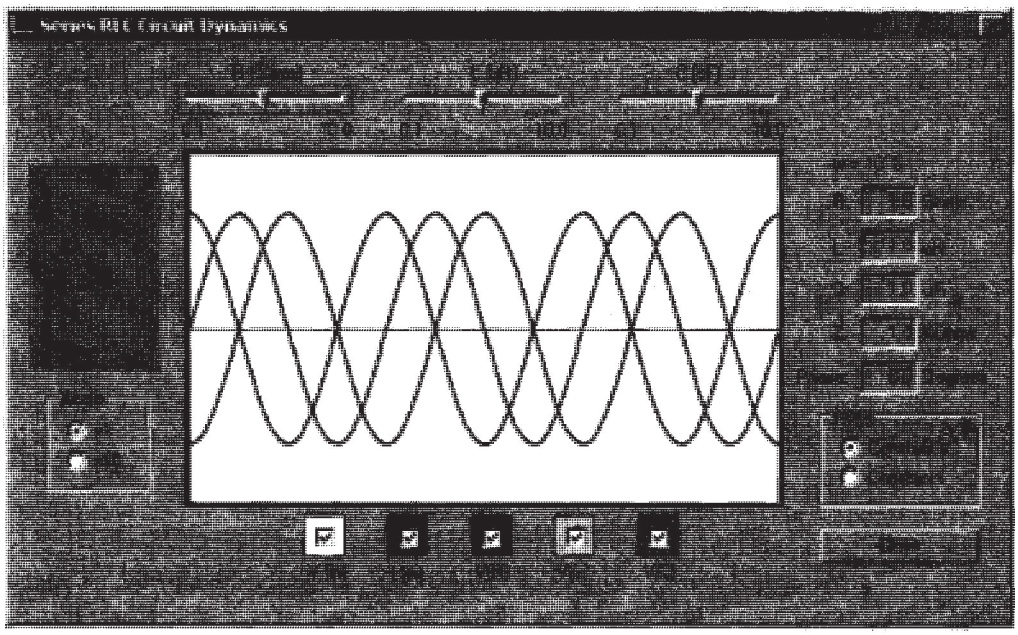


Figure D-1 UltraCAD's Phase Shift Simulator.

This simulation can be complicated to get into, so you might find it convenient to start with the following sequence. After completing this sequence, you can play with the simulation in other ways that interest you.

FIRST STEP

Start with the input set to constant current. Set all component values to unity. Select I(in), V(R), V(L), and V(C) for display. You probably see only three curves, and either the white or green one is missing. Why? Because the voltage across the resistor and the current are always in phase with each other. With these component selections, their magnitudes are also equal. So the curves lay on top of each other. Now deselect the I(in) curve for display.

Note how the V(C) curve lags the V(R) curve by 90 degrees and how the V(L) curve leads the V(R) curve by 90 degrees. Note that these phase relationships *will not change* throughout this entire simulation. The voltages across these three

INTEL-1010
U.S. Patent 7,151,430