Introduction:

Before describing the Ohmmapper, it is worthwhile to review just what an Ohm is and why you may want to map them.

In the late 1700's, when electricity was first being discovered, people knew that there was an electromotive force that tried to cause charges to move from one place to another. It caused "like charges" to repel each other and "opposite charges" to attract. Mr. Volta spent a fair amount of his time studying properties of electricity and got a unit of measure, the Volt, named after him as a result. The Volt is the unit of measure for the electromotive force. Often this is referred to as voltage or potential.

Scientists had also figured out that they could measure the rate at which the charge passed through a given point. Mr. Ampere realized that this charge would flow down a copper wire in a manner much like water flows in a pipe. He developed a device to measure this rate of flow and thus got his very own unit of measure named after him, the Amp. He may also be the person that coined the term "current" for the flow of electrical charges.

Now we come to the hero of our story, Mr. Ohm. It was only when Mr. Ohm came along that anyone succeeded in relating the electromotive force in Volts to the current in Amps. He realized that this flow of charge through a wire would not happen without causing some friction. He called this friction "resistance." After several false starts he came up with the simple formula:

$$I = E/R$$

This is the famous "Ohm's Law", which simply says that you divide the Electromotive force (measured in Volts) by the Resistance (measured in Ohms) to get the current (measured in Amps). There is a rumor that using "I" to stand for current is because it originally stood for "intensity".

People quickly saw that you could rearrange this to two other forms:

$$E = IR \quad \text{and} \quad R = E/I$$

With this simple set of equations most of the mysteries of electricity could be explained.

It was at about this same time that people realized that if you make an iron wire that is the same dimensions as a copper wire, the iron wire has a higher resistance. This difference in resistance is due to a property of the material. They decided to call this property "resistivity". Iron has a higher resistivity than copper.

Once folks realized that different materials have different resistivities, they figured out that this was a good way to identify materials. However, there was a problem:
If you make a copper wire twice as long, that makes the wire's resistance twice as high. The

resistivity of copper has not changed. All that has changed is the shape ("geometry") of the wire. This again is much like water flowing in a pipe. If water encounters some amount of friction flowing through a mile long pipe, it will encounter twice as much friction if another mile of pipe is added on the end. Just as making a pipe bigger around lowers its friction, making a wire bigger around lowers its resistance.

If we are going to determine what material an object is made of, based on its resistivity, the first step is to find out what the resistivity of that object is. The measurement must be made to depend only on the material the object is made of. We need a measurement that does not depend on the shape of the object. For example, doubling the length of a copper wire should leave this measurement unchanged. The measurement should also not depend on where the measuring device is placed (on the object). When measurements are corrected for both the shape of the object and the placement of the measuring device, they are said to be "corrected for the geometry." When values are corrected for the geometry, they are usually expressed in Ohm-Meters.

Up to here, the discussion has centered on electric currents flowing like water in pipes. When the object having its resistivity measured is the earth, the "water in a pipe analogy" no longer works. Instead, imagine a very large flat floor with a drain hole in it. We begin pouring water onto the floor. Some of the water will take the direct route straight to the drain. Most of it, however, will take a curved path. It will spread out, curve around and come back together at the drain.

In a similar way, the electrical current that a resistivity measurement system puts into the soil follows a curved pattern. A resistivity measurement system always has two "current electrodes" that it uses to force a current to flow through the soil. The current enters the soil at one "current electrode", spreads out, curves around, and comes back together at the other "current electrode". This is happening in three dimensions: the electrical current spreads out to the left, to the right and down into the soil.

Because the soil has resistance, whenever there is a current flowing in the soil, according to Ohms Law, there must be some voltage. From the $E = IR$ form of Ohm's Law, we know this voltage will depend on the current and the resistance. Our goal is to determine the resistance. If we knew the current and the voltage, we could use the $R = E/I$ form of Ohm's Law to find the resistance.

We can measure the voltage in the soil if we place two more electrodes on the ground and measure the voltage between them. These electrodes are called the "potential electrodes".

If we know the positions of the "current electrodes" and the "potential electrodes", we can use this and the calculated resistance to figure out a resistivity value. This resistivity isn't the resistivity of one small chunk of the soil, but is, instead, some sort of average of a large volume of soil under the resistivity measurement system. For this reason this measured resistivity is usually called "apparent resistivity".

The soil we are measuring must not all have the same resistivity. If it did the Ohmmapper would be of little practical value. The resistivity of the soils and rocks varies over a very wide range. Since the Ohmmapper measures the resistivity of the soil beneath it, a map can be made (from the data gathered while moving the Ohmmapper over the surface) that will show areas of higher and lower resistivities. Those resistivities can be used to identify geologic formations and other things of interest under the ground.

At this point, I expect some readers (who have actually seen an Ohmmapper) are wondering just how the Ohmmapper manages to put this electrical current into the soil and measure the voltage it produces without any obvious electrical connections to the soil. To explain this, we need to go back to those folks in the late 1700's.

Mr. Leyden knew that he could create an electrostatic charge by simply rubbing a glass rod with a piece of wool. To store an electrostatic charge, he invented a device now called the Leyden jar. Mr. Leyden took an ordinary glass jar and wrapped the bottom of its outside with metal and coated the inside with metal. He hung a wire from the lid into the inside of the jar. He found that this device could store a charge. It was soon discovered that larger ones would store more charge than smaller ones, if the voltages were all the same. This property of the Leyden jar is called its capacitance.

For some reason, Mr. Faraday's name got used for the unit of capacitance and the unit is called the Farad.

The modern version of Mr. Leyden's invention is called a capacitor. Modern capacitors are seldom built with glass. Any nonconductive material, sandwiched between two conductors, will work.

Lest anyone think that all of these people back in the late 1700's were outright geniuses: One day Mr. Leyden decided he wanted to find out what the magic fluid he was storing in his jars tasted like. He attempted to drink from a charged Leyden jar. When his lips made contact with both the inside and the outside of the jar he got quite a shock. We know this story because he lived to tell us what he did.

In Mr. Leyden's day, no one knew that atoms have electrons. Today, we do, so we can explain what is happening. Metals are good conductors because the electrons in them are free to move around. In metals, the electrons in the outer-most orbit around an atom are so far away from their atom and so close to the neighboring atom that you can't really say any given electron belongs to any given atom. The situation is more like a sea of electrons washing around the atoms.

Electrons repel each other. If you bring some electrons near one end of a piece of metal, those electrons will repel the electrons in the metal. The electrons on the surface of the metal will be pushed in to the metal. These electrons will repel their neighbors and push them away. Those neighbors will repel the next ones and so on, all the way through the metal.

For this reason, a charge put into a capacitor on one conductor seems to come out on the other conductor. It is not actually the same electrons. It is, instead, an equal number of electrons that were repelled by their neighbors, which, in turn, were repelled by their neighbors and so on, all the way back to the ones that were repelled by the electrons on the first conductor.

If we take a capacitor and put a charge on one of the conductors, and then take it back off, the electrons on the other conductor will be repelled only during the time the charge is on the first conductor. If we alternate applying the charge and removing it, an equal charge will flow in and out of the other conductor. Therefore, an AC (alternating current) voltage applied to one conductor of a capacitor will appear on the other conductor.

This ability to cause a charge to move in and out of one conductor by applying an alternating voltage to another nearby conductor is called "Capacitive Coupling". This principle is used in two places in the Ohmmapper. The transmitter uses capacitive coupling from its electrode to the soil to cause the current to flow within the soil. The receiver detects the voltages in the soil because of the capacitive coupling from the soil to the receiver's electrodes.

At this point I need to introduce two new words. The first is "reactance". Imagine that there is an AC current flowing back and forth through a capacitor. Because the capacitor is being charged up in one direction, discharged and then charged in the other direction, there is an AC voltage on the capacitor. Here we have an AC voltage caused by an AC current so there is a strong temptation to use the R = E/I form of Ohm's Law to find a resistance. However, we can't call what we calculate a resistance, because it is an effect of the capacitance. For some reason, lost in the mists of time, it is called "reactance" and the symbol used for it is "Xc".

The 3 forms of Ohm's law can each be written using Xc:

E = IXc   I = E/Xc and Xc = E/I

The voltage produced by a current flowing through a resistance happens as the current is flowing. In the case of capacitive reactance, however, it is quite different. At the point in time when the current stops flowing in one direction and is about to start flowing in the other, the capacitor has its highest charge and hence its highest voltage. The voltage is said to lag the current, because it reaches its peak value after the current does.

A smaller valued capacitor will reach a higher voltage when storing the same charge. The charge the capacitor must store is determined not only by the amount of current charging it. It also depends on how long the current is allowed to flow into the capacitor. If the frequency of an AC current applied to a capacitor is increased, the time the current flows in any one direction will be less and so the charge put on the capacitor will be less.

By stirring these facts around with a little mathematics, we can come up with an equation to get Xc:

Xc = 1/(2.PI.F.C)

Where:

Xc	is the reactance in Ohms
PI	is the usual 3.14159... etc. thing
F	is the frequency in Hz (Cycles per second)
C	is the capacitance in Farads

"Impedance" is a general term that can be used to mean resistance or reactance or some combination of the two. For our purpose, "impedance" is used only for the combination. If only resistance is being discussed, the word "resistance" will be used.

An Ohmmapper's electrodes:

The electrodes used on the Ohmmapper look, externally, like electrical cables. Internally, they have a twisted pair of two wires of modest gauge, a nonconductive filler to pad the diameter out to the desired size and then a copper braid wrapped over the filler. Over the copper braid, a tough outer insulation is placed. It is this copper braid that acts as the electrode for the Ohmmapper. The capacitive coupling from the copper braid to the soil couples the transmitter's current from its electrodes to the soil. The voltage on the soil is capacitively coupled into the braid on the receiver's electrodes.

In the traditional, or galvanic, soil resistivity measurement system, the electrodes are conductive rods that are driven into the soil. These electrodes are often called "point electrodes" because they only make contact to the soil at a single point.

The electrodes of the Ohmmapper are quite different. They capacitively couple to the soil over their entire length. For this reason they are called "line electrodes".

The explanation of why the transmitter's current couples into the soil evenly along its entire length, requires the following facts:

> The resistance of the copper braid in the electrode cable is very low.
>
> The diameter of the copper braid in the electrode cable is held very near constant by the manufacturing process.
>
> Even with small changes in the height of the electrode above the soil, the capacitance per meter of electrode is nearly constant down its length.
>
> The voltage required to cause the transmitter's current to pass through the capacitive reactance from the electrode to the soil is much larger than the voltage required to cause the current to flow within the resistance of the soil.

Given all of the above, it is the capacitive reactance from the electrode to the ground and the AC voltage on the electrode that control the current flowing into the ground. We know that the capacitance per unit length is very nearly constant down the length of the electrode. From this we know that the capacitive reactance per unit length must also be nearly constant. Since the resistance of the copper braid is very low, from the

$$E = IR$$

form of Ohms law, we know that the transmitter's current flowing down the electrode cable will produce very little difference in voltage from one end of the electrode cable to the other. Since the current flowing in to the soil is determined by two values that are very nearly the same at all points down the cable, it is logical that the "current per unit length" must also be very nearly the same at all points down the cable.

The receiver's electrodes receive a voltage that is the average of the voltages on the soil they are over. To explain why this is true, one more fact must be added to the above points.

> The receiver does not take any current away from its electrode cable. Any AC current that capacitively couples into the cable at one point along its length must capacitively couple out of the cable at another point along its length. The AC current coupling Out of the cable will only match the AC current coupling into the cable if the AC voltage on the cable is equal to the average of the AC voltage on the soil it is over.

The Ohmmapper transmitter:

The Ohmmapper's transmitter causes an AC current to flow in the ground. Because the only path for the AC current from one electrode to the other is via the soil, the transmitter can determine the current it is producing in the soil by measuring the current flowing into one of the electrodes. Using this measurement, the transmitter can quickly adjust its output voltage until some desired amount of current is produced.

The output of the transmitter is an AC current of 0.125, 0.25, 0.5, 1, 2, 4, 8, or 16mA. This current is very tightly regulated, (1%) so that the receiver only needs to know which range the transmitter is on. The transmitter automatically selects the correct range for the electrodes used and the soil conditions.

The transmitter encodes the current range onto the AC signal used for the measurement. It also encodes a 2Hz signal onto the AC signal. The encoded 2Hz signal is used by the receiver to get in step with the transmitter. It also serves as a reference frequency for the decoding of the current setting.

The operation of the power output stage is monitored. If the voltage required to cause the specified current gets close to the limit of the system, the transmitter will switch to the next lower current setting at the next 0.5 second increment. This places the change in current such that the receiver's cycle will start right at the change. In this way the transmitter can, change currents without causing an error in the receiver. The output stage is also checked to see if it is running at less than 1/3rd of its full power. If it is the power will be increased to the next higher setting.

The transmitter output is a good sine wave. Internally it is created by what is called a "class D" power amplifier.   This allows us to have good efficiency while transmitting a sine wave. The natural output impedance of the transmitter power stage is high. This prevents short-term changes in the output current. The actual transmitted current is measured by a current sensing circuit and this is used as a feedback to regulate the current. This gives an output impedance well above 10M Ohms.

The transmitter can produce output voltages of up to 1000VRMS and currents of up to 16mA, but not at the same time. The output power never exceeds 2 Watts.

The Ohmmapper receiver:

The input to the receiver has an impedance of greater than 10M Ohms at the operating frequency. This high impedance is required to prevent any current from flowing from the electrode into the receiver. At other frequencies its input impedance is lower. This was done to reduce the effects of cultural noise such as the 50 or 60Hz mains frequencies.

The difference in voltage between the two electrodes is filtered and amplified. The filters only pass a band of frequencies about 25Hz wide. These filters are tuned to the transmitter's frequency. There is an automatic fine-tuning system that matches the filter exactly to the transmitter's frequency. The amplifier’s gain can be selected by the processor in the receiver. The gain can be set to multiply the signal by 1, 4, 16, 64, 256, 1024, 4096 or 8192.

The gain of 1 allows the system to work with a 2V AC signal. In most environments the 256, 1024 or 4096 gain is used. The processor selects the gain that makes the signal going into the analog to digital converter as large as possible without risking exceeding the range of the converter.

The ADC is more than "20 times over sampled". This means that it measures the AC voltage much more often than is needed to accurately measure the amplitude of the received signal. The over sampling allows many more ADC readings to be used for each output number. This greatly improves the ability of the system to pick signals out of the noise. The samples are digitally filtered to give an overall system bandwidth of about 4Hz. This means that noise at frequencies more than 2Hz either way away from the transmitter's frequency are strongly rejected.

The receiver knows that it is in fact measuring the transmitted signal when it can detect the 2Hz and 4Hz modulation the transmitter places on the signal. The receiver phase locks its self to the transmitter's 2Hz so that 2 receivers measuring the same transmitted signal will run in exact lock step. This means that two receivers could be used with no electrical connection between them. The 4Hz modulation is used to indicate the current setting of the transmitter. In this way we do not need an electrical connection to the transmitter.

At power on the receiver does the following:

1 Say "Ver 1_0"
2 Adjust the gain
3 Say "Set gain"
4 Take 1/2 seconds worth of data
5 If the gain needs readjusting, restart from step (2)
6 Say "Phase A"
7 Take 1/2 seconds worth of data
8 Say "Phase B"
9 If the gain needs readjusting, restart from step (2)
10 Use all of the data so far to look for the transmitter
11 If we are not confident about the phase go to (4)

The receiver will stay in this part of the code and produce the "Phase A" and "Phase B" messages until it finds what it believes is the transmitter. It then does two more checks to make sure it really has a good solid lock on the signal from the transmitter. During this time it will say "Phase C" and "Phase D". After this it should begin to produce readings.

When running the receiver does the following:

12 Measure for 1/2 second
13 Determine if the gain needs to change
14 Find the phase of the 4Hz to determine the current
15 Find the phase of the 2Hz and trim the timing
16 If the transmitter's signal is gone go to (2)
17 Calculate the voltage that 1mA would have made
18 Send the results
19 Go to 12

Any gain change called for by any given cycle will not be applied until after the next reading. This means that it would take 8 seconds to step through all 8 of the gains.

Data is sent electrically out of the receiver and along the electrode cable.

Cables and stuff:

We are using a shielded cable with a strength member built into it and a cable jacket that takes wear well. The cables are a wear item. The receiver and transmitter electrodes are the same type of cable. A 2-meter electrode can be made by simply plugging 2 one-meter electrodes together.

The electronics bottles have an expendable plastic cover. This cover is held in place by the leading connector and the bottle.

The transmitter is only powered on when complete cables are connected. This means that the terminating plugs must be on the end of the electrodes before it will transmit, thus the 1000VRMS is never on an exposed surface.

Understanding "skin depth":

The "skin depth" effect places a limit on how far into the soil the Ohmmapper can see. In order to understand what causes this effect we need to introduce a few more of the things discovered by those clever folks back in the late 1700s.

It had been discovered that, whenever an electric current flows it creates a magnetic field. People first noticed the effect when they placed a compass near a wire. Every time they sent a current through the wires the needles of the compasses moved. If they turned off the current the needle went back to pointing north like compasses normally do. If they reversed the direction the current flowed the compass needle was deflected in the other direction. The magnetic field created by the current is always at right angles to the direction the current is flowing.

A man by the name of Mr. Faraday discovered an effect we today call "Faraday induction". Mr. Faraday discovered that if he swept a magnet past a wire, it produced a voltage in the wire. The faster he moved the magnet the higher the voltage he produced. If he swept the magnet past in the other direction the direction of the voltage also reversed.

A person by the name of Mr. Lens put the fact that electric currents create magnetic fields and Mr. Faraday's observations together. Mr. Lens thought to himself "When the current flowing in a wire is changed, the magnetic field around the wire will be changed. This changing magnetic field will induce a voltage in the wire." Experiments showed that changing the current in a wire did indeed cause a voltage to appear on the wire. The direction of the induced voltage was always such that it worked against any attempt to change the current. Today this is referred to as "Lens's Law". The tendency of a change in current to create a voltage is referred to as inductance. The unit of measure for inductance is the Henry.

Much like in the case of capacitive reactance, we have a voltage being caused by an AC current. In this case the effect is called inductive reactance and given the symbol "Xl". Unlike the case of the capacitor, the more rapidly the current changes back and forth the higher the voltage that will be produced. Xl can be calculated with the equation:

Xl = 2.PI.F.L

Where:

Xl	is the reactance in Ohms
PI	is the usual 3.14159... etc. thing
F	is the frequency in Hz (Cycles per second)
L	is the inductance in Henrys

As you may expect inductive reactance has its own form of Ohm's law.

E = IXl I = E/Xl and Xl = E/I

The Ohmmapper causes an AC current to flow in the soil. This current will cause a small voltage to be produced. From Len's law we know that this AC voltage will try to prevent the AC current from flowing.

Lens's law is not too hard to picture in the simple case of a wire. The AC current flowing in the soil is spread out through the soil, so it is a bit more complicated. The magnetic field created by an AC current flowing in one part of the soil will spread out through the soil and create voltages in other parts of the soil. The current that is flowing in the very top of the soil creates a voltage that tends to block the current from flowing in deeper parts of the soil.

In very resistive soils, the voltage created by any given current flowing in the resistance of the soil will be large compared to the voltage caused by the inductance of the soil. For this reason, the skin effect is not a problem in very resistive soils.

If the soil is not very resistive, the voltage created by the inductance can be large enough that it reduces the current flowing deep in the soil and crowds more of the current up near the surface.

Many people at this point will want to ask the question "Don't these voltages caused by the inductance offset the reading of the Ohmmapper?" The answer is, yes, but by a lot less than you might expect.

It is important to remember that the voltage caused by the inductance and the voltage caused by the resistance is out of step with each other. The voltage due to the resistance is produced when the current flows. The voltage due to the inductance is produced when the current changes. Its peak actually happens at the instant the current has stopped and is reversing directions.

There is a method to calculate the sum of two voltages that are out of step in this way. To add the two voltages you square each of them, add the squares together and then square root the sum. Proving this is the correct procedure is left as an exercise for the reader (i.e.: I'm too lazy to do it now so just trust me).

If we try a couple of examples, it should be obvious why the voltage due to the inductance is less of a problem than may be expected. Take the case where the voltage due to the inductance is 20% of the voltage due to the resistance. To make the math easy we can assume the voltage due to the resistance is just 1V. First we must square the two numbers. The square of 1 is 1. The square of 0.2 is 0.04. We add to get 1.04 and then square root to get 1.02. What at first looks like a major of 20% in fact only changes the reading by 2%.

In the real world, the voltage due to the inductance will never be as much as half of the voltage due to the resistance, but even for this extreme case the effect is not very large. One squared is one. A half squared is 0.25. The square root of 1.25 is 1.12 so even in this extreme example the error is only 12%. In practice, the receiver and transmitter would have to be so far apart that the receiver would not be able to detect the transmitter in order to get proportions as bad as these.