

Proton Precession Magnetometers, Revision 3

James A.R. Koehler

March, 2012

Comox, British Columbia, Canada

Copyright J.A.R. Koehler, Comox, BC, Canada 7 March, 2012

1 The Physics of Proton Precession and the Effects of Sensor Geometry

1.1 A few Words about Units

When I was a student in physics in the 1950s, the usage of units was confused to say the least. Most sub-branches of physics seemed to adopt a set of units which were convenient for that field (and they still do to some extent!) and, as they used to say in Rome, *scabes extremum occupet*, “May the itch catch the hindmost”. Physics was made even more difficult than it needed to be by the necessity to convert between cgs units, mks units, esu (electrostatic units) and emu (electromagnetic units) not to mention the continual frustration brought on by the additional occurrence of British units (lb-ft-sec). Indeed, looking back on it now, it seems to me that most of the problem sets assigned to us by the prof’s were concerned with converting all the disparate units into one self-consistent set appropriate to the problem! By the time I was teaching physics in the mid 1960s some sanity had come to the field and most textbooks, at least, were being written with a consistent set of units, most frequently the SI (Système International) set of units. However the usage in magnetic units has retained some of the confusion that used to be widespread in all other topics. The situation is better (much better!) but still not perfect.

Part of the problem is that there is still some controversy about which is the more fundamental property of magnetism - is it the “magnetic field”, commonly designated by the symbol, \mathbf{H} , or is it the “magnetic induction” or “magnetic flux density” commonly designated by the symbol, \mathbf{B} ? The two are related, in any material by the equation:

$$\mathbf{B} = \mu\mathbf{H}$$

where μ is the permeability of the material. In a vacuum or free space, the permeability is defined as μ_0 which is called the permeability of free space and has a value of precisely $4\pi \times 10^{-7}$ in SI units. I will not contribute to the discussion except to say that I have a slight preference for \mathbf{B} rather than \mathbf{H}

and I will tend to use it in the discussion that follows unless I am referring to materials by other authors who have chosen to use the other. In common with most workers in the field, I carelessly often call \mathbf{B} the magnetic field even though the precise term should be the magnetic induction.

The unit of \mathbf{B} in the SI system is the Tesla, named after the famous/infamous scientist/charlatan (take your pick of the four possible permutations) Nikola Tesla. The earth's magnetic field varies from location to location over the surface of the earth from a value of very roughly 60 microTeslas ($6 \times 10^{-5}\text{T}$) near the magnetic poles to very roughly 20 microTeslas near the magnetic equator. It is a vector which means that it has direction as well as magnitude. The convention is that \mathbf{B} , at any location, points in the direction that the north-seeking pole of a freely suspended magnet would point if suspended in that location. That means that in the northern magnetic hemisphere of the earth, the local magnetic induction, \mathbf{B} , points downwards and towards the north magnetic pole of the earth.

1.2 The Physics of Proton Precession

The proton has both angular momentum and a magnetic moment. Having angular momentum means that one can think of it as having some spin like a little top. The angular momentum is truly very small - just $5.27 \times 10^{-35}\text{kg}\cdot\text{m}/\text{s}$. The magnetic moment means that it behaves like a small (very small!) magnet. It has a magnetic moment of 1.41×10^{-26} A-m. The direction of the angular momentum and the magnetic moment is identical - they are parallel to one another. In one of those minor misfortunes in physics, the same symbol, $\boldsymbol{\mu}$, is used for the magnetic moment of the proton as is used for permeability. There just aren't enough Greek letters to uniquely define everything, I guess. However, remember that the proton's magnetic moment has both magnitude and direction - it is a vector - while permeability is just a scalar constant. Because the angular momentum, \mathbf{L} , and the magnetic moment, $\boldsymbol{\mu}$, are vectors which point in the same direction, they are related to each other by a constant, γ_p , in the equation:

$$\boldsymbol{\mu} = \gamma_p \mathbf{L}$$

γ_p is called the gyromagnetic ratio and has a value of 2.67512×10^8 in SI units. γ_p is dimensionless and is one of the fundamental constants of physics. If a proton is placed in an external magnetic field, \mathbf{B} , because of its own magnetic moment, it will experience a magnetic torque. Because it also has angular momentum, this magnetic torque will cause it to precess - this is called Larmor precession and the rate of precession depends on the magnitude of the external magnetic field. It turns out that the rate of Larmor precession is *independent of the proton's orientation with respect to the external field* and depends *only* on the magnitude of the external field. The rate of precession, ω , in radians per second, is given by the Larmor Equation:

$$\omega = \gamma_p B$$

where B is the magnitude of the external magnetic field. For $B = 50 \mu\text{T}$ (a nominal value for the earth's field over much of Europe and North America), ω is therefore 1.34×10^3 radians/second which corresponds to 2129 Hz. The basis of the proton precession magnetometer is to make an instrument to measure the precession frequency and then, using the Larmor Equation, calculate the magnetic field strength.

To do so, the previous equation may be rearranged to give the magnetic induction, in nT, in terms of the precession frequency:

$$B = 23.4875f$$

where f is the frequency in Hz. From this equation, one may get some idea of the precision required in the measurement of the Larmor frequency; one nT change in the magnetic field just causes a change of 1/23.4875 Hz. In other words, to measure the local magnetic induction to an accuracy of one nT (10^{-9}T), you must measure frequency with an accuracy of 0.0426 Hz!

In bygone years, for many branches of physics and geophysics, the cgs unit of magnetic induction was used; the Gauss. Commonly, in geophysics, an even smaller subunit, the gamma was used. One gamma was 10^{-5} Gauss. Expressed in the SI system of units, one gamma is exactly equal to one nT.

1.3 The proton precession magnetometer

There are protons in the nucleus of all atoms but there is only one atom that has a single proton and that is hydrogen. So it is hydrogen atoms which supply the protons necessary for a proton precession magnetometer. Other elements do not count. If all protons will precess in the presence of an external magnetic field, why can't we just put several loops of wire around a cup of water (or any other material containing hydrogen atoms) and measure the frequency of the ac voltage induced in the wire by all these little magnets precessing? The answer is that the protons, at room temperature, are all lined up in random directions. Therefore there will be as many turning in such a direction as to induce a positive voltage in the coil as there are in the opposite direction which will induce a negative voltage in the coil. In another words, because they are not all precessing in phase, the net induced voltage will be zero.

To make a magnetometer using proton precession, the protons must be "magnetized" (i.e., lined up in the same direction) and then all "let go" simultaneously. They will then all precess in phase and induce a voltage at the Larmor frequency which can be measured and used to calculate the ambient magnetic field.

Typically, the protons are subjected to a polarizing magnetic field which will line them up in its direction. Then the polarizing field will be turned off and the induced voltage detected. The simplest (but not necessarily the best!) configuration is a multi- turn solenoid containing a core of some substance which is rich in protons - distilled water, hydrocarbon fluids like kerosene, light oils or alcohol are commonly used. A large current is passed through the solenoid to generate a large magnetic field in the core in order to line up the protons.

The current is then switched off and the same solenoid acts as a sensor and is connected to a sensitive amplifier in order to detect the induced precession frequency. Typically, the induced voltage due to proton precession is of the order of microVolts. Because the precession of the protons will subsequently be randomized by thermal collisions, the induced signal decays exponentially with time and has a typical time constant, which depends on the particular proton rich substance used in the core, of a few seconds. The process is repeated to make another reading.

Since the polarizing current may be of the order of Amperes and since the polarizing field must be turned on long enough to line up the protons (this is also exponential and, typically, also has a time constant of a few seconds), the energy required for each measurement may be significant. This is no problem if sufficient electrical energy is available (in a vehicle, for example) but may be a limiting factor for back-pack battery operated prospecting equipment in the field.

1.4 Is a solenoid the best configuration?

Most amateur PPM's¹ (proton precession magnetometers) have been solenoidal; a solenoidal coil showing the magnetic field generated by it if a current flows through it is shown in Figure 1.1.

This configuration has one great advantage and several great disadvantages. The advantage is that it is very easy to wind a solenoid. The first disadvantage is that it is sensitive to external ac magnetic fields of which the most troublesome is the very large 50 or 60 Hz field over the surface of the earth. These will induce voltages which may be much, much greater than the desired proton precession voltage. Secondly, the sensor is orientation sensitive. If the polarizing field (along the axis of the solenoid) happens to be in the same

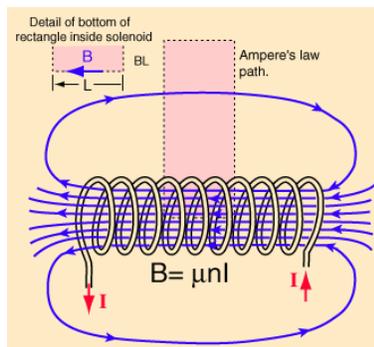


Figure 1.1: A solenoidal coil

direction as the external field which it is desired to measure, the induced precession voltage will be zero! This means that the user must always be sure that the solenoid is oriented correctly in order to get a sufficient induced voltage to measure. This may be a problem in magnetometers which are towed behind boats, for example, while turns are being made. The sensitivity of the instrument depends on orientation with respect to the earth's field. The induced signal strength is proportional to the sine of the angle between the earth's magnetic field and the axis of the solenoid. The first disadvantage of the solenoid, its

¹See, for example, The Amateur Scientist column, Scientific American, February, 1968 for a design which has been the basis of many later designs

sensitivity to external ac magnetic noise fields, can be ameliorated by making a double solenoid, with one winding in the opposite direction to the first. If they are very close together, the induced voltages due to external ac noise will be roughly the same and the two solenoids can be connected together so that the noise from each cancels resulting in a noise voltage which will be greatly reduced. However, the orientation problem still exists.

Many professional PPM's² are made using a toroidal core. A toroidal coil is shown in Figure 1.2

A toroidal coil is not very sensitive to external ac noise magnetic fields and so this reduces external noise greatly. Secondly, the toroidal configuration is not nearly as orientation sensitive. If the external magnetic field it is desired to measure is oriented in the least sensitive direction (in the plane of the toroid), the induced voltage is only reduced to one half of that in the most sensitive orientation (where the external field is aligned along the axis of the toroid). The signal *never goes to zero* and the worst case signal-to-noise ratio is still half that of the best. The only disadvantage of the toroid is that winding it is very labour intensive and difficult.

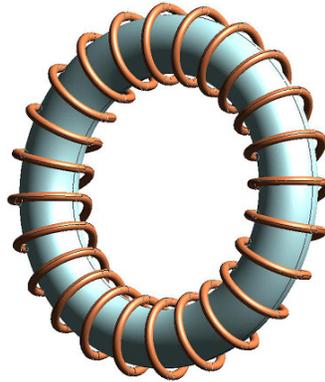


Figure 1.2: A toroidal coil

Finally, a toroid may be easily shielded electrostatically because the magnetic field is totally inside the winding. Therefore, a close fitting metal shield may be placed around the toroid. A metal shield may also be placed around a solenoidal coil but, in order to prevent it from acting as a shorted turn, it must be very much larger than the outside diameter of the solenoidal coil. This may be inconveniently large.

Since you only have to wind the core once while you use it many times, it seems to me that the extra work in making a toroidal PPM is more than repaid by its advantages - in my opinion, the toroidal configuration ought to be first choice.

²F. Prindahl, Scalar Magnetometers for Space Applications, Measurement Techniques in Space Plasmas: Fields, Geophys. Monograph 103, American Geophysical Union, 1998 - has many useful references. The paper is a useful summary of all the techniques for measurement of B in spacecraft - an environment where power and size are constrained. Therefore this summary is useful for those who are interested in portable instruments where, in general, the same constraints exist.

1.5 Signal strength calculation

I am a physicist and so I am interested in understanding exactly how the PPM works. I am therefore going to derive a simplified signal strength calculation using a model, for a toroidal sensor, which is physically intuitive and instructive but which contains some approximations. I am indebted to Dr. Fritz Primdahl for giving me a copy of his lecture notes for a presentation at the Danish Space Research Institute, 17 November, 1987. This derivation is based on these notes. There are more rigorous calculations which I will describe in a later section which, although mathematically correct, are much less illuminating.

Consider a sample of a proton rich substance in the presence of a polarizing field of magnitude, B_p . The magnetic energy of a proton in this field will be μB_p which is a scalar quantity. Here, μ is the magnitude of the magnetic moment of the proton. The thermal energy of the proton will be $\sim kT$ where k is Boltzmann's constant (1.38×10^{-23}) and T is the temperature of the fluid in Kelvins. Typically, μB_p will be many orders of magnitude smaller than kT which means that the thermal agitation of the protons prevents them from being completely aligned. For a typical PPM, a typical value of the ratio of the polarizing magnetic energy to thermal energy might be $\sim 10^{-7}$ so the fluid is only very slightly magnetized when it is being polarized.

One can think of this partial magnetization in one of two ways - either each and every proton is only partially magnetized by this amount or that this fraction of the protons is completely magnetized and the remainder are completely random. The latter way of thinking is convenient in the analysis that follows. Please note, however, that we have made an approximation in estimating the ratio of magnetic energy to thermal energy.

Let us now calculate the total magnetization of the fluid. Let the density of the fluid be $d \frac{kg}{m^3}$. Let the molecular weight of the fluid be w amu where each amu is equal to 1.67×10^{-27} kg. Let there be m hydrogen atoms in each molecule. Then the number, N , of protons per cubic meter of the fluid is:

$$N = \frac{md}{1.67 \times 10^{-27}w}$$

Putting in the appropriate numbers for water ($d = 1000$, $w = 18$, $m = 2$) for example, gives an N of 6.65×10^{28} protons per cubic meter. If all the protons were perfectly aligned with one another, the total magnetization (magnetic moment) would be $N\mu$. However, because only the fraction, $\frac{\mu B_p}{kT}$, is aligned, the total magnetization, M , will be approximately:

$$M \simeq N\mu \left\{ \frac{\mu B_p}{kT} \right\}$$

M is in units of $A - m^2$ per cubic meter or $\frac{A}{m}$. Now, let us consider that this fluid is inside a toroidal form which is wound with n turns of wire. This coil is used to generate the polarizing field and also to sense the induced voltage due to all the protons as they precess. The central polarizing field will be generated

by a polarizing current passed through the winding. A cross-section of such a toroidal coil is shown in Figure 1.3.

For a toroidal coil, the central magnetic field strength is related to the polarizing current by the equation given below where I_p is the polarizing current, n is the number of turns and R is the central radius of the toroid:

$$B_p \simeq \mu_0 \frac{nI_p}{2\pi R}$$

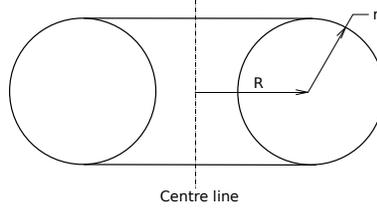


Figure 1.3: Toroidal coil cross-section

The net magnetization in the fluid is then given by this value of polarizing field, B_p , substituted into the equation for M . M will be aligned along the central core of the toroid. The above equation is an approximation since the magnetic induction inside a toroidal core varies with the radius being greater than this closer to the inner side of the toroid and weaker than this near the outer side. However, for toroids which are thin compared to their outside diameter, the differences between the inner magnetic field strength and the outer magnetic field strength is very small so the above value is a very good approximation to the average field strength inside the toroid.

Let us now suppose that the external field which we wish to measure is directed along the axis of the toroid (imagine the toroid as a wheel - then the axis is in the direction of the wheel axle). Then, when the polarizing field is switched off, all the aligned protons will precess about this axis and each one will generate a tiny voltage in the winding according to Faradays Law. The amplitude of the precessing magnetic moment per unit volume is M at the instant the polarizing field is switched off and so the subsequent axial component of M may be written as a function of time:

$$M_{axial} = M \cos \omega t$$

The magnetic field (induction) from this fluctuating magnetic moment is thus:

$$B_{axial} = \mu_0 M_{axial} = \mu_0 M \cos \omega t$$

The total magnetic flux, ϕ , inside the core is just $B_{axial}A$, and from Faradays Law, the time varying induced voltage in n turns of wire surrounding this flux is given by:

$$e(t) = -n \frac{d\phi}{dt} \simeq -nA \frac{dB_{axial}}{dt} \simeq nA\mu_0\omega M \sin \omega t$$

where A is the cross-sectional area of each loop of the coil. Since the radius of the toroid core is r , then $A = \pi r^2$. The rms (root mean square) amplitude of the induced voltage is just the amplitude in the above equation divided by the square root of two. Therefore, making all the appropriate substitutions, the

rms output voltage, e_s , at the Larmor frequency will be:

$$e_s \simeq \frac{\omega N (\mu\mu_0 nr)^2 I_p}{2\sqrt{2}kTR}$$

Note that the output voltage depends on the magnitude of the polarizing current and on the square of the number of turns. It also depends on the geometry of the toroid since both r and R are in the equation.

1.6 How bad are the approximations?

One approximation used in the section above was that the average value of the magnetic induction inside a toroid due to a current flowing through it was equal to the value at the centre of the toroidal core. It is, as stated before, stronger at smaller radii and weaker at larger radii and the average value is not exactly equal to the value at the centre. To get a proper value, one needs to integrate the elemental values of the magnetic induction over the entire core and then calculate the average. This has been done (see the next section). However, the error in the approximation is not too significant as long as $r \ll R$.

Secondly, in considering the problem of magnetizing the medium, the previous section derives an equation relating the magnitude of the total magnetic moment of the medium due to a polarizing field and this contained an approximation:

$$M \simeq N\mu \left\{ \frac{\mu B_p}{kT} \right\}$$

In practice, it is possible to measure the magnetization of a substance experimentally and the equation used is:

$$\mathbf{M} = \chi \mathbf{H}$$

χ is called the magnetic susceptance of the substance and the equation above gives the relationship between the resulting magnetization, \mathbf{M} , produced by a substance by an external magnetic field, \mathbf{H} . For non-ferromagnetic substances, this equation can be rewritten (in terms of the magnitudes since the directions of the vectors is the same):

$$M = \frac{\chi B_p}{\mu_0}$$

which means that, using the approximate value for B_p , the following equation should be true.

$$\chi \simeq \frac{N\mu^2}{kT} \mu_0$$

Well, χ is a measured quantity and the values for a large number of substances (unfortunately in cgs units!) can be found in the Handbook of Physics and Chemistry. For example, the value for water at 300 Kelvins is 4.26×10^{-9} (converted to SI units). The product of all the numbers on the right side of the equation which I derived approximately is 4.02×10^{-9} - an apparent difference of just about 6%. This is not bad considering the broadness of the

approximation! However, the close agreement might be too good to be true. The other atoms in any molecule have some magnetic susceptibility also and so the bulk measurement of the magnetic susceptibility of some substance will include those as well. In the case of water, the other atom is oxygen and for other commonly used hydrocarbons, there are carbon atoms as well. Therefore, in conclusion, it is not possible to (easily!) take into account the effect of other atoms and so using the measured bulk values of susceptibility is, itself, just an approximation!

So, given the fact that I derived an equation on physical grounds which I stated were just approximate and that the measured values are also not really appropriate, exactly what value should one use? There is a good answer to this question because it is possible to correctly evaluate that part of the total magnetic susceptibility of a substance that is due to the hydrogen atoms (i.e., protons) alone. *This is really what we want!* ... and it is given by the Curie equation:

$$\chi = \left\{ \frac{j + \frac{h}{2\pi}}{3j} \right\} \frac{N\mu^2}{kT} \mu_0$$

where j is the angular momentum of the proton and h is Planck's constant. For a proton, j happens to be equal to $\frac{h}{4\pi}$ and so the quantity inside the first set of brackets is exactly equal to 1. Therefore, my approximate equation turns out to be exactly correct! This means that it is possible to calculate the appropriate value for the magnetic susceptibility to use in signal strength formulas for any substance provided you know how many hydrogen atoms there are in each molecule, the molecular weight of the substance and the density of the substance. For water,

$$\chi = 4.015 \times 10^{-9}$$

in SI units. The values for some common liquids, relative to that of water, is given in the table below:

Substance	Molecular Weight	H atoms per molecule	Density	Chi relative
Benzene	78	6	876.5	0.607
Methanol	32	4	791.4	0.890
Ethanol	46	6	789.3	0.927
Isopropanol	60	8	780.9	0.937

1.7 A more detailed study of the toroidal configuration

In a paper a number of years ago, Acker³ did a detailed study of the toroidal PPM sensor and derived an equation taking into account the variation of magnetic induction over the cross-section of the toroid as discussed above. He also took into account the effect of an external magnetic field being in a direction other than the axis of the toroid. He also generalized the problem to include toroids with elliptical cross-sections. This is a very sophisticated analysis and

³Acker, F.E., Calculation of the Signal Voltage Induced in a Toroidal Proton-Precession Magnetometer Sensor, IEEE Trans. Geosci. Electronics, GE-9, 98-103, 1971

includes a section on the cross-sectional shape of the windings when the core is so completely wound that the central hole is filled with wire. For the purposes of this discussion, we shall assume that the toroidal cross-section is circular and then just produce his equation for the rms sensor voltage:

$$e_s = \frac{\chi\omega\mu_0}{2\sqrt{2}} (n^2 I_p) \left(R - \sqrt{R^2 - r^2} \right) (2 - \sin^2 \alpha)$$

This is essentially the same equation as derived in section 1.5 except that the term, $R - \sqrt{R^2 - r^2}$, replaces the previous approximation, $\frac{r^2}{2R}$. (It is left as an exercise for the student to show that these two terms approach one another when $r \ll R$.) The other addition is the quantity, α , which is the angle between the measured magnetic field and the axis of the toroid. In my derivation, α was assumed to be zero. Since α can vary between zero and 90 degrees, the quantity in the brackets can vary from 2 to 1 respectively. This means that the weakest sensitivity (when α is 90 degrees) is equal to one-half of the strongest signal. The signal never goes to zero - as it can with the solenoidal configuration.

1.8 The solenoidal configuration

The solenoidal configuration has been studied by Faini and Svetlo⁴. A longitudinal cross-section of a solenoidal sensor is show in Figure 1.4.

The situation is more complex, mathematically, because the magnetic field inside a solenoid of finite length is not everywhere parallel to the axis of the solenoid. Calculating the net induced voltage requires integration over a geometry which is not simple. Nevertheless, it is possible and they derived the following equation:

$$e_s = \frac{\chi\omega\mu_0}{\sqrt{2}} \left(\frac{n^2 I_p}{b^2} \right) \eta V_i \sin(\alpha)$$

where b is the length of the solenoid, V_i is the internal volume of the solenoid (the inside cross-sectional area times the length) and α is the angle between the field being measured and the axis of the solenoid. It is assumed that the entire volume within the solenoid is filled with the working liquid; if it is not then a correction for the percentage which is filled must be applied. Note that if α is zero (i.e., the magnetic field being measured is exactly

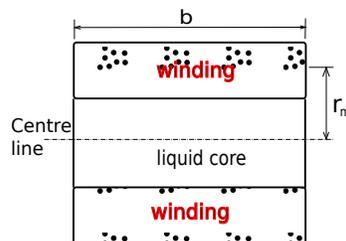


Figure 1.4: Solenoidal sensor

⁴Faini G. and Svetlo O., Signal-to-Noise Considerations in a Nuclear Magnetometer, Suppl. al Vol 23, Serie 10 del Nuovo Cimento, 55-65, 1962

parallel to the long axis of the solenoid), the signal amplitude goes to zero. The quantity, η , is called the filling factor and depends on the geometry of the coil. It is the calculation of this filling factor which is mathematically complicated.

For an infinitely long solenoid (i.e., in practice, a solenoid that is much, much longer than its diameter), the filling factor is unity. For a short solenoid, it is less than unity and has a value which depends on the ratio of the mean solenoid radius to the solenoid length and which they just showed in in a graph. The graph looks roughly exponential and, indeed, I have found that the equation below is a fairly good fit to it:

$$\eta = e^{-1.42a}$$

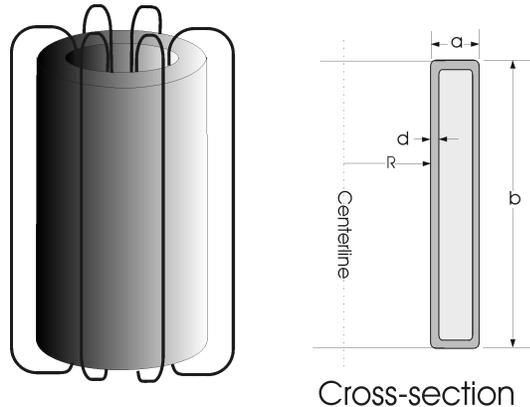
where a is the ratio of the mean radius of the solenoid, r_m , to its length. I have made a slight change of notation from the original in their paper. They gave the symbol α to the ratio of mean radius to length; I have changed it to the letter, a , in order to avoid confusion with the angle between the solenoid axis and the local magnetic field.

On the practical side, the authors considered the one of the consequences of having a multi-turn solenoidal sensor which was operating continuously; the heat generated by the losses during polarization. This heat is only dissipated through the surface of the winding. They found that the optimum shape of a solenoidal sensor was one in which a was just about 0.4 and the ratio of the thickness of the multi-turn winding to the inside radius of the core was just 0.2.

1.9 The cylindrical solenoid

There is another configuration which is potentially useful for PPMs designed to be towed behind boats. It is a toroid with rectangular cross-section (see Figure 1.5) and of arbitrary length and as with the ordinary toroid, the worst-case orientation still gives you a signal of half the best-case orientation. It has the advantage that the length can be made quite long so as to increase the quantity of active fluid. I have not seen this configuration described in the literature and so have derived the signal strength below for the case where the local magnetic field is along the axis of the cylinder.

The magnetization produced by the polarization field inside a toroidal



R is the inside radius
 d is the wall thickness
 b is the outside cylinder height
 a is the cylinder thickness

11 Figure 1.5: Cylindrical toroid

form is:

$$\mathbf{M}(r) = \chi \mathbf{B}_p(r)$$

where H_p is the polarization magnetic field which is a function of radius. I

have shown both \mathbf{M} and \mathbf{B}_p with a bracketed r to show that they are functions of r . We have discussed the valuation of χ earlier. \mathbf{B}_p then has a magnitude given by:

$$B_p = \mu_0 \frac{nI_p}{2\pi r}$$

where I_p is the polarization current and n is the total number of turns. The total flux inside the toroid is given by the integral:

$$\phi = \int M dA$$

where dA is the elemental area in a strip of constant r inside the cylindrical toroid. Thus ϕ becomes:

$$\phi = \chi(b-2d) \int B_p dr$$

integrated over r from $R+d$ to $R+a-d$. Putting in the value for B_p , we get:

$$\phi = \frac{\chi\mu_0(b-2d)nI_p}{2\pi} \ln \left[\frac{R+a-d}{R+d} \right]$$

Again, the generated rms voltage can be shown to be:

$$e_s = -n \frac{d\phi}{dt} = \frac{\chi\mu_0\omega}{\sqrt{2}} n^2 I_p \frac{(b-2d)}{2\pi} \ln \left[\frac{R+a-d}{R+d} \right]$$

In most cases, d will be much less than b and a will be much less than R so, for those cases, the equation can be approximately written as:

$$e_s \simeq \frac{\chi\mu_0\omega b n^2 I_p}{2\pi\sqrt{2}} \ln \left[\frac{R+a}{R} \right]$$

1.10 Summary of signal strength calculations

Looking at the equations for the signal strength for the toroid, the solenoid and the cylindrical solenoid, we see that there is a number of common constants in the equations for rms voltage amplitude. We can replace these by a single constant, κ , given by:

$$\kappa = \frac{\chi\mu_0\omega}{\sqrt{2}}$$

This term will vary with the value of the local magnetic field since ω is the angular precession frequency but, for a typical value over most of Europe and

North America, the local magnetic field has a magnitude of around 50 microT and, using the value of χ for water, κ will be approximately 4.77×10^{-11} .

For the case where the orientation of the sensors is that which gives the maximum signal, we can rewrite the equations for these three configurations in a general form as:

$$e_s = \kappa G n^2 I_p$$

where G is different for each of the three configurations. They are given in the table below

solenoid	$\frac{\eta V_i}{b^2}$
toroid	$R - \sqrt{R^2 - r^2}$
cylindrical toroid	$\frac{b-2d}{2\pi} \ln \frac{R+a-d}{R+d}$

1.11 Signal-to-noise considerations

The equations discussed in the previous section give the rms signal output voltage one can expect from the sensor. However, because of its resistance, the sensor will also produce an rms random noise voltage which has a magnitude given by:

$$e_n = \sqrt{4kTR_w\beta} = 1.29 \times 10^{-10} \sqrt{R_w\beta}$$

where R_w is the resistance of the wire used in the sensor and β is the bandwidth of the subsequent electronic detection system. The right hand side of the equation has been evaluated for a temperature of 300 Kelvins.

It is the ratio of signal-to-noise that is important, not just the magnitude of the signal. This ratio is just:

$$\frac{S}{N} = \frac{e_s}{e_n}$$

and, for the types of sensors discussed here, you just put in the appropriate value for e_s .

The wire used to make sensors is often labelled “magnet wire”. The term means that the insulation around the wire is just a thin coat of, traditionally, enamel or, more often, a thin plastic coating of some sort. The wire resistance will depend on the material used to make the wire and its dimensions. Since the wire will almost certainly be made from copper which has a resistivity, in SI units, of 1.73×10^{-8} , one may easily calculate the total resistance of the coil by:

$$R = \frac{1.73 \times 10^{-8} \lambda}{A}$$

where λ is the length of the wire making up the coil and A is the cross-sectional area of the wire. The length of the wire depends linearly on the number of turns of wire making the coil and its diameter. The cross-sectional area depends on the wire size.

We may rewrite R_w as:

$$R_w = \frac{3.46 \times 10^{-8} nr}{r_w^2}$$

where r_w is the radius of the wire and r is, for the toroid, the radius of the core and, for the solenoid, the radius of the solenoid core. n is the number of turns. For multi-turn windings, a mean value for r should be used.

In addition to this dc resistance, there is the skin effect which has to be taken into account. The skin effect gets its name from the fact that ac currents flowing in a wire do not flow with equal density in all parts of the wires cross-section. The higher the frequency, the more the current is concentrated near the outside or skin of the wire. At radio frequencies, this effect is very pronounced and the actual resistance of a wire at these frequencies is much, much greater than its dc resistance. At the low frequencies of the proton precession signal due to the earth's magnetic field (about 2 kHz), the effect is small but not negligible. Faini and Svelto (ibid.) found that, at 2 kHz, the actual resistance of one experimental solenoid was about 20% greater than the dc resistance. This factor, although just the value for that particular solenoidal coil, is probably a reasonable estimate for the kinds of coils likely to be wound for PPM sensors. We may therefore take into account the skin effect by replacing the number four under the square root sign of the equation for rms noise voltage with a five. This extra 20% in resistance gives about a 10% increase in rms noise voltage.

At a temperature of 300 Kelvins (about room temperature), the noise voltage thus becomes:

$$e_n = 2.68 \times 10^{-14} \frac{\sqrt{nr\beta}}{r_w}$$

Finally, note one caution. For a PPM which is used almost continuously and which has high currents flowing through it during the polarization periods, the temperature may rise to be considerably greater than 300 K.

Some high-grade professional magnetometers use aluminum wire. Aluminum has the advantage that the weight of the sensor is greatly reduced compared to a sensor made with copper wire. In addition, copper wire sometimes has ferrous impurities and these can effect the accuracy of the measurement of B if the amount of these impurities is significant. Aluminum wire does not often have ferrous impurities. Unfortunately, aluminum magnet wire is difficult to find. Also, some care must be taken in making connection to aluminum wire; it is difficult to solder. Nevertheless, if aluminum magnet wire can be found, it is preferable to copper. If any reader of these notes knows where a hobbyist can find small quantities of aluminum magnet wire, the author (jark@shaw.ca) would be grateful to hear of it!

1.12 Optimization

If one just wants a large signal to noise ratio with no other considerations, then the equations for signal amplitude and noise indicate that you need to:

1. have a large sensor (i.e., lots of the proton rich material inside the windings),
2. use a very large polarizing current,
3. have a very small detection bandwidth and,

4. use large diameter wire to minimize the noise voltage by lowering the resistance of the winding.

For example, one might design a toroidal PPM sensor with $R=30$ cm, $r=7.5$ cm, wind it with 3 layers (2036 turns) of wire with a diameter of 2.0 mm. If the polarizing current were 10.0 Amperes, the rms signal voltage due to the earth's magnetic field in the most favourable orientation would be about 19 microV. The noise voltage would only be about 3.4 nV (assuming a bandwidth of the detector of 100 Hz) giving an outstanding S/N ratio of about 5600!

Of course, this would be an extremely big (and heavy!) toroid containing 8.33 kg of fluid (if the fluid is water) in the core and it would require 27.9 kg of copper wire to wind it. The dc resistance of the winding would be about 5.5 Ohms so to get a polarizing current of 10 Amperes through it, you would need to have a 55 Volt supply. The unit would therefore consume about 550 W for the period (a few seconds) that it was being polarized. Such a unit would obviously be very impractical for a portable instrument (or, for that matter, for an observatory instrument). Imagine carrying about 36 kg of sensor plus 55 V of batteries capable of supplying more than half a kW for the duration of each polarization and with sufficient capacity for the many measurements one might like to make in a survey lasting several days!

So the question is: what is the optimum configuration (size, number of turns, etc.)? To answer this, one needs to decide in what way it needs to be optimized. Faini and Svelto (ibid.) spent some effort in calculating how one would need to optimize the solenoidal sensor if the principal consideration was heat dissipation in a continually operating unit. In this case, the heat conducted away depends on the outside surface area of the solenoid and, when they had gone through the appropriate equations, they came up with sensors which were rather short - only roughly two times the mean radius - even though this results in a rather small filling factor.

For a back-packed sensor, one of the critical factors is the total energy required to make one measurement since the energy source is a battery. One might therefore come up with a figure of merit, FOM, for such a sensor which is just the ratio of signal-to-noise divided by the power dissipated:

$$FOM = \frac{\frac{S}{N}}{I_p^2 R_w}$$

For a sensor designed to be towed behind a vehicle, power consumption is not as critical a concern and some other factor such as size and difficulty in construction might dominate the other considerations. In the "olden days" (before computers), one would write an expression for the desired figure-of-merit and take the differential with respect to the desired variable and equate it to zero in order to find the maximum value.

However, in a practical case, there may be many such considerations - the size of the core, the availability of wire and many other factors - which one might want to vary and the resulting equations become very cumbersome. Nevertheless, there are computer programs (e.g., Mathematica, Maple) for writing

symbolic expressions which allow one to write these complicated expressions, do the differentials and solve for the maxima. In my experience, with modern computers and software, it is much easier to just make up a spreadsheet. You can then vary many parameters and look at the resulting S/N's. I have done this, using Excel, for all three configurations. The file is named ppm.xls. This spreadsheet has six pages; two for each of the three configurations, the two being in metric or Imperial (lb-ft-sec) measure.

Most of the world measures wire by its diameter in mm but in the UK and parts of the British Commonwealth, Standard Wire Gauge (SWG) is commonly used to express wire diameter rather than an explicit metric size. In North America, American Wire Gauge is used (AWG); SWG and AWG differ slightly. In these spreadsheets, I include a column showing the approximate AWG for the wire sizes.

In these spread sheets, I have assumed that the polarizing current comes from a battery and so the value of I_p is just the battery voltage divided by the dc resistance of the winding. This battery voltage is then one of the variables in the spreadsheet. I have also, for the toroidal and solenoidal configurations, not made the number of turns one of the variables but, instead, the number of layers of turns. The spreadsheet then calculates the approximate number of turns and the total length of the wire needed. When winding either a solenoid or a toroid, it is most convenient to wind the coils as complete layers rather than to count turns. For the toroid, a layer is defined as when the inside of the winding is close wound; the outside is therefore more loosely spaced. The cylindrical toroid spreadsheet does not use layers of wire but rather the number of turns.

What sort of accuracy can one expect using these calculations? Faini and Svelto (ibid) state that their experimental values of S/N for a particular solenoidal PPM sensor agreed to about within 10% of the calculations. Presumably, one could expect similar agreement here. However, it is one of the maxims of experimental physics that things always turn out for the worst - that is, all the little unaccountable errors and disregard for second order effects, etc. are always in such a direction as to make the desired quantity less measurable. In the case of winding multi-turn toroids and solenoids, it is very difficult to actually get the number of turns one expects for a close wound layer. Secondly, the actual volume of the proton rich fluid will be less than the inside volume of the solenoid or toroid form because of the thickness of the walls. For these reasons, it would be prudent to discount the calculations by some factor - say 20% in order to be conservative.

1.13 Running the spreadsheets

The spreadsheet is basically used to see what effect on the S/N ratio is caused by changing wire size, number of layers of wire, dimensions of sensors, etc. There is a separate pages in the spreadsheet for each of the possible configurations in either metric or Imperial measure. Let us take, as an example, how one might use the spreadsheet for toroidal sensors in metric measure. Each spreadsheet is

similar in how it is used.

The spreadsheet itself is organized in rows of increasing wire diameter. The user only puts numbers in the small area, with a coloured background, in the bottom left of the spreadsheet. The various columns then give you the weight of wire needed, the signal voltage, the noise voltage and the signal-to-noise ratio. At the bottom right, there is a small box which gives some calculated values such as the volume of the fluid in the core, the G factor and the κ factor. The user should not type anything in these boxes.

If you put the values for the toroidal sensor mentioned in the previous section into the coloured area, you will see the value for the S/N described there. If you replace the values for toroid central core radius or the toroidal cross-sectional radius, you will see the table recalculate itself and new values will appear for all the derived values. Usually, you are only interested in the right-most column which is signal-to-noise ratio. For example, replace the toroidal radius with 10 cm and the toroidal cross-sectional radius with 2.5 cm, the number of layers with 4, and the battery voltage with 12, you will see the maximum S/N in the bottom right corner go from its previous value of about 5600 to 1081 but the resistance of the sensor using 2.0 mm diameter wire is far too low resulting in a needed polarization current of 16.1 A which is, probably, unacceptably high. But, if you look at the row with wire of 1.0 mm diameter, you see a polarizing current of only about 1.9 A results in a signal-to-noise ratio of about 227.

If you have access to aluminum magnet wire, you should replace the two values for resistivity and density of copper with those of aluminum. I have tabulated those respective values in the bottom centre of the spreadsheet.

For the spreadsheets in Imperial measure, the numbers are put into the darker blue boxes at the right of the blue boxes used for metric measure. *NOTE* ... put numbers *only* into these darker blue boxes as the spreadsheet then calculates the equivalent metric measurements and inserts them automatically into the light blue boxes; the rest of the spreadsheet is then unchanged. Also, in the Imperial measure spreadsheets, the wire sizes are shown in integral AWG sizes and the equivalent metric sizes are then calculated for the appropriate column.

2 The Measurement of Magnetic Field Magnitude

2.1 The basic operation of a proton-precession magnetometer

A block diagram of a PPM is shown in Figure 2.1.

A PPM runs in two basic modes; the first of these is the polarization mode in which the working substance is subjected to a strong field in order to magnetize (i.e., line up) the protons. The second mode is the actual measurement of the precession frequency in order to determine the external magnetic field strength. As we have seen, both modes are normally done using the same winding as both the polarizing electromagnet and as the precession sensor. The function of the controlling electronics is to switch between the two modes. The remainder of the electronics then performs the two functions. In practice, the division between the three major blocks may not be as clear as shown in the block diagram (the polarizing and control electronics may be essentially the same circuit) but they are, in principle, separate functions.

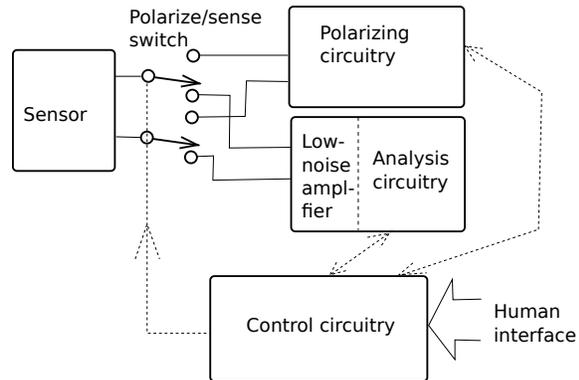


Figure 2.1: Block diagram of a PPM

2.2 Measurement accuracy

Before considering the circuitry in more detail, let us discuss in a general sort of way exactly what we want to measure and how best to do it. The basic quantity to be measured is the frequency of the precessing protons. This appears as an ac audio-range voltage at the terminals of the sensor. This voltage has to be amplified and the mere presence of the amplifier will add more noise. Obviously the amplifier needs to have as low a noise factor as possible.

Then, the frequency has to be measured as accurately as possible. Since the signal decays exponentially with time (typical time constant is a few seconds or less), the measurement period is limited. The signal-to-noise ratio calculated in the previous chapter is the value at the instant the polarizing field is turned off. Since the signal amplitude decays, the signal-to-noise ratio gets worse as time goes on.

Finally, we need to consider exactly what we want to measure. In a magnetometer, we want to measure the absolute value of the magnetic field. In a gradiometer, on the other hand, we just want to measure the difference in

frequency between two sensors and that may be entirely a different problem. It is intuitive that the ultimate accuracy in measuring the frequency must depend in some way on the signal-to-noise ratio. So let's consider what the theoretical limitations are - we can then discuss ways of making the measurements later and see how the closely various techniques come to the theoretical limit. Let the amplitude of the amplified sensor output be given by:

$$e(t) = A \sin(\omega t) + \frac{a}{\sqrt{2}} \{ \cos(\omega t) + \sin(\omega t) \}$$

where A is the amplitude of the desired signal and a is the amplitude of the noise (the sum of the sensor noise plus the additional noise added by the amplifier). $\frac{A}{a}$ is the actual signal-to-noise ratio achieved by the system. I have shown the noise as the sum of an in-phase component (the sine term) and an quadrature component (the cosine term) because the noise will have a random phase with respect to the desired signal. Let's assume, further, that A is much greater than a - that is, that the final signal-to-noise ratio is reasonably high (the following analysis will be fairly accurate if it is greater than about 10). Then, the above equation can be rewritten:

$$e(t) \simeq \sqrt{A^2 + a^2} \sin(\omega t + \phi)$$

where

$$\phi = \arctan\left(\frac{a}{\sqrt{2}A}\right)$$

Now, from the measurement point of view, this random phase added by the noise can be thought of as causing a change in the period of the sine wave since, in measuring frequency, we are really measuring zero crossings. This is illustrated in Figure .

In this figure, I have shown the phase error, ϕ , as a crossing time error, ∇t . The two are related by:

$$\nabla t = \frac{\phi}{\omega}$$

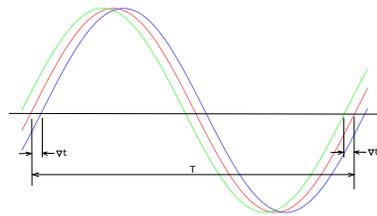


Figure 2.2: Sine wave with phase changes

Now, let us suppose that we are going to measure the frequency over some interval of time, T . There will be an uncertainty in the first zero-crossing of ∇t at the beginning of the interval and a similar uncertainty of ∇t at the end of the interval. The standard

deviation corresponding to plus or minus one unit is about $2/3$ of a unit and, since there are two such uncertainties (one at the front and one at the end), the overall standard deviation will be about 2 times this $2/3$ of a unit - altogether about one unit of ∇t . The relative error, therefore, in measuring the frequency

will be $\frac{\sqrt{t}}{T}$. This is the error due to the noise alone and is really the theoretical best that can be done. Other errors will add to this.

How big is this minimum error in a realistic case? A table of accuracy versus S/N is shown below.

S/N	Accuracy	Error in 50 μ T (nT)
1	9.2×10^{-5}	4.6
2	5.1×10^{-5}	2.5
5	2.1×10^{-5}	1.1
10	1.05×10^{-5}	0.5
20	5.28×10^{-6}	0.26
50	2.11×10^{-6}	0.11
100	1.06×10^{-6}	0.05
200	5.28×10^{-7}	0.03

As an example, suppose a frequency of approximately 2130 Hz was measured over a total period of one second and that, at the end of this interval, the overall S/N of the system was 10 (it will be better at the beginning because the signal decays exponentially). Then, $\frac{\sqrt{t}}{T}$ works out to about 1.05×10^{-5} . If the field being measured were 50 μ T, then the error would be just ~ 0.5 nT! This looks really good since, from the S/N calculations in the previous chapter, it looks like S/N values of ~ 10 are fairly easily obtainable.

However, now we need to tackle the subject of actually measuring the frequency; what techniques are there and what has been done by others? First, let us consider the obvious (and simplest) method and that is to measure the frequency using a frequency counter. If we had a measurement interval of one second, the count (assuming the frequency to be 2130 Hz) would be 2130 plus or minus 1. The plus or minus one means a standard deviation of about 2/3 of a unit so the relative error would be 2/3 parts in 2130 corresponding to an average error of one part in about 3000. This corresponds to about 15 nT in a 50 μ T field. The error for any individual measurement will be larger than the rms average error; it would be plus or minus one part in 2130 which is about plus or minus ~ 25 nT in a 50 μ T field. This is pretty poor.

Of course, we could measure the frequency over a longer period but we are limited to the fact that the signal decays exponentially with a time constant of a few seconds so periods of greater than a few seconds are just not feasible. Suppose, we had sufficient S/N so that the signal was still well above the noise after ten seconds. Then our average accuracy would be ten times better (about 1.5 nT) and the error for an individual measurement would be plus or minus about 2.5 nT. This is still not very good; it is about 100 times worse than the ultimate limit of 25 pT for a S/N ratio of 100 measured over 1 second and, the measurement would take 10 seconds in addition to the polarizing time.

So, just using a counter to measure the precession frequency is not going to be very good. The standard method for improving the accuracy of frequency measurement for low frequency signals is not to measure frequency but to measure period instead. In principle, you can think of this as being done in the following way. Suppose we take our magnetometer signal and use it to close and

open an electronic gate. Suppose the gate is opened at the rising edge of the signal sine wave and closed at the next rising edge. This gate connects a very high frequency clock to a counter. The counter will therefore count the number of clock pulses during one period of the signal sine wave. Of course, the count is only accurate to plus or minus one count but, if the clock frequency is very high, the relative error can be made very small.

We can make it even smaller but leaving the gate open for some integral number of signal periods so that the total count is greater but the error is still just plus or minus one count. It is easy to show that, if the gate is open for N periods of the signal which is at frequency f_p and the clock frequency is f_c , then the period being measured is accurate to one part in $\pm \frac{Nf_c}{f_p}$. For example, if the gate time is 2048 periods of a signal at 2130 Hz and the clock frequency is 1 MHz, then the period can be measured accurate to approximately one part in $\pm 10^6$. This corresponds to approximately ± 50 pT for a 50,000 nT field.

From this discussion, it appears obvious that the accuracy can be increased indefinitely simply by having a clock at higher and higher frequencies. For example, if the clock were at 100 MHz, the counting error would only be one part in about $\pm 10^8$. But, in order to achieve this low error in measuring the precession frequency, the previously discussed limit due to the signal-to-noise ratio at the end of the approximate one second interval would have to be about 10,000 to one! Therefore, there is no advantage to going to a higher clock frequency thinking that it will always increase the accuracy unless the signal-to-noise ratio at the end of the counting interval is also sufficient. In this example where the counting frequency is 1 MHz, the counting accuracy produces an error similar which is about the same as would be produced by a signal-to-noise ratio of about 100. Therefore, unless the signal-to-noise ratio is better than 100 at the end of the counting interval, there is no improvement gained in going to a higher counting frequency.

The actual measured frequency and hence the measured magnetic field will be measured by the number of counts of the clock during the N periods of the precession signal. The measured magnetic field magnitude is given by:

$$B = 23.4875 f_c \frac{N}{n}$$

where n is the number of counts made during N precession periods and f_c is the clock frequency. The quantities in the numerator are usually constant for a given system and so can be evaluated just once and then the local magnetic field is given by the equation:

$$B = \frac{\epsilon}{n}$$

where ϵ is just $23.4875 N f_c$. This method is the basis of most commercial magnetometers.

In the olden days this method of measuring the period was implemented in hardware using strings of logic and counter ICs. Nowadays, it can be done using microprocessors very easily and it is really the method of choice. The absolute error in this form of measurement depends on the stability of the clock. Since

even the most inexpensive crystal controlled clock available in hybrid circuit form will have a stability and/or accuracy of better than about one part in 10^5 , the ultimate measurement accuracy and hence magnetometer sensitivity will depend mostly on this stability. One part in 10^5 of the earth's field corresponds to about 0.5 nT.

Finally, I should mention one other method proposed as a gradiometer⁵. A gradiometer gives the difference between the magnetic field at two different sensor locations. If PPMs are used for both sensors, then the difference between the magnetic fields at the two locations will appear as a difference in frequency. The simplest way of observing a frequency difference is to observe the beat frequency between the two. The beat appears when the two signals are electronically added or multiplied together; what could be simpler than that? The answer is not much - but how sensitive is this beat frequency to differences in magnetic field?

The beat frequency appears as a modulation of the signal from the two magnetometers. The signal decays exponentially with time but, superimposed on this, is a rise and fall of amplitude at the beat frequency rate. A difference of 1 Hz in frequency between the two gives a modulation of one beat per second. If the difference is very small, say 0.1 Hz, then the beat will take 10 seconds. So the above question becomes the following; how slow a beat can one detect in a signal which is decaying towards zero with a time constant of a few seconds? This is a matter of judgement but, in my experience, it would be very difficult to detect a beat any slower than about 5 seconds. Using this number, the minimum detectable frequency difference would be about 1/5 Hz. As we have seen earlier, a change of 1 Hz corresponds to about 23.5 nT so a frequency difference of about 1/5 Hz corresponds to a magnetic field difference of about 5 nT.

This is much poorer than the errors one can expect in a system where period is measured so the technique, while simple, is not really very sensitive. It simply does not take advantage of the inherent sensitivity of the PPM. Therefore, even in a gradiometer, it seems desirable to use the period measurement to get the ultimate in accuracy. In fact, for a gradiometer where we are measuring the difference in two periods, both measured with the same clock, the stability of the clock is no longer as important as it only has a second order effect on the measurement.

In summary, the theoretical limit of sensitivity is limited by the signal-to-noise ratio of the system where the noise is the sum of the sensor noise plus the amplifier noise. For reasonable S/N ratios, this theoretical limit is very small and the measuring errors in any practical system will be much larger. The most common system for measuring sensor frequency is one in which the period is measured using a high-speed clock. The absolute errors in measurement then depend mostly on the clock frequency (and its accuracy - except in a gradiometer) used in the period measuring system. Since stable frequencies of several MHz are easily produced by simple circuitry, a magnetometer or

⁵Delta Magnetometer by M.L. Dalton Research, 6035 Aberdeen Dallas, TX, 75230. Ph: (214) 691-4925 27

gradiometer for the earth's magnetic field which is accurate to at least ~ 1 nT is readily achieved.

There are refinements to the clock counting method of measuring frequency and there are other, more complex methods for accurately measuring frequency; these will be discussed in a later section of this paper.

2.3 Electronics functions

2.3.1 Control

The controller must switch the unit between its polarization and measurement modes and, supplementally, provide the interface with the operator. This latter function, at its simplest, consists of going through one cycle of polarization and measurement when directed by, for example, a push-button switch. In addition, the controller may also operate a display which gives the magnetic field output (or, in the case of a gradiometer, the difference between the magnetic field at two sensors) in some easily interpretable units. It may also provide some mechanism for storing results electronically. Indeed, the controller may easily be the most complicated part of a completed instrument even though its basic function of switching between two modes is very simple.

2.3.2 Polarization

The polarization part of the cycle is very simple in concept. A polarizing current must be passed through the coil in order to magnetize the substance. The substance takes some time to reach its final magnetization state - it approaches it exponentially with a time constant, called the spin-lattice relaxation time as shown in Figure 2.3.

T_1 is 2 to 3 seconds for water; about 0.5 seconds for kerosene, for example. In order to get close to full magnetization, the polarizing current needs to flow for several times this relaxation time. For example, after 4 time constants (say, 10 seconds, for water), the magnetization will have approached about 98% of its final value. After one time constant, it will only have approached 63% of the final value. Magnetizing the substance takes energy and the amount depends on the size and geometry of the sensor. Primdahl (ibid.) states, as a general rule of thumb, that it takes ~ 50 Joules per measurement. The actual number for any particular case can be found from the spreadsheet. Simply multiply the polarizing current, the battery voltage, and the polarizing time in seconds to get the energy required in Joules.

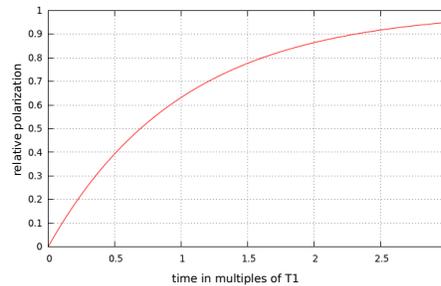


Figure 2.3: Polarization as a function of time

Remember that the longer the polarizing current is on, the closer the substance gets to its final value but there isn't much point in keeping it on for more than a few time constants.

After the polarizing current is switched off, the precession signal amplitude will decay exponentially with a time constant called the spin-spin relaxation time, T_2 . This is shown in Figure 2.4.

T_2 also depends on the substance and is about 2.1 seconds for distilled water with oxygen dissolved in it (i.e., normal distilled water). If the oxygen is removed, this time is increased to about 3.1 seconds. Oxygen may be removed from distilled water by bubbling dry nitrogen or helium through it or by boiling it and letting it cool down in an oxygen free environment.

There is one other critical factor to be considered. That is, one must turn off the polarizing current rapidly. If it is not, the protons will start to precess before the polarizing field has gone to zero and the final signal amplitude may be much reduced. How rapid is rapid? The answer is that, so as not to perturb the proton alignment, the polarizing current should go to zero in a time much shorter than the precession period.

The precession period is about 0.5 mS (1/2130 of a second) in the earth's magnetic field. Therefore, one would like to turn off the polarizing field in a time shorter than this - say, less than 0.05 mS which is 50 microS. Turning off the field that rapidly would be very simple if it were not for the fact that the polarizing coil has some inductance. The energy stored in its magnetic field is equal to $\frac{LI^2}{2}$ where L is the inductance of the sensor and this energy must be dissipated in that short turn-off period. If a switch or a relay is used to open the circuit and there is no other elements in the circuit, the inductive voltage spike will cause an arc which will dissipate some (most) of this energy.

It is best to somehow snub the voltage - this means to limit the magnitude of the inductive spike - and the usual way of doing this is to have a Zener diode across the coil. This diode must not conduct while the polarizing voltage is applied across the coil but should conduct at a voltage much lower than that required to cause an arc at the contacts. Typical Zener diode voltage ratings are 1.5 to 2 times the polarizing voltage. We will discuss this in more detail later.

If a semiconductor switch is used instead of a relay, then normally, the turn-off will be very abrupt (typically less than a microsecond) and the inductive voltage spike may or may not be absorbed in the semiconductor switch. The transient behaviour of the switch and its associated circuitry then becomes important. The nature of the switch and how it is turned off are important

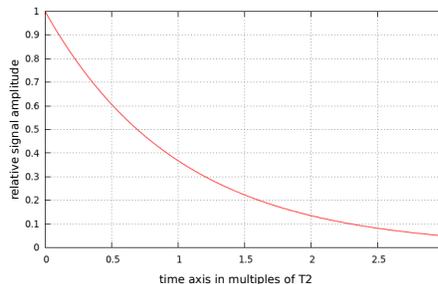


Figure 2.4: Signal decay after polarization is over

design considerations in the overall scheme of things. We will also look at this in more detail later when we consider actual circuits.

2.3.3 Amplification

The desired signal is normally very small - of the order of microvolts only. Therefore it is very important that the amplifier not add any more noise than necessary. Fortunately, the study of low-noise amplifiers in the audio range is well advanced because of the preponderance of audio devices being manufactured. Moreover, low noise operational amplifiers and/or transistors are readily available so it is just a matter of designing an appropriate circuit to take advantage of them.

The signal source, the sensor, is a low impedance device with typical source resistances of ~ 10 Ohms. It is also inductive with typical inductances of a few tens of mH. As stated before, the typical output voltage levels are a few microVolts. Commonly, the sensor circuit is designed with some external capacitance in order to resonate the sensor at the precession frequency. I will discuss some additional consequences of resonating the signal in this way in a later section.

All inductors have some resistance and resistance is a lossy circuit element. The “quality” of an inductor is described by a number, Q , which is a measure of the lossiness; a high Q means a low loss. Since the overall Q of the L-C circuit is determined almost entirely by the sensor (capacitors have very low losses) and since these sensors typically have Q s in the range from 10 to 100, the first limitation on the bandwidth of the system is determined by this front-end resonant circuit. The higher the Q , the narrower the bandwidth. There is a very simple relationship between Q and bandwidth; namely:

$$BW = \frac{f}{Q}$$

where f is the resonant frequency. BW is defined at the frequency difference between the -3 dB points in the voltage response curve for the resonant circuit.

In principle, the circuit may be arranged in two ways depending on whether one has a low impedance amplifier or a high impedance amplifier. These two alternatives are shown in Figure 2.5.

In both circuits, the capacitor is resonant with the sensor at the signal frequency (2130 Hz for a 50 μ T field). Consider the parallel resonant circuit. If the resistance of the coil is r Ohms at the signal frequency (approximately 20% greater than the dc resistance of the coil), then the parallel effective source resistance of the parallel configuration is Q^2r . For example, if the dc resistance of the coil,

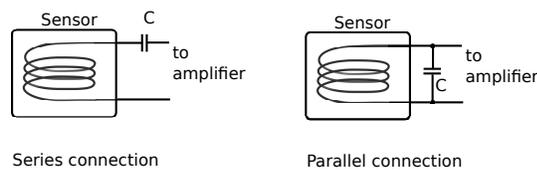


Figure 2.5: Series or parallel sensor configuration

R_w , were ~ 8 Ohms, the effective resistance, r , at 2.13 KHz would be ~ 10 Ohms. If the coil Q were ~ 25 , then the source resistance in the parallel case would be $\sim 625 \times 10$ which is $\sim 6.3K$. For the series resonant circuit, the source resistance would just be R_w plus the additional 20% due to the skin effect.

The topic of low-noise amplification is a very broad one and, without doing too much violence to the facts, it is possible to say that the parallel configuration almost always gives better results. This is because the lowest noise audio amplifiers tend to have high input impedances. National Semiconductor Corporation has a series of excellent application notes describing low-noise techniques. It is possible, however, to use an audio transformer in the series configuration to transform the low impedance sensor source into a higher impedance one.

The usual way of characterizing an amplifier's noise characteristics is to think of it as something which is generated inside the amplifier at the input. The total signal-to-noise ratio is then the ratio of the desired signal amplitude divided by the sum of the amplifier noise plus the noise from the sensor itself. When adding noise amplitudes, it must be remembered that it is not the simple addition of these two component's amplitudes because the phase is random. However, the powers do add so the final amplitude is the square-root of the sum of the squares of the two noise amplitudes. The noise of an amplifier is usually stated, in manufacturer's data sheets, as a noise voltage density; an amplitude per square-root of bandwidth; for example, the Linear Technology LT1037 operational amplifier has a typical voltage noise density of $2.8 \frac{nV}{\sqrt{Hz}}$. Obviously, the lower the noise from the amplifier, the better.

For PPMs, we need an amplifier with a low noise in the audio range - fortunately, there has been considerable development in low-noise audio amplifiers and there is a range of amplifiers available. The actual noise generated by an amplifier usually depends on the source impedance - in this case, the PPM sensor is the source. In the series configuration, the source impedance will be of the order of tens of Ohms - in the parallel configuration, it will be in the range of several thousand Ohms. In the example given earlier, the parallel source impedance worked out to be about 6K. This is just about optimum for the National Semiconductor LM837 quad op-amp. For a parallel input circuit as input to this op-amp and with a reasonable sensor, the calculated amplifier noise adds about 40% to the sensor noise meaning that the S/N from the sensor is degraded by just 40%. National Semiconductor application note AN-104 gives a good description of how to do these sort of noise calculations. I recommend it for anyone interested in finding out more about amplifier noise. Linear Technology also make several low-noise op-amps specifically designed for the audio range of frequencies. The LT1037 is one such and it can be used in a similar circuit. A reader (thanks Charles) has suggested that there are even better op-amps with lower noise figures; the Texas Instrument OPA211, the Linear Technology LT6200 and Analog Devices AD797.

Although I have mentioned some low-noise op-amps as a suitable input amplifiers, some discrete transistors can be used to make low-noise amplifiers. One such is the National Semiconductor LM394. National Semiconductor has an

application note for this device, AN-222, which is very helpful. Unfortunately, this device is obsolete but quantities of it are still stocked by retailers who specialize in selling obsolete semiconductors. Analog Devices makes a replacement for this dual transistor; the MAT12.

In typical circuits, the voltage gain in the low-noise amplifier is typically about 1000 thus bringing the sensor signal up into the mV range. It is advisable not to have too much gain in the first stage of amplification because, normally, one wants to do some band-pass filtering before further amplification. After the signal is in this voltage range, the noise contributed by further stages of amplifiers will be negligible.

Finally, before leaving the topic of low-noise amplifiers, it is worth mentioning that normal carbon resistors are noisier than metal film resistors. Therefore, the resistors in the earliest stages of amplification should all be metal film.

The band-pass filter is an important part of the amplifier since it will define the system bandwidth and hence the output S/N. Part of the band-pass is defined by the sensor-capacitor parallel (or series) combination at the front - the bandwidth is inversely proportional to the Q of the L-C network. For typical Q's of ~ 25 , the bandwidth will be ~ 90 Hz at 2.1 KHz. However, the amplifier noise is added subsequent to this L-C circuit so it is important to have an additional band-pass circuit later. The question then becomes: what bandwidth is desirable?

The answer is: as narrow as feasible considering the task to be done. If the magnetometer is to be used for surveying the magnetic field over a small region of the earth's surface, then the magnetic field is not going to vary very much and the bandwidth of the filter can be made very small. If the magnetometer is to be used in a spacecraft flying over the earth's surface, the variation will be large and so either the filter needs to be very wide or it needs to be adaptive to the value of the magnetic field during the flight.

For the hobbyist building a proton-precession magnetometer, the low-noise amplifier and band-pass filter is probably the most difficult part of the whole instrument to design and fabricate. Low-noise circuitry requires careful circuit layout to avoid oscillation and to avoid ground loops which may increase the noise level. The proper design of these circuits is not trivial!

3 Practicalities of making a Magnetometer or Gradiometer

3.1 Overall Considerations

A single magnetometer will give a measurement of the total magnetic induction at the location of the sensor and at the time that the measurement is made. On the surface of the earth, the local magnetic field not only varies with location but it also depends on the time. These time variations are due to the presence of currents in the upper ionosphere which are caused by the interaction between the earth's magnetic field and the solar wind. These variations are greatest in the auroral zones of the earth (roughly at latitudes of ~ 60 degrees) but do occur at all latitudes.

In the auroral zones, it is common to see magnetic sub-storms in which the local magnetic field will vary by some hundreds of nT over a period of a few hours. Even at lower latitudes, variations of tens of nT over periods of several minutes are common. A single magnetometer head will allow one to monitor these variations; to have one's own magnetic observatory. In prospecting, however, one wants to detect the small variations in magnetic field caused by ferro-magnetic materials in or under the soil. These small variations can easily be masked by the time dependent variations described above. That is, if one makes a series of observations at a number of different locations over the surface of the earth, each one will be different. However, one does not know whether these differences are due to artifacts under the soil or, perhaps, they were caused by the whole magnetic field in that vicinity changing in time over the period that the measurements were made.

For this reason, gradiometers are often used; especially in regions where one might expect significant short-term magnetic variations. In a gradiometer, there are two sensor heads. Local time variations due to upper atmospheric currents will cause the same change in both sensors; both will increase or decrease essentially the same amount because the cause of the change is very far away. The difference between the sensors will be due to magnetic material in the vicinity of the two sensors alone. The two sensors might be mounted at the top and bottom of a non-magnetic rod which is held vertically at a location where a measurement is being made but, because sensors are large and heavy, this would be a very cumbersome instrument. Another technique is to locate one of the two sensors in a fixed location and just have one mobile sensor used to make a survey over some area. In principle, the data from both sensors needs to be recorded simultaneously and only the difference used. In my experience, this is best done by having two independent magnetometers set to operate at fixed time intervals and each recording the data. These data are then subtracted subsequently and presented in some human readable form. However, one might consider linking the two instruments by radio and just recording the differences. Yet another configuration is to mount the two sensors horizontally at the ends of a long non-magnetic boom. This can then be mounted on a vehicle or towed

behind a boat and measurements taken as the system is moved along.

With this background, let us now look at what qualities are important in an instrument. From the viewpoint of sensitivity alone, big is good. The larger the sensor, the more working fluid there will be and the more turns you can wind. Both of these increase signal-to-noise ratio. For a field instrument meant to be carried on a back-pack, there are obvious constraints on the size and weight and, in general, you want the sensors to be as small and light as is commensurate with a reasonable signal-to-noise ratio. The secondary consideration must be the power required to polarize the sensor. Again, for a back-packed instrument, you are constrained by the size and weight of the batteries needed to operate the instrument. Dry cells give the greatest power to weight ratio at a reasonable cost however, dry cells are usually cased in steel containers and so are highly magnetic. Ni-Cd batteries are magnetic as are some Lithium batteries because of the way they are encased. In practice, you will likely want to use a gel-cell rechargeable battery because it is not magnetic and still gives a reasonable energy per unit weight. In recent years some non-magnetic Lithium-ion cells have become available and, despite some cautions about how to charge them safely, these certainly give the highest power per unit weight.

Let me work out one example. Suppose you want to make a portable instrument suitable for back-packing. You decide to make a sensor which have an all-up mass of several kilos and which consumes about 5 Amperes at 12V when polarizing. If the polarization takes place for one second at intervals of two seconds it will consume about 60 Joules of energy during this two-second period. We know that the remainder of the electronics will consume some power. However this will, usually, be small compared to the total power consumed in the polarization portion of the measurement cycle. A small 1.4 A-hr 12V gel-cell battery can produce about 60,000 Joules of energy ($12 \times 1.4 \times 3600$). Therefore, you might think that you can make about 1000 measurements on each charge of the battery. In reality, you won't be able to do this because the battery's rated capacity is given for a much lower discharge rate; nominally they are rated at 1/10 the rated discharge rate. That is, a 1.4 Amp-hour battery is rated at this capacity for a discharge current of 1.4/10 which is just 140 mA. Polarizing the sensors takes a very high current and at a high current, the capacity of the battery is considerably less than the number written on the battery case. So, let's estimate that the thing will be good for at least 250 measurements, perhaps as many as 500. You have to decide whether one to two hundred measurements is sufficient. If you're staying in a motel every night and just walking around nearby during the day, you can recharge the battery every few hours and so it might be acceptable. If you're back-packing into some rough country for several weeks, it clearly wouldn't be good enough. In that case, you can choose a larger capacity (and heavier!) battery - or, you can choose a bigger core, wind more turns on it and still get the same S/N with a smaller current drain. Or, you can decide that a S/N ratio of 100 is good enough and so get by with a smaller and lighter sensor. In short, you have to optimize the design depending on what you consider to be the most important constraints.

3.2 Winding the coil

3.2.1 Solenoidal sensors

Winding a solenoid is relatively easy. The working fluid can be put into a bottle which will slip inside the form and so may be readily changed and refilled if necessary. A convenient container is the more or less standard laboratory 0.25 litre bottle; it has a diameter of about 6 cm and a length, not counting the neck, of about 11 cm. A solenoid is most conveniently wound on a lathe and, with a little care, it is possible to wind the number of layers very uniformly and smoothly. The form can be made from PVC plastic irrigation tubing which is available in a number of diameters and for which there are very good adhesives. A quite usable sensor can be made by using two such solenoidal sensors connected so that induced background signal sources cancel out. A dual solenoidal sensor of this type is shown in Figure 3.1.

The two solenoids are made as similar to one another as possible and wound in the same direction. It is possible to wind them in opposite directions but it is not necessary since the background noise cancellation can be done by paying attention to how the two solenoids are connected together. The solenoids can also be put together into a single unit end-to-end rather than beside one another. There should be a space of at least half the diameter between the two solenoids if this is done; a full diameter is better. Assuming the two solenoids were wound in the same direction, then the connections for both the side-by-side and end-to-end are shown in Figure 3.2.

The dual solenoid suffers from the primary disadvantage of all solenoids which is its directional sensitivity. Nevertheless, except near the magnetic equator, the local magnetic field will have a large vertical component so, holding the sensor so that the solenoids are horizontal, the signal will never go to zero and the sensor is quite usable. The noise-cancelling effect of the two solenoids, connected as shown, is quite good and so the effect of the large 50 or 60 Hz background signal is greatly reduced.

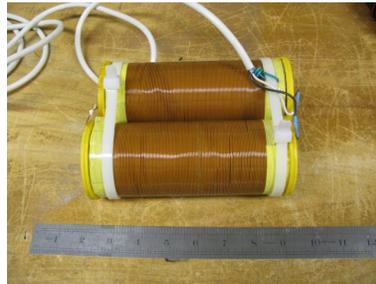


Figure 3.1: Photo of a dual solenoidal sensor

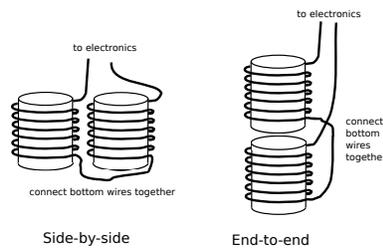


Figure 3.2: Dual solenoid connections

3.2.2 Toroidal sensors

Assuming one can find a suitable form, then it is much much better to use a toroidal core because of its overwhelming advantages. However, making a toroidal sensor has a number of difficulties. The first problem is finding that suitable form! In previous revisions of this document, I described winding a sensor using the hollow blue ring which is the largest ring on the Fisher-Price Rock-a-Stack toy. This blue ring from the Fisher toy produced a sensor which was very marginal as it only has a volume of about 0.18 litres of fluid; this was a really too small. At the time I wrote the first revision of this document, a reader sent me an example of a fibre-glass fishing net float which was toroidal in shape but considerably bigger than the blue Fisher ring. I regret that I have misplaced the name of this person who sent it to me and I do not recall where he got it; if he reads this, I would appreciate it if he could contact me again as I think it might be the basis for a very good sensor.

Even having a suitable form, the next difficulty in making a toroidal sensor is winding the requisite number of turns on the form. In previous revisions of this document, there was a picture showing the blue toroidal form, filled with distilled water (with no dissolved oxygen), a reel of #20 AWG wire and a shuttle. The shuttle was made from some thin plywood and was about 0.7 meters in length. The shape was slim enough to pass easily through the centre of the toroid even when the toroid was nearly completely wound. The shuttle was wound with about 200 meters of the wire. It was convenient to put a rubber band around the wire at each end of the shuttle - this kept the wire from falling off the shuttle too easily as it is passed through the centre of the toroid form. It was easy to wind the first layer of the toroid uniformly, closely wound on the inside and evenly spaced on the outside of the toroid, because the form was smooth and uniform. The second layer was more difficult because it did not fit nicely over the first layer on the inside. Each successive layer got more and more difficult to wind uniformly. However, if you take your time, it is possible to do a fairly neat job.

It is very important to avoid kinks in the wire - when you get a small loop (and you will get many of them!), carefully straighten it out - do not pull it tight. The finished sensor contained about 2000 turns of the #20 AWG enameled wire and had a total dc resistance of 7.4 Ohms. I estimate that it took about 8 hours of winding to create this sensor - done a bit at a time while watching TV in the evenings or while having morning coffee.

A fatter toroid core would be better as it would hold more of the working fluid. I have considered using an inner tube from a small scooter tire. These can be obtained with outside diameters of the order of 15 cm and with a radius of the core of about 2 cm. The problem here is how to make them rigid enough to wind on the wire. One possibility that occurred to me was to fill the core with water and then freeze it. If the wire size used is greater than about 1 mm or so, the wire itself becomes rigid enough after a few layers have been wound. Therefore, after winding the core, it can be allowed to thaw. I have not tried this - it is just an idea. An alternative might be to inflate the inner tube to the

desired shape and then wrap it with fiberglass tape and give it a coat of resin to harden it. Again, I have not tried this.

One final caution: don't keep the wire too taut as it is wound - on either a solenoid or a toroid. When the wire is tight, it will tend to compress the core. If you are using a soft plastic form, because of the number of turns, it is possible to collapse the coil form if the wire is wound too tightly. It is necessary to make each turn snug in order to wind the coil uniformly but you should not wind it any more tightly than is necessary.

3.3 Other Features in a Field Instrument

An instrument to be used in the field for prospecting or exploration should have several features. I have already discussed the weight-sensitivity trade-off. Another very desirable feature is to provide some means of automatically recording the values measured. For example, when prospecting, it is often desirable to make a series of measurements in a grid over the ground - perhaps on a one meter grid. It is annoying and time consuming to have to write down the value of the gradient in a notebook at the time of each measurement. It is much better to have some internal electronic storage in the instrument to record the successive value as they are measured. Then, it is a simple matter to note where the grid started and its orientation and to recover the actual values later. Fortunately, large scale digital storage is easily and cheaply available and adding this feature to the instrument is relatively trivial. Readily available and inexpensive serial EEPROM memory devices are capable of storing several thousands of multi-byte data values and keep this data even when the power is turned off. The latter aspect is important since you don't want to lose data if the batteries fail.

More recently, FLASH memory devices such as memory cards designed originally for cameras and/or USB "thumb drives" have become widely available. They are extremely inexpensive, are easy to configure in order to write data to them in the magnetometer and easy to subsequently read the data in a personal computer.

3.4 Overall Block Diagram

The overall block diagram for a modern magnetometer is shown in Figure 3.1. There are, of course, many ways of implementing the required functions and the above block diagram is just one such way. However, there will always be an analogue section which amplifies the low-level precession signal and a digital section where the analogue signal is ultimately converted into a number representing the magnitude of the magnetic field.

3.5 The Controller

The controller provides the interface between the sensor and the human operator. The human interface is some sort of method for providing input; switches or a touch sensitive display. It will almost certainly have some sort of menu

system to select options. It goes, almost without saying, that the controller should be constructed around a microprocessor rather than being built out of fixed logic devices.

3.5.1 The Display

The display gives the operator information about the operation of the unit. For example, it must show the measured values after each measurement has been made. It must also give some information about the internal operation of the instrument as a whole - perhaps some measure of the battery capacity remaining, how much internal memory capacity remains unused, etc. Fortunately, very rugged and very low power alpha-numeric LCD displays are available and they are well suited to the task. Graphical data is not really necessary in a field instrument but LCD versions of these are also available. LED display devices are not really suitable since they tend to be difficult to see in sunlight.

3.5.2 Switches or touch-sensitive display

The switches or touch-sensitive display provides the human input to the instrument. This portion needs to be carefully thought out. Any field instrument has to be extremely simple to operate. However, people get tired after a day in the field so it should not be too easy, for example, to erase the instruments internal memory. Since the interface is to a microprocessor, it is possible to use a menu method to access the various functions of the instrument. This requires only a few switches or touches but gives the greatest flexibility. The standard modes of operation should require the fewest human interactions while more dangerous or complex tasks (e.g., memory erasure) should require a not-so-simple sequence of human actions.

3.6 The Polarizing Circuit

The polarizing circuit deserves considerably more discussion. On the surface, it appears simple enough - you just have to connect the sensor to the battery for the polarizing period and then turn it off. The first part really is simple enough. When the voltage is applied across the sensor coil, the current rises from zero to the final value exponentially with a time constant of $\frac{L}{r}$. Typically L is several tens of mH and r is a few Ohms so this time constant is of the order of a few to perhaps a few tens of mS.

The latter part, the turning off the current, needs to be described in more detail. We have said before that the polarizing field needs to be turned off in a time much less than the precession period. This means that the current must be switched off in less than about 50 microseconds. Unfortunately, the sensor has considerable inductance and it takes some time to dissipate the energy stored in the sensors magnetic field by the polarizing current. Lets start this discussion but considering what we mean exactly by turning off the polarizing field.

Consider, for a moment, a situation where the polarizing field and the local earth's magnetic field (which is what you want to measure) are at right angles to one another as shown in Figure 3.3. In this figure, B_e is the earth's field, B_p is the polarizing field

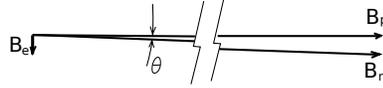


Figure 3.3: Polarizing field and the earth's field

and B_r is the resultant of the two. Typically, when the polarizing current is flowing, B_p is much, much greater than B_e and so the angle, θ , is very small - a fraction of a degree; it is exaggerated in the drawing to make it visible. So B_r , the resultant polarizing magnetic induction is essentially parallel to B_p . As the current is turned off, as it collapses, the resultant stays more or less parallel to B_p until the polarizing field is only a few times larger than the magnitude to B_e and then it starts to move towards B_e . It is really just this last period that must be smaller than the precession time since, prior to that, there is no significant change in the direction of the polarizing field. Let us arbitrarily say that when θ becomes ~ 10 degrees is the beginning of the effective turn-off period. Let this period be T_s ; the switching period. It starts when θ is ~ 10 degrees which is when:

$$\frac{B_p}{B_e} \simeq \cot(10^\circ) \simeq 5.7$$

Therefore, we may consider the turn-off switching time to start when B_p is $\sim 6B_e$ which is $\sim 300 \mu\text{T}$. Let us call the current flowing through the inductor at the moment to be I_s . Let T_s be the time from when the polarizing current goes from I_s to zero. It is this time which we wish to be much shorter than the precession period. Now, we cannot carry this discussion forward without considering the exact circuit used for polarizing. The simplest possible circuit, shown in the Figure 3.4A, is the worst.

There is no problem when the switch is set to connect the sensor to the polarizing voltage; the problem is when it is switched the other way to turn off the current. At the instant the polarizing current is interrupted, the negative induced voltage at the top end of the sensor will rise to a very high negative value and an arc will appear briefly across the switch contacts. Every time this happens, the switch contacts will get dirtier and dirtier due to residue for the arc and eventually, the switch will fail. The only way to prevent this, if a mechanical switch is used, is to somehow "snub" the voltage produced by the sensor so that an arc does not form. A common way of doing this is to connect a Zener diode across the sensor as shown in Figure 3.4B; the conventional diode

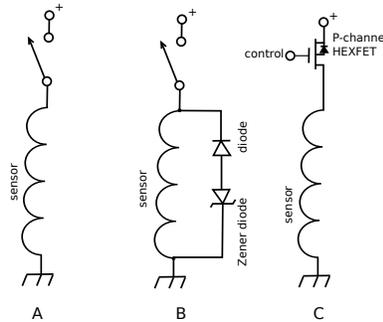


Figure 3.4: Possible polarization circuits

is used to prevent current flow through the Zener diode when the polarizing voltage is connected. Another method is to use a HEXFET as shown in Figure 3.4C and to use it to turn the polarizing voltage on and off. When the HEXFET is turned off, no current flows through it until the HEXFET is turned on. HEXFETs are commonly used because they have such a low on-resistance that the losses through them are negligible. When the HEXFET is turned off, induced sensor voltage when the current is broken will cause the HEXFET to temporarily break-down at its rated breakdown voltage; typically 100 Volts or so. This temporary breakdown of the HEXFET does not damage it as long as the total energy dissipated during the breakdown does not cause it to get overly hot. With the proper choice of HEXFET, this is not a problem.

Consider the case for a HEXFET circuit, using, for example, an IRF5210. When the HEXFET is turned off, the magnetic field of the sensor starts to collapse. This induces a voltage across the sensor which is negative at the HEXFET end with respect to ground. The drain of the HEXFET starts to break down at its rated break-down voltage which is, for the IRF5210, about -100 Volts. The current rapidly goes down towards zero during this breakdown period and, when it reaches zero, the induced voltage goes to zero and the switching period is over.

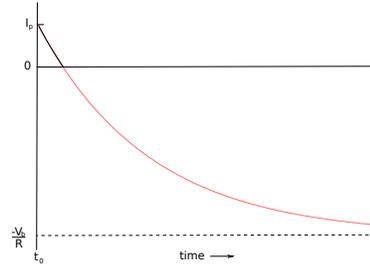


Figure 3.5: Graph of I vs time when HEXFET is turned off

Figure 3.5 shows the current through the sensor if the HEXFET is turned off at a time, t_0 . In this figure, V_b is the negative voltage at which the HEXFET breaks down. The exact same equivalent circuit is applicable for the case where the voltage is snubbed by a Zener diode which breaks down at that voltage. The equation for I in the equivalent circuit is:

$$V_b + IR + L \frac{dI}{dt} = 0$$

where R is the dc resistance of the sensor, and L is its inductance. Given the initial condition that the current is I_p when $t = 0$, just prior to the instant that the switch is opened (or the HexFET is turned off), the solution for I is then:

$$I(t) = \left\{ I_p + \frac{V_b}{R} \right\} e^{-\frac{R}{L}t} - \frac{V_b}{R}$$

We are only interested in the interval between when the HEXFET is first turned off and when it reaches zero current; this is shown in more detail in Figure 3.6. Here, T_s is the total time for the current to go from I_s to 0. The time at which I is zero is t_s and the above equation may be solved to find this time:

$$t_s = \frac{L}{R} \ln \left\{ 1 + \frac{I_p R}{V_b} \right\}$$

For example, considering the dual solenoid shown in Figure 3.1. I_p was 1.7 A, R was 7.0 Ohms, V_b was 100 V and L was 31 mH. Substituting those numbers into this equation gives a calculated value for t_s of 0.5 mS. The value observed on an oscilloscope was in good agreement with this calculated value. However, we do not really care what t_s is; rather we need to know the time from when the current has fallen from I_s to when it is zero. This is a much shorter time, T_s . In order to estimate this time, we can calculate the rate at which the current is decreasing in the vicinity of t_s and use this to estimate the time taken to go from I_s to zero. The rate of change of current, at any time t, is given by:

$$\frac{dI}{dt} = -\frac{R}{L} \left\{ I_p + \frac{V_b}{R} \right\} e^{-\frac{R}{L}t}$$

When $t = t_s$, this equation will give us the rate of change of current at the time when the current is zero. Substituting the values for the same toroidal core, $\frac{dI}{dt}$ is about -3200 A/sec and, since I_s is about 50 mA, T_s turns out to be about 15 microS⁶. This is sufficiently small to satisfy the criterion that the turn-off period must be much smaller than the precession period which is about 500 microS.

These equations are applicable for circuits using a Zener diode; in these cases, the breakdown voltage is just the Zener voltage. It is interesting to apply these equations for those cases where V_b is very small. For example, some published circuits for PPMs just show a reversed diode across the sensor (sometimes, a few in series). That is, in the second circuit, the Zener diode is replaced with a short circuit.

In this case, V_b might be as low as just one diode voltage drop (~ 1 V at the typical currents of a ~ 1 to ~ 10 A) and, t_s may be many mS. Therefore T_s is also very long - perhaps too long!

Finally, although strictly speaking not part of the polarizing circuit, we must discuss how to connect the sensor to the amplifier after the polarization cycle is complete and we want to amplify the proton precession signal. The voltage at the output end of the sensor goes from the positive polarizing voltage (12V in this example) during the polarizing portion of the cycle to the negative breakdown voltage (-100 V in this example) of the HexFET at the end of the cycle. If we were to just capacitively couple the sensor output to the front end of our low noise operational amplifier, we would be putting a voltage to the

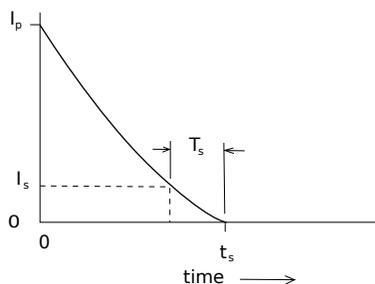


Figure 3.6: Detail of current decay near zero current

⁶The value of I_s may be calculated by knowing that the internal field should be, in this case, about 6 times the earth's field; i.e., 300 nT. The formula for calculating I_s depends on the configuration. For a toroid,

$$I_s = \frac{2\pi R \times 6B_e}{n\mu_o}$$

input with a swing of about 110 V. This would certainly not do the amplifier any good!

We could lower the voltage swing by using a lower voltage Zener as mentioned above - but then, the time taken to collapse the field would be longer. To get the voltage swing low enough not to damage the amplifier might result in an unacceptably long time to collapse the field. Obviously, simple capacitive coupling is not going to work unless the amplifier input can handle a fairly wide voltage swing. Typical low-noise amplifiers cannot.

The traditional answer to this problem has been to use a relay to connect the sensor to the amplifier after all the high voltage stuff has settled down. This has the disadvantage that all relays have a soft iron core and so will perturb the local magnetic field which we are trying to measure. However, it is possible to find some very small relays with very small cores and, if the relay is located at some distance from the sensor, the problem is not too serious. In the next section, I will discuss a more complex semiconductor switching circuit which eliminates the need for a relay.

4 Some Other Considerations

4.1 To resonate or not to resonate

In the discussion about amplification, we assumed that the sensor be coupled with a resonant circuit; either a series circuit for a low-impedance amplifier or a parallel circuit for a high-impedance amplifier. As we shall see below, this confers some advantages to the overall system but it has the disadvantage that the value of the resonating capacitor will depend on the value of the precession frequency and that, in turn, will depend on the local value of the magnetic field. This means that the correct value for the resonating capacitor will depend on where, on the earth, the measurements are to take place. Over the earth, the value of the magnetic field varies between about 3 and 6 microT. This is a considerable range and means that the value of the resonating capacitor will also need to be changed as one goes from one region to another. This is a problem only, of course, if the same magnetometer were to be used, say, in South Africa where the field is near the lower limit and in North America where it is near the upper limit. For most users, this is not a significant problem but it does mean that the magnetometer must be tuned for the region where it is to be used. The advantage of the resonating capacitor is that it provides some free, noiseless amplification!

Consider the parallel tuned circuit shown in Figure 4.1. The circuit shows the sensor as a voltage source (the precession signal), V_p , in series with the sensor resistance (including the skin effect) and inductance; r_s and L_s respectively. There is, in general, a coupling capacitor, C_c , between the sensor and the amplifier. The amplifier is shown as the resonating capacitor, C_r , in parallel with the amplifier input resistance, R_{in} . The

voltage appearing at the input of the amplifier is V_a . We shall analyze this circuit making some simplifying assumptions. Firstly, the coupling capacitor, C_c , is assumed to be much, much larger than the resonating capacitor. This means that its reactance will be much, much lower than that of the resonating capacitor and so, for the ac equivalent circuit, it may be regarded as a short circuit. Secondly, we shall assume that the amplifier input resistance is much larger than the resonating capacitor reactance at the precession frequency. These approximations are usually good ones as the coupling capacitor is usually at least ten times larger than the resonating capacitor. The input resistance is typically at least 10^5 Ohms (for the LM394, for example) while the resonating capacitor reactance at the precession frequency is typically between 10^3 and 10^4 Ohms. With these approximations, it is possible to show that the voltage appearing at the amplifiers input, V_a , is related to the precession voltage, V_p , by the following

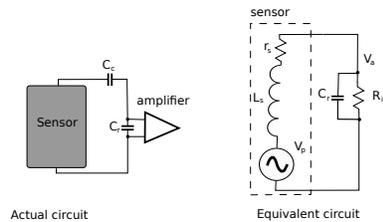


Figure 4.1: Sensor and first amplifier circuit

equation:

$$V_a = \frac{\omega L_s}{r_s} V_p = Q V_p$$

ω is the angular frequency ($2\pi f$) and Q is the ratio of inductive reactance of the sensor to the losses in the sensor. Q has typical values between 10 and 100. This means that the precession signal has been amplified by the factor Q . Unfortunately, this does not increase the S/N ratio since the noise at the precession frequency is also amplified. The series circuit for a circuit using a low-impedance amplifier may be analyzed the same way and gives the same result: the S/N is unchanged but the signal is amplified by the value of Q of the sensor. So, given the fact that you get free amplification, should you always choose to add a resonating capacitor?

The answer depends on how marginal the signal is to start with. If you have a rather small sensor, the extra free gain is helpful and may compensate for some of the low-noise preamplifier noise. For a large sensor such as those towed behind boats in ocean based magnetometry, the signal may already be quite large (several tens or hundreds of microV) and so the extra gain is not necessary and the fact that the magnetometer has to be tuned for different locations may be a nuisance.

4.1.1 Measuring Inductance

If you do decide to use a resonating capacitor, to determine its value you need to know the inductance of the sensor. Since sensor Q 's are typically something in the neighbourhood of 20 to 50, it means the value of inductance must known to an accuracy of about $\frac{1}{Q}$ in order to calculate the needed capacitance; that is, 2% to 5%. The spreadsheets mentioned in a previous section do have a column giving the estimated inductance for any given sensor configuration but these estimations are not likely to be accurate enough. Sensors also have some considerable stray capacitance (perhaps several hundred pF) so most of the types of inductance meter to be found in the average amateur's workshop will not give a valid value for the inductance of the sensor.

The only way I know to get a good value for the inductance requires a good, accurate signal generator in the audio range plus an oscilloscope.

The signal generator is connected to a coupling loop which is placed in the vicinity of the sensor. The sensor is connected to the oscilloscope and a pair of test leads is connected to the sensor to allow connection of a test capacitor in parallel with the sensor. This is shown in Figure 4.2.

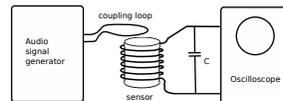


Figure 4.2: Measuring the inductance

Note that the coupling loop, made by using the test leads on a cable going to the signal generator, has a resistor in series with the loop; it should have a value equal to the output impedance of the signal generator. This coupling

loop is draped over one solenoid of the dual-solenoid sensor previously shown in Figure 3.1. In the background of the photo, you can see the display of my signal generator showing it set to 2131 Hz. You should set it to a value in the vicinity of the expected precession frequency for your area. Then, using just a pair of clip leads, substitute values of the capacitor, C, until you see the maximum signal on the oscilloscope. As a starting point, assume the inductance is going to be the value calculated from the spreadsheet. Then the value of C will be

$$C = \frac{1}{4\pi^2 f_p^2 L}$$

where f_p is the expected precession frequency. If you have a double solenoid sensor, remember that the total inductance will be close to twice the inductance of each solenoid. Once you have a value for C which gives you close to the maximum signal level, then vary the frequency of the signal generator to and fro slightly to more precisely find the frequency which gives you the maximum value. Then, disconnect C from the circuit and measure it carefully with a capacitance meter. If the frequency of the maximum is f_m , and the value of the capacitor which gives you a maximum at this frequency is C_m , then the inductance of the sensor is:

$$L = \frac{1}{4\pi^2 f_m^2 C_m}$$

4.2 Relay-less polarization

The presence of a relay in the instrument is always a problem. Relays have iron cores and so perturb the values of the local magnetic field. This is not a problem with a shipboard magnetometer towing the sensor at some distance behind the ship but it is a problem with hand-held magnetometers. However, in addition, a relay is an electro-mechanical device and so is prone to mechanical failure, dirty contacts, etc. If, the relay is switched hot (i.e., the polarizing voltage is always present at the relay contacts), there will be arcing when the relay is switched off and this will certainly dirty the contacts very rapidly.

Since the desired precession signal is only of the order of microV, any dirt on the contacts will add unnecessary noise to the signal. Finally, relays are normally rated to operate for several hundred thousand closures. However, the relay in a shipboard magnetometer may be switching once per second while it is operating and since there are 3600 seconds in an hour so it may not take long to get to the manufacturers rated lifetime.

There are semiconductor devices which are called HEXFETs which can act as switches and so may be used to

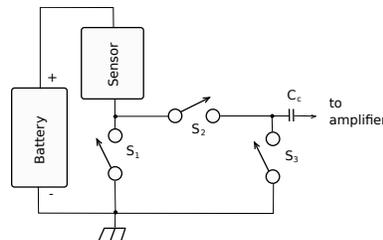


Figure 4.3: Simplest HEXFET switching circuit

replace the relay. They are MOSFETs in which the gate is insulated from the conducting channel. They are designed to be non-linear and have very low gate-source turn-on voltages and very low on-resistance. Typically, they can be turned on and off by 5V logic level signals and have ON resistances of a fraction of an Ohm and OFF resistances of several megOhms. One can imagine a simple circuit for switching the polarizing current to a sensor and then switching the precession voltage to an amplifier.

Consider the circuit diagram shown in Figure 4.3. I have shown the HEXFETs as if they were simple switches. To polarize the sensor, switch S1 is closed and, to isolate the sensor from the amplifier during the polarization, S2 is opened and S3 is closed. To connect the sensor to the amplifier after polarization, S1 is opened, S2 is closed and S3 is opened.

All three of the switches can be replaced by HEXFETs (specifically N-channel) and the opening and closing of the switches is done by applying logic-level voltages to their gates. Since the voltage at the junction of S1, S2 and the sensor will go up to the breakdown voltage of S1 (which might be $\sim 100V$) when it is opened, it is important that S2 not break down at this voltage - it must have a higher rating than S1. This circuit has one problem and that is that S1, when opened has a leakage current which flows through the sensor. This will produce a small magnetic field which may be significant. Even worse, the leakage current will be temperature sensitive so it will not even be constant during the operation of the magnetometer.

Ideally, we would like the DC voltage across the sensor to be zero during the measurement phase of the operation. This can be done by adding some other HEXFET switches as shown in Figure 4.4. In this circuit, two additional HEXFETs have been added; S4 and S5. During polarization, S4 is closed and S5 is open. After polarization, S4 is open and S5 is closed. This connects the top end of the sensor the ground and so no leakage current flows through the sensor. S4 would be a P-channel HEXFET while S5 would be another N-channel one.

As a final touch, immediately after polarization stops, because the sensor

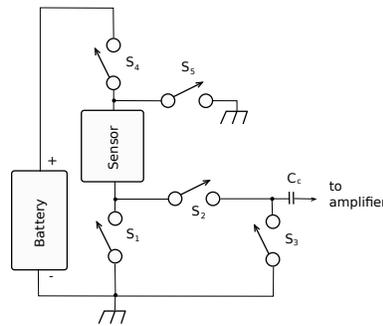


Figure 4.4: More complex HEXFET switching circuit

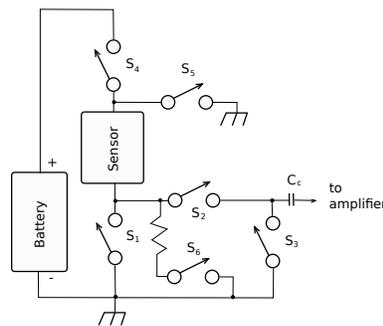


Table 1: Final HEXFET switching circuit

would otherwise ring at the self-resonant frequency of the inductor, it is useful to connect yet another HEXFET switch to the junction of S1, S2 and the sensor to ground through a resistance to act as a squelch of this transient. The final resulting circuit is shown in Figure 1. Here, S6 is a N-channel HEXFET and the resistance is of the order of 10K. The controller has a more complicated job in switching these HEXFETs but it is not difficult. A typical sequence might be as follows:

1. turn on S3 and turn off S2 - this isolates the amplifier from the rest of the circuit, also turn off S5,
2. wait a few mS,
3. turn on S1, S4 - this turns on the polarization current, also turn on S6,
4. wait for the polarization time,
5. turn off S4 and S1, turn on S5, leave S6 on. wait about 10 mS to let the ringing die down,
7. turn off S6,
8. wait a mS or two and then turn on S2 and turn off S3,
9. wait another 10 or 20 mS for amplifier transients to die away and then start the measurement of the period.

There is one minor complication with this technique and that is that the microprocessor logic signals appear on the gates of the HEXFETs and will contain some microprocessor noise which may be capacitively coupled into the input of the low-noise preamplifier. If the microprocessor operates at 10 MHz, for example, there will be some 10 MHz noise on the logic output lines. However, this can be eliminated with a low-pass filter on these logic lines going to the gates of the HEXFETs. I have found that two or three sections of RC filter, each of a series 470 Ohm resistor and a 0.1 microF shunt capacitor to ground are sufficient to completely eliminate the microprocessor noise. This filter will cause a delay of approximately 0.5 mS between when the logic gate is put high or low and the corresponding HEXFET is turned on or off for each filter section; if you use three sections, the total delay will be about 1.5 mS. This is not a problem as long as you dont want to switch the gates more rapidly than this time delay; which is the case for the sequence of operations stated above.

4.3 Advanced Signal Processing

The discussion in previous sections assumed that the method of determining precession frequency was to measure the time taken for a number of cycles of the precession signal. In this method, the precession signal is essentially squared and a high frequency signal is counted in order to determine the total time for that number of cycles of the precession signal. This method is very susceptible to errors caused by noise spikes from any exterior sources. An error of one count will give a very large error in the calculated magnetic field.

Even if there are no such errors due to external noise spikes, this method also throws away a lot of information because all that is important is the time difference between some starting signal zero-crossing and a final one. One can

improve the statistics by, for example, also measuring the time for some intermediate signal crossings and, although they might have lower statistical weight than the largest time difference, they still provide some additional information. If an analysis system is designed so that the time for every zero-crossing is recorded, then all these data can be used to reduce the uncertainty significantly.

In a letter written to the editor of the Proceedings of the IEEE, Farrell and Grosch⁷ have shown that, if the time of each of N zero-crossings of the precession signal are recorded, then, using an algorithm they develop in the paper, the estimated period of the signal will have a variance which is reduced by a factor of $\frac{N(N+1)}{6(N-1)}$ compared to the single measurement of the time from start to finish of the N zero-crossings. Since the square root of the variance is a measure of the probable error of the calculated period, this means the probable error is reduced by this factor. For example, if N is 1024, then the probable error is reduced by the square-root of 171 which is about 13. This is a significant improvement and the method is easily applied if modern microprocessors are used to do the signal analysis.

There are other methods of determining the frequency or period of the precession signal which require that the signal be sampled at some high data rate and stored in digital form. One such method, called the phase-slip method, was described by Koehler⁸. Zhang *et al.*⁹ have described a different method for the calculation of the frequency of a sampled signal. Their algorithm, used in modern sampling oscilloscopes, again relies on having the sampling frequency being much greater than the precession frequency. In both these two methods the digitized data is used in such a way that the amplitude of the signal and its DC offset does not enter into the calculation. This is a desirable feature since the precession signal has an amplitude which decays during the period of the measurement.

Although all the previously described methods all do not require any knowledge of the amplitude of the signal, in that sense, they are not using all the information which is inherent in the sampled digital values. With such a digital sample sequence, it is possible to employ other, more general, mathematical techniques to determine precession frequency with the necessary precision. Again, the amplified precession signal is digitized at a high rate (typically at least 10 times the precession frequency) and the resulting time-sampled waveform then analyzed using techniques which are collectively called "Digital Signal Processing" or DSP. Modern personal computers are powerful enough that these mathematical techniques may be done in a time approximately equal to the polarization time. So, for example, using a sensor filled with kerosene, one can polarize for half a second, collect the data for another half-second and ana-

⁷E.J. Farrel and C.B. Grosch, "Determination of Period from the Times of Zeroes", Proc. IEEE, **53**, 1965, p2163

⁸James Koehler, "Proton Precession Magnetometer", Circuit Cellar, **202**, May 2007, pp14-23

⁹Jian Qiu Zhang, Zhao Xinmin, Hu Xiao and Sun Jinwei, "Sinewave Fit Algorithm Based on Total Least-Squares Method with Application to ADC Effective Bits Measurement", IEEE Transactions on Instrumentation and Measurement, **46**, August 1997, pp1026-1030

lyze during the next polarization. The result is one complete magnetic field measurement every second.

What are the mathematical techniques? For a starter, simple Fourier analysis is not going to be good enough. If you sample a waveform with N samples over a time period, T, you will get a spectrum with a resolution of 1/T Hz. It does not matter how many samples are taken! So if we were to sample for half a second, we could derive a spectrum with a resolution of 2 Hz which would give an uncertainty in magnetic field about 50 nT. To get a resolution of 1/25 Hz (in order to get a magnetic field accurate to about 1 nT), you would have to sample the waveform over 25 seconds and this is clearly too slow and the signal would have decayed significantly in 25 seconds.

The best mathematical technique is probably going to be some version of a best fit algorithm; a technique which uses the fact that you have some a priori knowledge of the signal. In the case of the precession signal, you know the voltage is going to conform to the equation:

$$v_s = A + Be^{-Ct} \sin(Dt + E) + noise$$

It is an equation with five unknowns and so it can, in principle be solved by a least-square best-fit algorithm. The general method for solving least-square best-fit to arbitrary equations with an arbitrary number of unknowns is called the Levenberg-Marquardt¹⁰ method. The variable, C, is going to be a roughly constant from one measurement to another; it just depends on the spin-spin relaxation time and so can be measured for any given sensor just once and then used in all subsequent calculations. Similarly, B, is going to be roughly constant for each successive measurement and so the value for it can be constrained to be within certain limits. One possible method of analysis is, then, to just take an average of the signal to get a value for A and to then subtract this from the signal. This can be done very quickly. Then, the resulting numbers can be multiplied by e^{Ct} to give a simplified equation of the form:

$$v'_s = B' \sin(Dt + E) + noise + A'$$

where A' is the DC offset error due to our crude method for determining A. Now we need only solve for B', D, E and A'; a best-fit analysis problem with just four variables. Best-fit algorithms require some initial estimate of the variable before starting. Here A' is going to be close to zero. B' may be fairly constant from one measurement to another unless the orientation of the sensor changes between measurements. In the case of the above equation, we usually have a pretty good idea of the approximate value of the precession frequency. Indeed, if you are prospecting either on land or over the ocean, you are likely to already have a GPS system to determine where you are and, knowing your location, you can easily calculate what the local unperturbed magnetic field magnitude is¹¹. There are a number of programs to calculate the earth's local magnetic field at

¹⁰William H. Press, Brian P. Flannery, Saul A. Teukolsky and William T. Vetterling, "Numerical Recipes in C", Cambridge University Press, 1988, p542ff

¹¹<http://www.ngdc.noaa.gov/geomagmodels/IGRFWMM.jsp>

any location on the earth's surface if you do not have access to the internet in order to go to the above site. The variable, E , is totally unknown and so is A' .

Most modern lap-top computers are fast enough to do this sort of analysis quickly and, if the best-fit is defined as a correlation to a sine wave of arbitrary frequency and phase, the errors can be made very small. I have not tried this technique myself but have done some computer simulations which are encouraging. This suggests that a complete PPM might be made by having just a fast A/D converter in the magnetometer and sending these digitized samples to a lap-top computer which does the subsequent analysis. USB interfaces are fast enough that transfer time can be made small enough. Alternatively, the precession signal can be put directly into the lap-top audio port and digitized by the lap-top's own A/D converter and subsequently analyzed.

Finally, there are specialized microprocessors which are specially designed for DSP analysis and they may be incorporated into a complete, self-contained, instrument to measure precession frequency with the required accuracy.

4.4 Other interference

The sensor may act as a radio receiver, especially if it is not toroidal. Almost everywhere on the surface of the earth, there are radio stations broadcasting in the HF bands and the induced voltages in a solenoidal sensor may be hundreds of microV in amplitude and even, in some locations, several mV. Because the low-noise preamplifier is biased at a very low current, it may be driven non-linear by these radio signals and hence demodulate them. Because the demodulated audio is in the same range as the desired signal, this can cause a problem. The solution is to put a low-pass filter at the input of the pre-amplifier to reduce the amplitude of these signals. Even as simple a filter as a small value capacitor (for example, 1000 pF) across the sensor terminals may be enough to do this. Remember that this capacitor has to withstand the voltage developed across the sensor when polarization is stopped.

For sensors being towed in the ocean, this radio interference is not as severe a problem because salt water attenuates signals in the HF region very effectively.

5 Projects for the amateur

5.1 A geomagnetic observatory

A proton-precession magnetometer can be the basis of a school science project or just a simple geophysical observatory. The sensor should be placed several tens of meters from other buildings or regions where people walk about. This requires more space than can be found adjacent to most people's residence. Nevertheless, many urban backyards, while not ideal, are large enough to allow interesting geophysical measurements. Everywhere on the earth, there are small fluctuations in the local magnetic field called micro-pulsations. These have amplitudes of a few to tens of nT and show periodic fluctuations with periods from a few seconds to hundreds of seconds. They are caused largely by interactions between the earth's magnetic field and the solar wind. In regions close to the auroral zones of the earth, the amateur can also observe magnetic storms and sub-storms with good precision. These are often associated with auroral phenomena.

For an observatory, the sensor is ideally put inside a shielded enclosure and the sensor connections brought inside the house to a room where the electronics is located. Two coaxial cables made from cable sizes like RG-8A/U will provide both shielding and will have a centre conductor of such a large size that the resistance will not be significant compared to the resistance of the sensor.

The front-end part of the electronics is probably the most difficult part of the system for an amateur to build because it requires a very low-noise preamplifier and a very large voltage-gain. The data analysis to measure the magnetic field magnitude from the precession frequency can either be built as part of the instrument using a microprocessor or, alternatively, the audio precession signal can be connected to the microphone input of a computer and can be digitized and further analyzed in software in the computer; the experimenter can try any number of techniques such as those mentioned in section 4.3 for accurately measuring the signal frequency and thereby determining



Figure 5.1: Shielded enclosure my backyard

the magnetic field magnitude.

Figure 5.1 shows the shielded enclosure I have in my back yard. The area is wooded and the enclosure is located only about three meters from a concrete slab which is my back patio. The concrete contains reinforcing rods which undoubtedly distort the nearby magnetic field. However, this will mostly be in the form of a permanent offset so that although the measured magnetic field will not be equal to the undisturbed value for this neighbourhood, nevertheless the temporal variations due to micro-pulsations, etc. will still be observable. The sensor in this enclosure is the same one shown in the Circuit Cellar article (Footnote 8). The sensor is located in about the middle of the enclosure; it sits on top of a plastic pot of the type used for bedding plants. One must be very careful when making the enclosure to make sure that there is no iron or steel used. The coaxial cable used is two lengths of RG-8A/U; it is available from electronics supply houses.

If the audio signal is connected to the microphone input of a computer, the precession signal can be digitized by a number of programs. One such program which is both useful and free is Audacity¹². Figure 5.2 shows the screen of Audacity showing the precession signal from my backyard sensor in its enclosure. The computer used to run Audacity was an elderly laptop. The electronics used to get this signal employs the HEXFET switching described in section 4.2 and has an LT1037 as the low-noise preamplifier with a voltage gain of about 4000. There is no band-pass filtering used here. If you listen to a .wav file made with this setup, you can hear some 60 Hz interference. Narrow band-pass filtering will remove most or all of such low-frequency interference, however, it is wise to keep the electronics away from noise generating equipment like switching power supplies, computers or computer displays. Light dimmers are particularly notable for generating switching noise which then propagates through the wiring in a house.

Because it is difficult for the average amateur to design and build a front-end amplifier/polarization circuit such as that used to produce this signal, I am considering making and selling just this as a single printed-circuit board. This board would be functionally equivalent to the front-end board used in the MkII magnetometer design produced by myself and Willy Bayot¹³. In this design, there is short-circuit protection of the HEXFET's so even a short circuit of the sensor will not cause the HEXFET's to fail. There is a DIP switch to select capacitor values for tuning the sensor. The polarization sequence is initiated by a single line where +5V turns it on and 0V turns it off. A small microprocessor built into the board does all the sequencing of the HEXFET's. The board size is approximately 150mmx75 mm (approximately 6"x 3"). The board is designed to require a 12V battery or power supply as the source for both the polarization current as well as for operating the board during the signal amplification period. The user has to provide a sensor and a +12 Volt power source (battery or AC power supply). If there is sufficient interest (send an email to jark@shaw.ca), I

¹²<http://audacity.sourceforge.net/>

¹³<http://users.skynet.be/fa352591/index.htm>

will go ahead with this design and production. The selling price will be \$200 CDN plus shipping.

5.2 A field survey instrument

The main body of this document describes the main features of a field magnetometer and I will not repeat them here. There are two primary considerations:

1. the instrument must be easy to use so the human interface requires careful design,
2. there must be some method of non-volatile data storage.

Many people have the capability of modern microprocessor design and microprocessor based digital systems are easy to lay out and produce in printed-circuit board form. The more difficult design of the analogue circuitry necessary for the front-end of such an instrument can be circumvented by the purchase of a front-end board as described above which then just requires a digital back-end to provide the human interface and to display the measured magnetic field in nT.

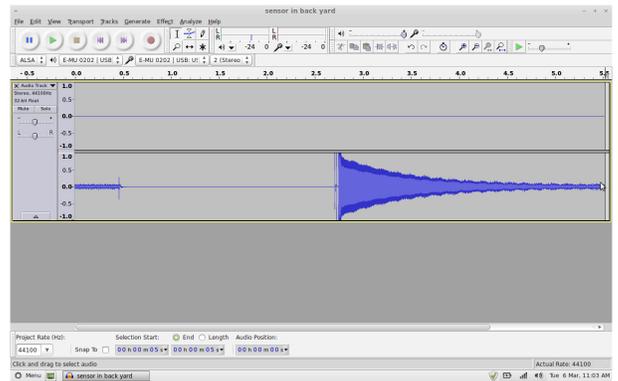


Figure 5.2: Audacity screen showing precession signal

References

- [1] Acker, F.E., Calculation of the Signal Voltage Induced in a Toroidal Proton- Precession Magnetometer Sensor, IEEE Trans. Geosci. Electronics, GE-9, 98-103, 1971
- [2] Aiken, M.J. and Tite, M.S., A Gradient Magnetometer using Proton Precession, J. Sci. Instrum., 39, 625-629, 1962 Bloch, F., Nuclear Induction, Phys. Rev., 70, 460-485, 1946
- [3] Burrows, K., A Rocket Borne Magnetometer, J. Brit. IRE, 19, 769-776, 1959 Cahill, L.J.
- [4] Van Allen, J.A., High Altitude Measurements of the Earths Magnetic Field with a Proton Precession Magnetometer, J. Geophys. Res., 61, 547-558, 1956
- [5] Faini, G., Fuortes, A. and Svelto, O., A Nuclear Magnetometer, Energia Nucleara, 7, 705-716, 1960
- [6] Faini, G. and Svelto, O., Signal-to-Noise Considerations in a Nuclear Magnetometer, Suppl. al Vol. 23, Serie 10 del Nuovo Cimento, 55-65, 1962
- [7] Haggard, L, Proton Precession Ferrous Metal Detector, Practical Electronics, 782-793, October, 1970
- [8] Packard, M.E. and Varian, R.H., Free Nuclear Induction in the Earths Magnetic Field, (abstract only), Phys. Rev., 93, 941, 1954
- [9] Primdahl, F., Scalar Magnetometers for Space Applications, Measurement Techniques in Space Plasmas: Fields, Geophys. Monograph 103, American Geophysical Union, 1998
- [10] Serson, P.H., A Simple Proton Precession Magnetometer, Dominion Observatory, Ottawa, Canada, 1962
- [11] Serson, P.H., U.S. Patent No. 3,070,745, 1962
- [12] Sigurgeirsson, Th., A Continuously Operating Proton Precession Magnetometer for Geomagnetic Measurements, Scientia Islandica, 2, 64-77, 1970
- [13] Stuart, W.F., Earths Field Magnetometry, Rep. Prog. Phys., 35, 803, 1972
- [14] Waters, G.S. and Phillips, G., A New Method of Measuring the Earths Magnetic Field, Geophys. Prospecting, 4, 1-9, 1956
- [15] Waters, G.S. and Francis, P.D., A Nuclear Magnetometer, J. Sci. Instrum., 35, 88- 93, 1958

- [16] Wadsworth, Nicholas, - a description of how to build a PPM in
The Amateur Scientist column of Scientific American, February,
1968 p62

Index

- B_p , 7
- accuracy, 18
- Acker, 9
- Advanced Signal Processing, 42
- Amplification, 25
- angular momentum, 2
- arc, 34
- Audacity, 47
- auroral zones, 28

- band-pass filter, 27
- basic operation, 18
- Bayot, 47
- beat frequency, 22
- Boltzmann's constant, 6

- Control, 23
- Controller, 32
- coupling loop, 39
- Curie equation, 9
- cylindrical solenoid, 11

- Danish Space Research Institute, 6
- Display, 33
- DSP, 43
- Dual solenoid connections, 30
- dual solenoidal sensor, 30

- E.J. Farrel and C.B. Grosch, 43
- earth's magnetic field, 2
- Electronics, 23
- external ac magnetic fields, 4

- Faini and Svelto, 10
- Field Instrument, 32
- figure of merit, 15
- filling factor, 11
- Fourier analysis, 44

- γ_p , 2
- gel-cell, 29
- geomagnetic observatory, 46
- GPS, 44

- gradiometer, 22

- HEXFET, 35
- human input, 33

- interference, 45

- Jian Qiu Zhang *et al*, 43

- Koehler, 43

- Larmor Equation, 2
- Levenberg-Marquardt, 44
- Lithium, 29
- LM394, 26
- LM837, 26
- local magnetic field, 44
- low-noise techniques, 26
- LT1037, 26

- magnetic moment, 2
- magnetic sub-storms, 28
- magnetization, 6
- magnetometer, 3
- Measuring Inductance, 39

- Ni-Cd, 29

- Optimization, 14
- orientation problem, 5
- Overall Considerations, 28

- permeability, 1
- Physics of Proton Precession, 2
- Polarization, 23
- Polarizing Circuit, 33
- polarizing magnetic field, 3
- Possible polarization circuits, 34
- PPM's, 4
- Practicalities, 28
- Primdahl, 6

- Q, 25

- Relay-less polarization, 40

resonate, 38
resonating capacitor, 39

Series or parallel sensor configuration, 25
SI, 1
Signal strength calculation, 6
Signal-to-noise, 13
snub, 34
solenoid, 4
solenoidal configuration, 10
Solenoidal sensors, 30
spin-lattice relaxation time, 23
spin-spin relaxation time, 24
spreadsheet, 16
Summary of signal strength calculations, 12

T1, 23
T2, 24
Tesla, 2
toroidal configuration, 9
toroidal core, 5
Toroidal sensors, 31

William H. Press, *et al*, 44

Zener diode, 34