

Proton Precession Magnetometers, Rev 2

J.A. Koehler

Comox, BC, Canada

November, 2004

Chapter 1 The Physics of Proton Precession and Sensor Geometry

A few words about units

When I was a student in physics in the 1950's, 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 1960's 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.

The unit of \mathbf{B} in the SI system is the Tesla; named after the famous/infamous (take your pick) Serbian scientist/charlatan (take your pick) Nikola Tesla. The earth's magnetic field has a value of about 50 microTeslas (5×10^{-5} T) but varies from location to location over the surface of the earth being (in general) strongest near the magnetic poles and weakest 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. 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 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 5.27×10^{-35} kg-m²/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-m². 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, $\boldsymbol{\mu}$, (note the bold, italics here) 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 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. A derivation of the Larmor Equation is given in Appendix A for those interested.

For $B = 50 \mu\text{T}$ (a nominal value for the earth’s field), ω is therefore 1.34×10^4 radians/second which (divide by 2π) corresponds to 2130 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, 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.

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' as quickly as possible. 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 is 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.

Is a solenoid the best configuration?

Most amateur PPM's¹ (proton precession magnetometers) have been solenoidal. 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 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.

Most 'professional' PPM's² are made using a toroidal core. 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.

The first disadvantage of the solenoid, its 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 (but not too close!), the induced voltages due to external ac noise will be greatly reduced. However, the orientation problem still exists.

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.

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

² F. Primdahl, 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.

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.

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 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. 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 there is an approximation here.

Let us now calculate the total 'magnetization' of the fluid. Let the density of the fluid be $d \text{ kg/m}^3$. Let the molecular weight of the fluid be $w \text{ amu}$ where each amu is equal to $1.67 \times 10^{-27} \text{ kg}$. 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,

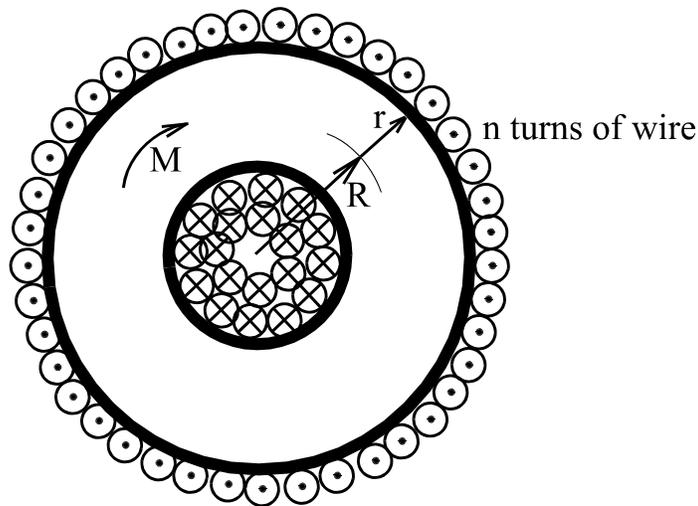
$$\frac{\mu B_p}{kT}$$

is aligned, the total magnetization, M , will be approximately (see *caveat* above):

$$M \approx N\mu \left(\frac{\mu B_p}{kT} \right)$$

M is in units of $A\cdot m^2$ per cubic meter or 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.



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

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

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 as shown in the figure. Note that the above equation is an approximation. The magnetic induction inside a toroidal core varies with the radius so the above approximation is valid, as an average value, provided that r is much, much smaller than R . In the figure of the toroid, the usual convention is followed in showing the direction of the current in the windings by a point if it is coming out of the paper and by a cross if it is going into the paper.

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 Faraday's 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:

$$M_{\text{axial}} = M \cos(\omega t)$$

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

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

From Faraday's Law, the time varying induced voltage is given by:

$$e(t) \cong -nA \frac{dB_{\text{axis}}}{dt} = nA\mu_0\omega M \sin(\omega t)$$

where A is the cross-sectional area of each loop of the coil. Let the radius of the toroid core be 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:

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

Note that the output voltage depends on the square of the number of turns and on the magnitude of the polarizing current. Also, r should be large compared to R which means that a 'fat' toroid is better than a 'thin' one. Of course, if the toroid is made too 'fat', the central hole will be too small to allow many turns of wire.

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, in fact 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 total magnetic moment of the medium due to a polarizing field and this contained an approximation:

$$M \approx N\mu \left(\frac{\mu B_p}{kT} \right)$$

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

$$M = \chi H$$

χ is called the ‘magnetic susceptance’ of the substance and the equation above gives the relationship between the resulting magnetization, M , produced by a substance by an external magnetic field, H . For non-ferromagnetic substances, this equation can be rewritten:

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

which means that the following equation should be true:

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

How far wrong were we in the approximation we made? 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 susceptance also and so the bulk measurement of the susceptance 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 susceptance 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 susceptance of a substance that is due to the hydrogen atoms (i.e., protons) alone. This 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 $h/4\pi$ 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 susceptance 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. The values for some common liquids is given in Appendix C.

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 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 as the equation derived in the section above but with the substitution that $(R - (R^2 - r^2)^{1/2})$ replaces $r^2/2R$. (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

³ 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

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. This is a much better situation than in the case of a solenoidal sensor where the weakest signal is zero!

The solenoidal configuration

The solenoidal configuration has been studied by Faini and Svelto⁴. 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:

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

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. The quantity, η , is called the ‘filling factor’ and depends on the geometry of the coil. It is the calculation of this 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 solenoid radius to the solenoid length and which they just show in the form of the graph reproduced below.

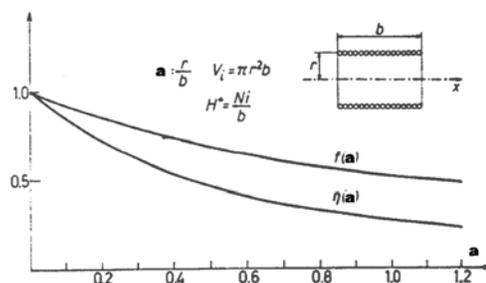


Fig. 3. - Approximate values of $\eta = (1/V_i) \int [(H_x^2 + H_y^2)/H^{*2}] dV$ for a single-layer coil. In the figure also value of $f(a) = (1/V_i) \int H_x^2/H^{*2} dV$ are given. To a very good approximation it can be assumed $\eta(a) = f^2(a)$.

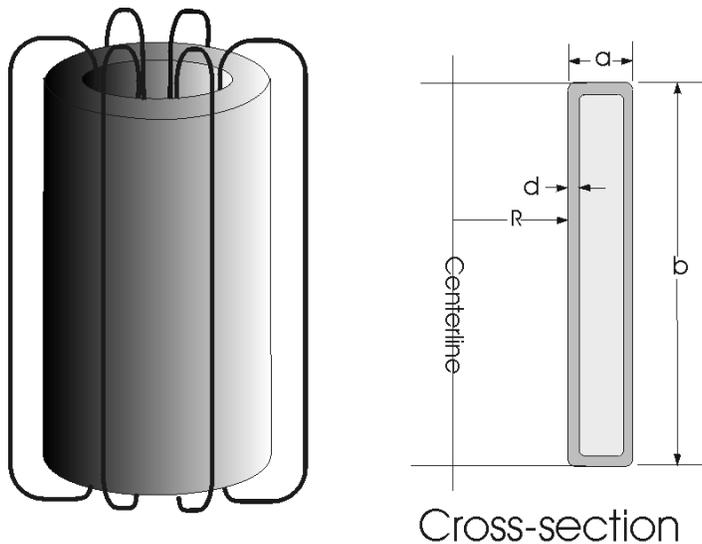
⁴ Faini G. and Svelto O., Signal-to-Noise Considerations in a Nuclear Magnetometer, Suppl. al Vol 23, Serie 10 del Nuovo Cimento, 55-65, 1962

The abscissa of this graph is the ratio of radius to length for the solenoid.

I have made a slight change of notation. They gave the symbol α to the ratio of 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.

A cylindrical toroidal configuration

There is another configuration which is potentially useful for PPM's designed to be towed behind boats. It is a toroid with rectangular cross-section 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.



R is the inside radius

d is the wall thickness

b is the outside cylinder height

a is the cylinder thickness

The magnetization produced by the polarization field inside a toroidal form is:

$$M = \chi B_p$$

where B_p is the polarization field which is a function of radius. We have discussed the valuation of χ earlier. B_p is 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:

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

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

$$\phi = \frac{\chi \mu_0 (b - 2d) n I_p}{2\pi} \int_{R+d}^{R+a-d} \frac{dr}{r}$$

which is evaluated as:

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

Again, assuming that the precession is sinusoidal, the generated voltage can be shown to be:

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

Summary of signal strength calculations

Let us simplify these equations for the toroidal, solenoidal and cylindrical toroidal sensors. Firstly, let's just calculate the value for the best orientation – where the external field which it is desired to measure is along the axis of the toroid and cylindrical toroid or perpendicular to the axis of the solenoid. Then, all the equations can be expressed by a single equation:

$$e_s = \frac{\chi\omega\mu_0}{\sqrt{2}} n^2 I_p G = \kappa n^2 I_p G$$

where κ is a constant and G is a geometric factor which is different for the solenoid, the toroid and the cylindrical toroid, n is the number of turns of wire and I_p is the polarizing current in Amperes. The constant, κ , which is equal to $\chi\omega\mu_0/\sqrt{2}$, has a value of 4.87×10^{-11} in SI units for water in a location where the external magnetic field is 50 μT . Its value is proportional to the magnetic field strength which the PPM is expected to measure.

$$G_{\text{toroid}} = R - \sqrt{R^2 - r^2}$$

where R is the central radius of the toroid (from the central axis to the middle of the 'donut' core) and r is the radius of the cross-section and,

$$G_{\text{solenoid}} = \frac{\eta V_i}{b^2}$$

where b is the length of the solenoid, η is the filling factor and V_i is the internal volume of the solenoid.

For the cylindrical toroid, the value of G is:

$$G_{\text{cylinder}} = \frac{(b - 2d)}{2\pi} \ln\left(\frac{R + a - d}{R + d}\right)$$

where R is the inside radius of the cylinder, a is the thickness of the cylinder (the difference between the inside and outside radii), b is the length of the cylinder and d is the wall thickness. Normally, d will be small compared to the other values and, if so, G_c becomes simplified to an approximation:

$$G_{cylinder} \approx \frac{b}{2\pi} \ln\left(\frac{R+a}{R}\right)$$

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:

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

where R_w is the resistance of the wire used in the sensor and Δ 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 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:

$$R_w = \frac{1.73 \times 10^{-8} \lambda}{\sigma}$$

where λ is the length of the wire making up the coil and σ 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. For multi-turn windings, the mean value 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 wire's 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 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:

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

Finally, note that there are some approximations here. There is the use of the mean radius in multi-layer windings and, 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 300⁰ K.

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, if one designed a toroidal PPM sensor with $R = 0.10\text{m}$, $r = 0.05\text{m}$ (for our American readers, this is approximately an outside diameter of the 'donut' core of 12 inches, with a 4" diameter donut thickness), wound it with 2000 turns of wire with a diameter of 1.67 mm (#14 AWG). If the polarizing current were 10 Amperes, the rms signal voltage due to the earth's magnetic field in the most favourable orientation would be about 40 microV. The noise voltage would only be about 3.5 nV (assuming a bandwidth of the detector of 100 Hz) giving an outstanding S/N ratio of about 11,500!

Of course, this would be a pretty big (and heavy!) toroid containing several kg of water and it would require several kg of wire to wind it. The dc resistance of the winding would be about 6 Ohms so to get a polarizing current of 10 Amperes through it, you would need to have a 60 Volt supply. The unit would therefore consume about 0.6 kW for the period (a few seconds) that it was being polarized. Such a unit would obviously be very impractical for a back-pack gradiometer where you need two such sensors. Imagine carrying around 60 V of batteries capable of supplying 1.2 kW for each measurement 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 time in discussing 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, η .

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 product of the polarizing voltage and polarizing current:

$$\text{FOM} = \frac{\frac{S}{N}}{V_p I_p}$$

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’, one would write an expression for the desired figure-of-merit with the dominating consideration as the denominator, and take the differential with respect to the denominator 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. 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 both the solenoidal and the toroidal PPM sensor. The file is named *ppm.xls*. I have also made a spreadsheet for the cylindrical toroidal shape and that spreadsheet is named *cyl.xls*. In these spreadsheets, one can vary the dimensions of the sensor, the number of turns of wire, the size of the wire, etc and look at

the calculated S/N. I have used an approximate expression giving the radius of the wire (in m) as a function of the AWG. I have also used an expression, for the solenoid, giving the approximate value of the filling factor as a function of mean solenoid radius divided by solenoid length. These are the only approximations used in the calculations and should be accurate to better than about 5%. 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 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.

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.

Running the Spreadsheets

Wire AWG Radius	# turns	Length (m)	G factor	L (mH)	Res. Ohms	Q	C (nF)	Signal (uV)	SN	Polarizing Current (A)	R _{mean}	
14	6.35E-04	726	107.4	6.639E-03	4.40	0.85	58.0	1270.2	2.4	1825	14.18	2.35E-02
16	6.65E-04	970	135.2	5.735E-03	6.91	1.69	45.7	809.4	1.9	1019	7.12	2.20E-02
18	5.19E-04	1325	172.1	5.029E-03	11.09	3.52	35.1	503.8	1.4	543	3.41	2.07E-02
20	3.96E-04	1901	221.6	4.488E-03	18.29	7.67	26.6	305.7	1.1	278	1.56	1.96E-02
22	3.04E-04	2440	287.2	4.087E-03	30.57	17.09	19.9	182.8	0.8	140	0.70	1.87E-02
24	2.35E-04	3231	367.8	3.808E-03	49.96	36.62	15.2	111.9	0.6	73	0.33	1.81E-02
26	1.92E-04	4024	448.2	3.638E-03	74.03	87.24	12.3	75.5	0.5	43	0.18	1.77E-02
28	1.74E-04	4470	493.2	4.185E-03	105.07	90.22	13.0	53.2	0.5	40	0.13	1.76E-02

Toroidal Sensor Wire radius = $2.749e-3 - 1.813e-4 \cdot \text{AWG} + 3.19e-6 \cdot \text{awg}^2$

Beath (T) = 5.00E-06 NOTE - above formula is only approx. correct for wire sizes less than about #28
 Ch = 4.02E-09
 Frequency = 2.13E+03 r = 0.016
 Kappa = 4.77E-11 R = 0.045 Numbers applicable to toroid

V = 12 Density (kg/cu. m) = 1000 6.65336E+28
 Bandwidth = 100 Molecular weights = 18 Numbers for spreadsheet as a whole
 Layers = 9 Protons/molecule = 2

Wire AWG Radius	# turns	Length (m)	Filling Factor	L (mH)	Res. Ohms	Q	C (nF)	Signal (uV)	SN	Polarizing Current (A)	R _{mean}	
14	6.35E-04	538	104.5	0.48	37.83	1.53	27.5	147.7	3.8	2172	7.82	5.75E-02
16	6.65E-04	877	238.1	0.46	56.69	2.87	213	95.6	3.2	1304	4.04	5.60E-02
18	5.19E-04	867	297.8	0.47	88.66	6.09	162	63.0	2.6	741	1.97	5.47E-02
20	3.96E-04	1128	379.6	0.47	144.22	13.14	122	38.8	2.1	402	0.91	5.36E-02
22	3.04E-04	1479	489.9	0.48	240.05	26.14	82	23.3	1.6	212	0.41	5.27E-02
24	2.35E-04	1913	626.4	0.48	392.41	62.37	70	14.2	1.3	114	0.19	5.21E-02
26	1.92E-04	2348	763.1	0.48	582.41	114.49	57	9.6	1.0	70	0.10	5.17E-02
28	1.74E-04	2593	840.0	0.49	705.61	153.65	51	7.9	1.0	55	0.08	5.16E-02

Solenoidal Sensor Filling factor = $\exp(-1.4 \cdot r_{\text{mean}}/\text{length})$

Length = 0.10
 Inside radius = 0.05
 Gsol = 0.07654 Numbers applicable to solenoid

The spreadsheet, *ppm.xls*, shown above has two sections – one for toroidal sensors and the other for solenoidal ones. There is a large section at the left middle in the ‘Toroidal’ area where various constants are shown – these apply to both sections. For example, one might choose to use a fluid with a different value for density and molecular structure and the appropriate values could be entered here. Here is also where one can vary the number of layers of wire, the polarizing voltage and the detector bandwidth and observe the changes in S/N for the two configurations.

Above it and to the right is a small area which fixes the geometry for the toroidal sensor – there are only two numbers here and those are the central radius of the toroid, R, and the radius of the toroidal cross-section, r. The values given the spreadsheet at the moment are those of the blue toroid making part of the Fisher-Price ‘Rock-a-Stack’ toy.

At the bottom of the solenoidal section, there is the similar section where one may enter the values for the solenoidal geometry. These numbers are the length of the solenoid and its internal radius. There is also shown, in this vicinity, a value for the G factor appropriate to the solenoid but this is a calculated value, calculated from the length and radius and so not a number which can be independently entered.

Finally, the spreadsheet calculates the approximate inductance of the winding and also the approximate capacitance needed to resonate at whatever value has been entered for the earth’s field. The word approximate means just that!

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 spreadsheet used for the cylindrical toroidal configuration. It is named *cyl.xls* and can be used similarly to those described above. A screen snapshot of this spreadsheet is shown below. Here, the number of turns of wire is used rather than the number of layers.

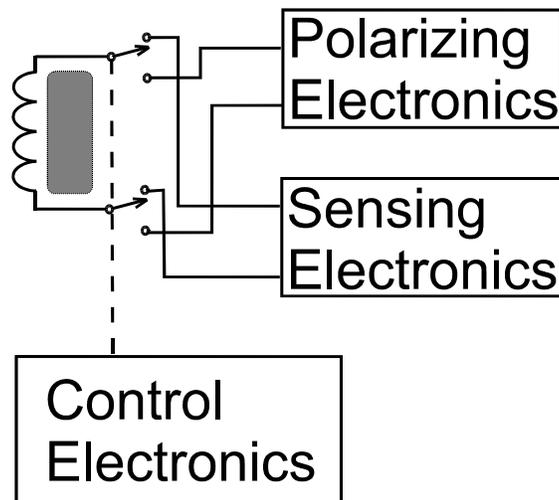
	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	Wire	Wire	# turns	Length	G	L (mH)	Res.	Q	Wire weight	Signal	S/N	Polarizing	Thickness	
2	AWG	Radius		(m)	factor		Ohms		(kg)	(uV)		Current (A)	(m)	
3	14	8.36E-04	1400	414.2	1.059E-02	30.46	3.27	104.0	8.1	3.6	714	3.67	7.94E-03	
4	16	6.65E-04	1400	406.3	1.059E-02	31.33	5.07	68.9	5.0	2.3	370	2.37	5.12E-03	
5	18	5.19E-04	1400	400.9	1.059E-02	32.00	8.19	43.5	3.0	1.5	180	1.46	3.16E-03	
6	20	3.99E-04	1400	397.3	1.059E-02	32.47	13.75	26.3	1.8	0.9	83	0.87	1.89E-03	
7	22	3.04E-04	1400	395.1	1.059E-02	32.77	23.50	15.5	1.0	0.5	37	0.51	1.10E-03	
8	24	2.35E-04	1400	393.9	1.059E-02	32.94	39.22	9.4	0.6	0.3	17	0.31	6.61E-04	
9	26	1.92E-04	1400	393.2	1.059E-02	33.03	59.00	6.2	0.4	0.2	9	0.20	4.39E-04	
10	28	1.74E-04	1400	393.0	1.059E-02	33.06	71.89	5.1	0.3	0.2	7	0.17	3.61E-04	
11	Cylindro-toroidal Sensor								Wire radius = 2.749e-3 - 1.813e-4*AWG + 3.19e-6*awg^2					
12	Bearth (T) =	5.00E-05				d =	0.003		NOTE - above formula is only approx. correct for wire sizes less than about #28					
13	Chi =	4.02E-09				b =	0.1							
14	Frequency =	2.13E+03				a =	0.04							
15	Kappa =	4.77E-11				R =	0.03		NOTE: Thickness is the approximate thickness of the outer winding.					
16														
17	V =	12			Density (kg/cu. m)=	1000		6.65336E+28						
18	Bandwidth =	400			Molecular weight=	18								
19	Number of turns	1400			Protons/molecule=	2								
20														
21														

Chapter 2 The Electronics and Measurement of Magnetic Fields

The basics of the operation of a PPM

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 block diagram form, this is shown in the figure below.



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.

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), 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:

$$e(t) = A \sin(\omega t) \pm a \left(\frac{1}{\sqrt{2}} \cos(\omega t) + \frac{1}{\sqrt{2}} \sin(\omega t) \right)$$

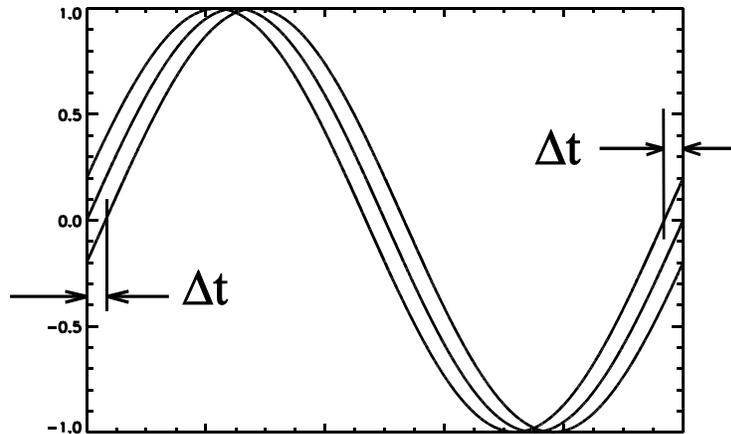
where A is the amplitude of the desired signal and a is the amplitude of the noise (the sum of the sensor noise plus the additional noise added by the amplifier). A/a is the actual signal-to-noise ratio achieved by the system. I have shown the noise as the sum of an in-phase component (the sine term) 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:

$$e(t) = \sqrt{A^2 + a^2} \sin(\omega t \pm \varphi) \cong A \sin(\omega t \pm \varphi)$$

where

$$\varphi = \tan^{-1} \left(\frac{a}{\sqrt{2}A} \right) \cong \frac{a}{\sqrt{2}A}$$

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 the drawing below.



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

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

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 Δt at the beginning of the interval and a similar uncertainty of Δt 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 $\sqrt{2}$ times this $2/3$ of a unit – altogether about one unit of Δt . The relative error, therefore, in measuring the frequency will be $\sim \Delta t/T$. 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 in Appendix D. As an example, suppose frequency (2130 Hz) was measured over a total period of $1/2$ second and that, at the end of this interval, the overall S/N of the system was 100 (it will be better at the beginning because the signal decays exponentially). Then, Δt works out to about 0.5 microS (assuming the frequency to be 2130 Hz and the overall best case error will therefore be about one part in 1 million. If the field being measured were $50 \mu T$, then the error would be just 50 pT! Wow!

This looks really good since, from the S/N calculations in the previous chapter, it looks like S/N values of ~ 100 are fairly easily obtainable. However, now we need to tackle the subject of actually measuring the frequency – what techniques are there and what have 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 just be plus or minus one part in 2130 – about plus or minus \sim 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.

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 – 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 plus or minus one part in;

$$\frac{Nf_c}{f_p}$$

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 plus or minus one part in $\sim 10^6$ – this corresponds to plus or minus 50 pT for a 50,000 nT field. This is about twice the theoretical limit for a S/N of 100 measured over the \sim 1 second that the 2048 periods of the signal takes.

Because we don't want to measure frequency but rather the magnetic field in nT, we may rearrange the above equation to give us the measured precession frequency. Then, using the relationship between this frequency and local magnetic field, we get:

$$B = 23.4875 \frac{Nf_c}{n}$$

where N is the number of precession periods over which the clock is counted, f_c is the clock frequency and n is the number of counts of the clock frequency observed. 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:

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

where α is just $23.4875Nf_c$. Again, using the previous example, N would be 2048 and f_c would be 10^6 giving us an α of 4.81024×10^{10} .

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 IC’s. 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 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 in use in some gradiometers. A gradiometer gives the difference between the magnetic field at two different sensor locations. If PPM’s 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 two⁵. The beat appears when the two signals are electronically added 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, the stability of

⁵ see, for example, the ‘Delta² Magnetometer by M.L. Dalton Research, 6035 Aberdeen Dallas, TX, 75230. Ph: (214) 691-4925

the clock is no longer as important – 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 optimum system for measuring sensor frequency is one in which the period is measured using a high-speed clock. The 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 gradiometer for the earth's magnetic field which is accurate to at least ~ 1 nT is readily achievable.

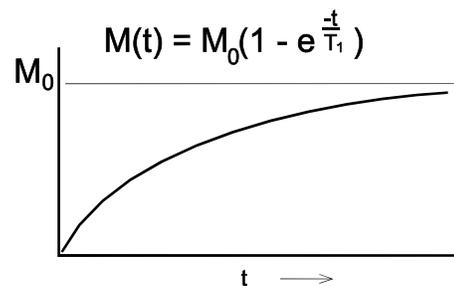
Electronics functions

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.

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 below.

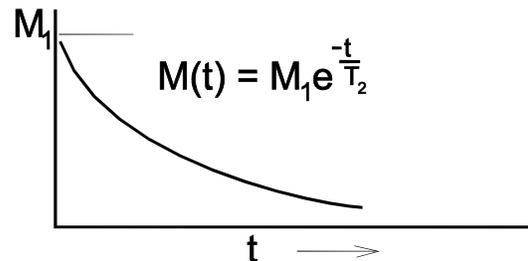


T_1 is 2 to 3 seconds for water; about 0.5 seconds for kerosene, for example. In order to get close to the desired magnetization, the polarizing current needs to flow for several times this relaxation time. For example, after 4 time constants (say, 10 seconds, for water), the magnetization will have 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. 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 below:



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.

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

$$\frac{LI_p^2}{2}$$

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.

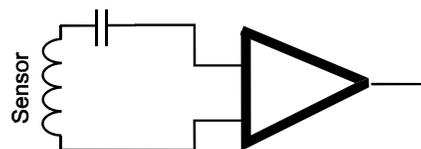
The nature of the switch and how it is turned off are important design considerations in the overall scheme of things. We will look at this in more detail later when we consider actual circuits.

Amplification

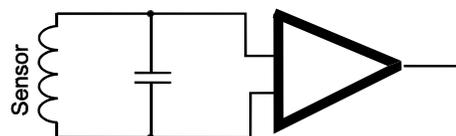
The desired signal is normally very small – of the order of microvolts only. Therefore it is very important that the amplifier not add any more noise than necessary. Fortunately, the study of low-noise amplifiers in the audio range is well advanced because of the preponderance of audio devices being 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 ~ten Ohms. It is also inductive with typical inductances of a few ten's of mH. As stated before, the typical output voltage levels are a few microVolts.

Since the signal needs to be coupled to an amplifier anyway, the coupling capacitor is normally chosen to have a value which resonates with the desired signal frequency. I will discuss the effect of resonating the signal in this way in a later section. Since the overall Q of the L-C circuit is determined almost entirely by the coil and since these 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 circuit may be arranged in two ways depending on whether one has a low impedance amplifier or a high impedance amplifier.



**Series configuration
needs low input
impedance amplifier**



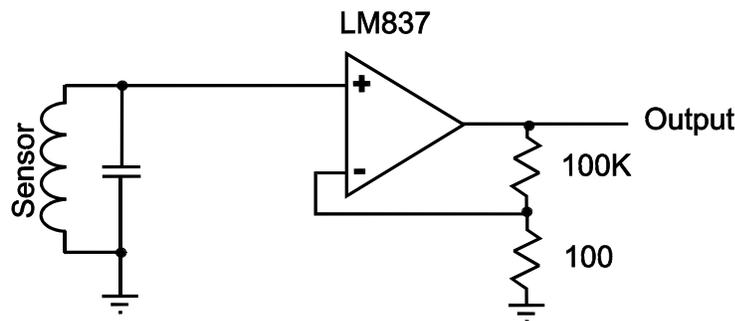
**Parallel configuration
needs higher input
impedance amplifier**

These two alternatives are shown in the figure below.

In both circuits, the capacitor is resonant with the sensor at the signal frequency (2130 Hz for a 50 μ T field). 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 were ~ 7 Ohms, the effective resistance 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 ~ 6.3 K.

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 Corporation⁶ 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.

It is possible to characterize the noise performance of an amplifier by defining a noise figure, NF. The NF is defined as the input S/N divided by the output S/N; it is usually expressed in dB. Obviously, the lower the NF, the better. For PPM's, we need an amplifier with a low NF 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 NF of an amplifier 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 the circuit shown below, the calculated NF is about 1.4 meaning that the S/N from the sensor is degraded by just 40%.



⁶ AN-104 gives a good description of how to calculate the noise figure. I recommend it for anyone interested in finding out more about amplifier noise.

The circuit has a voltage gain of 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.

Although I have discussed a particular low-noise op-amp as a suitable input amplifier, some discrete transistors offer an even lower noise figure. The ultimate in low-noise preamplifiers can be built using specially constructed discrete transistors such as the National Semiconductor LM394. National Semiconductor has an application note for this device, AN-222, which is very helpful.

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? If the PPM is to be used in the field, the local value of the earth's field can vary between ~45 to ~60 μ T. The corresponding sensor frequencies will therefore vary from about 1.9 KHz to 2.6 KHz. From the viewpoint of getting a signal at some random location on the earth's surface, then, one would like the bandwidth to be no smaller than 700 Hz. However, from the viewpoint of enhancing the output S/N, one would like the bandwidth as narrow as possible. Thus, we have two conflicting requirements.

This can be solved in the following way. Using a microprocessor, one can design a system with a variable band-pass filter which is digitally controllable. Then, a preliminary sample of the data with a wide bandwidth can be used to determine the approximate frequency. Then the bandwidth can be reduced to make the final measurement. As I have showed in a previous section, it is possible to make a measurement of the precession signal period with an accuracy corresponding to better than 1 nT in about 0.12 seconds. Since the precession signal decays with a time constant of several seconds, there is sufficient time to make a preliminary estimate of the frequency, then set the filter to the correct frequency, narrow the bandwidth and then make another measurement with the narrow bandwidth all within one polarize/measure cycle.

Chapter 3 Practicalities of making a Magnetometer or Gradiometer

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 ten's 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 used. 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. Typically, 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. The observed difference between the magnetic field measured at each sensor will then be a measure of the vertical gradient caused by material under the soil in that vicinity.

Another technique is to locate one of the two sensors in a fixed location and to connect the other sensor through a long cable. The other sensor can then be located at a number of places in the vicinity of the fixed sensor thus mapping the magnetic differences around the fixed sensor.

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. For the purposes of this discussion, let us assume that we want to make a sensitive gradiometer.

Perhaps the most important consideration is selection of the size and weight of the finished magnetometer head(s). For a gradiometer, there will be two of them. 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(s). 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 want to use a 'gel-cell' rechargeable battery because it is not magnetic and still gives a reasonable energy per unit weight.

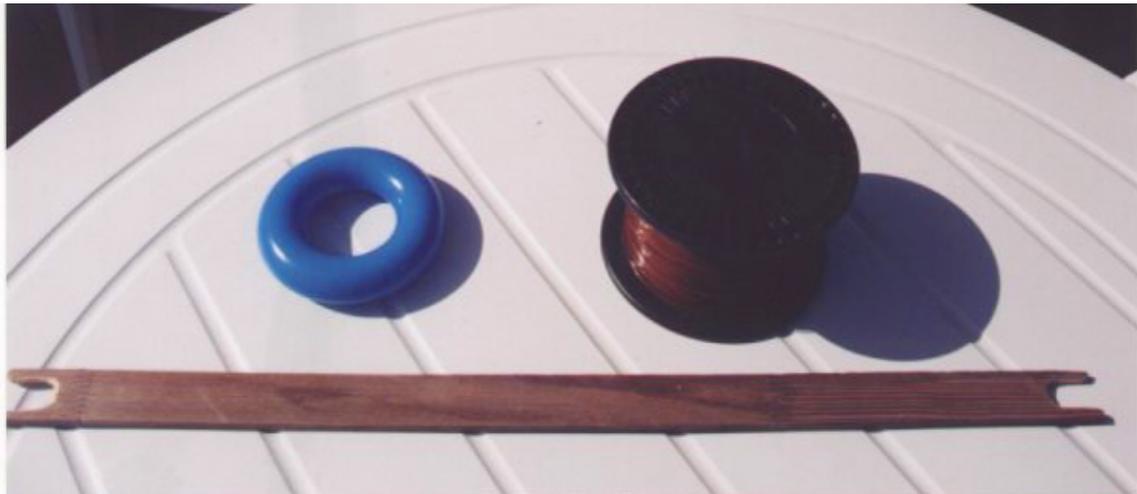
Let me work out one example. Suppose you want to make a portable gradiometer suitable for back-packing. You decide to make toroidal sensors wound on cores made from the 'blue' toroid in a 'Rock-a-Stack' toy filled with distilled water. If you wind about 10 layers on these cores with #20 AWG copper wire, and use a polarizing voltage of 12 V, you can expect a S/N of about 300. Each core, after being wound, has a total mass of ~1.2 kg. The dc resistance is about 7.5 Ohms so the current drawn is about 1.6 A and, if the time you polarize the core is 6 seconds, it will consume about 115 Joules of energy during this period. Because there are two cores in the gradiometer, the instrument will consume about 230 Joules per measurement. 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 can make $\sim 60,000/230$ measurements on each charge of the battery; about 260. In reality, you won't be able to make even this many because the battery's rated capacity is given for a much lower discharge rate. 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 100 measurements, perhaps as many as 200.

You have to decide whether one to two hundred measurements is sufficient. If you're staying in a motel every night and just hiking out during the day, you can recharge the battery every night 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 with a much smaller wire 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.

Winding the coil

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. It is my opinion, however, that it is much much better to use a toroidal core because of its overwhelming advantages.

Making a toroidal sensor has a number of difficulties. The first problem is finding a suitable form. The best thing I have been able to come up with is the hollow blue ring which is the largest ring on the Fisher-Price 'Rock-a-Stack' toy. It can be filled with about 0.18 litres of fluid; this is a bit small, a greater volume would have been more desirable. The next difficulty in making a toroidal sensor is winding the requisite number of turns on the form.



The above picture shows 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 is convenient to put a rubber band around the wire at each end of the shuttle – this keeps the wire from falling off the shuttle too easily as it is passed through the centre of the toroid form.

It is 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 is smooth and uniform. The second layer is more difficult because it will not fit nicely over the first layer on the inside. Each successive layer gets 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 next picture shows the sensor when completed with about ten layers of wire.



This sensor contains about 2000 turns of the #20 AWG enameled wire and has 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 #20 AWG 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.

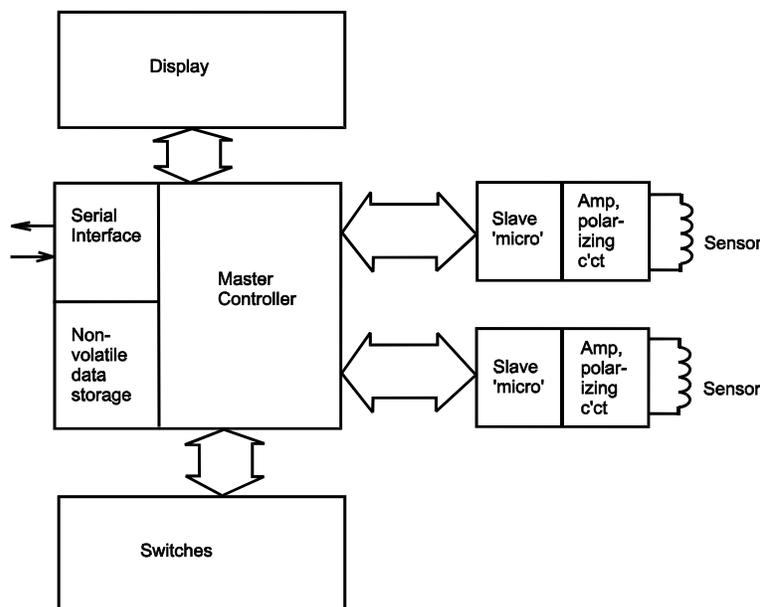
Other Features in a Field Instrument

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.

Overall Block Diagram

The overall block diagram for a modern gradiometer is shown below.



There are, of course, many ways of implementing the required functions and the above block diagram is just one such way. For example, many modern 'single-chip' microprocessors have enough internal capacity to implement most, if not all, of the

functions in a single unit. However, it is simpler both in concept and in implementation to divide up the various sub-functions in the way shown.

The Slave Units

The function of the slave unit is to make a measurement (i.e., polarize the sensor, measure the period) upon the command of the master controller. This unit contains its own microprocessor, a digitally controlled band-pass amplifier, the low-noise preamplifier and the polarizing circuit. The interface between the slave and the master is very simple: the slave receives a single command to make a measurement and returns a single value which is the measured magnetic field strength in nT

The Master Controller

The master controller provides the interface between the slave units and the outside world. It has a serial interface so that data from its internal storage can be sent out to, for example, a laptop computer for further analysis. It sends the commands to the slaves and receives the measured values from them. It interacts with the human operator by sensing the state of the switches and by providing output to the display.

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

Serial Interface

This can be a simple RS-232⁷ interface to a personal computer. If the internal storage memory of the instrument itself is 64 KB, then at 9600 baud, the entire memory contents can be dumped out into a personal computer in about a minute. That means that an extremely simple interface protocol can be devised: perhaps as simple as the reception of a single byte by the instrument from the external computer triggering a memory dump.

⁷ RS-232 really refers to a 25-pin interface. Most modern personal computers have a simpler 9-pin interface which implements some of the RS-232 protocol – and this is more suitable for this application.

Switches

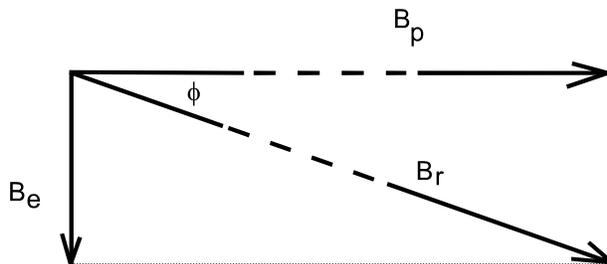
The switches provide the human input to the instrument. This portion needs to be carefully thought out. Any field instrument has to be extremely simple to operate. However, people get tired after a day in the field so it should not be too easy, for example, to erase the instrument's internal memory. 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 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.

The Polarizing Circuit

The polarizing circuit deserves a bit 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 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 L/r – typically $L \sim 10$ mH and $r \sim 10$ Ohms so this time constant is of the order of a few 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 sensor's magnetic field by the polarizing current.

Let's 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 the figure below.



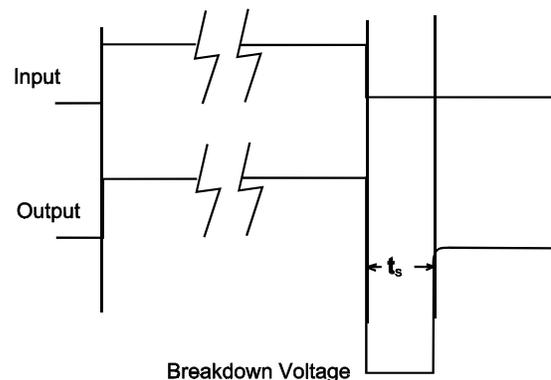
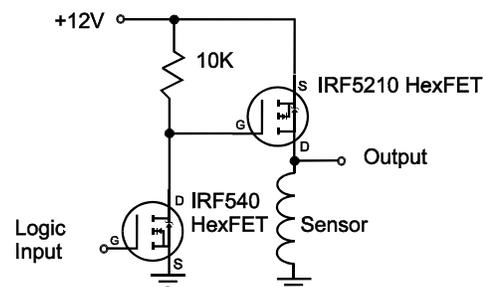
Here, B_e is the earth's field, B_p is the polarizing field and B_r is the resultant of the two. Typically, when the polarizing current is flowing, B_p is much, much greater than B_e and so the angle, ϕ , is very small – a fraction of a degree. So B_r is more or less parallel to

B_p . As the current is turned off, 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 ϕ becomes ~ 10 degrees is the beginning of the ‘effective’ turn-off period. Let this period be T_s – the ‘switching’ period. It starts when ϕ is ~ 10 degrees which is when:

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

Therefore, we may consider the turn-off switching time to start when B_p is $\sim 6 \times B_e$ which is $\sim 300 \mu\text{T}$. Let us call the current flowing through the inductor at the moment to be I_s . T_s is thus the time from when the polarizing current goes from I_s to zero.

Now, let us look at a possible circuit for switching the polarizing current on and off. HEXFET’s are used because they have such a low on-resistance that the losses through them are negligible. The first switching transistor could just be a simple bipolar NPN switching transistor instead of a HEXFET but would require a few more resistors in the base circuit and would require slightly more drive.



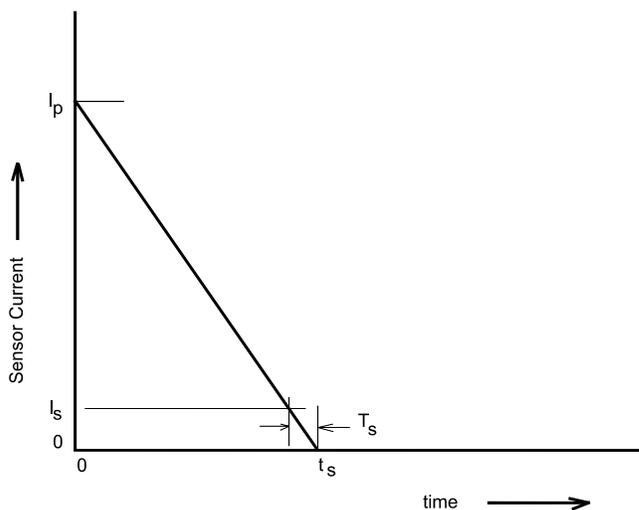
What happens when the current is switched off? When the IRF5210 is turned off, the magnetic field of the sensor starts to collapse. This induces a voltage across the sensor

which is negative at the ‘Output’ end with respect to ground. If it were a physical switch instead of transistor, the voltage would go high enough to cause some arcing across the terminals of the switch. In the circuit above, the drain of the IRF5210 starts to break down at its rated break-down voltage which is, for this particular transistor, 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. During this period, the entire stored energy in the inductor is dissipated in the switching transistor and in the dc resistance of the sensor. The final change in current when the voltage across the sensor goes to zero is so abrupt, that the sensor will ‘ring’ at its self resonant frequency (the resonant frequency of its inductance along with its’ stray parallel capacitance). This ringing is normally quenched by placing a resistance of a few hundred times the dc sensor resistance in parallel with the sensor – typically, a few K for a ~ 10 Ohm sensor. This large resistance in parallel with the sensor will have no significant effect on the sensitivity.

During the breakdown period, the voltage across the inductor is approximately constant and so the rate of change of current with time is approximately linear (the approximation is because we’ve neglected the resistance of the sensor). For an inductor, the terminal voltage is related to the current through the device by:

$$V = -L \frac{dI}{dt}$$

So we would expect the current to decrease to zero approximately as shown in the graph below.



t_s is the total time to collapse the field. The question is: what is T_s , the time for the sensor current to go from I_s to zero?

Let's calculate this for the example toroidal sensor described in this chapter. Here I_p was 1.6 A. Observing the turn-off with an oscilloscope, t_s is approximately 0.2 mS⁸. From the equation relating the centre field in a toroidal inductor magnetic field to the current flowing through it, the current flowing when B_p has fallen to 300 μ T is:

$$I_s = \frac{2\pi R}{\mu_0 n} B_p \approx \frac{2\pi \times 0.05}{4\pi \times 10^{-7} \times 2000} \times 300 \times 10^{-6} \approx 20 \text{ mA}$$

Since the current is decreasing at 1.6 A in 0.2 mS (8000A/s), it will take about 2.5 μ S to fall the last 20 mA. This easily meets the requirement that the time be shorter than about 50 μ S.

Many existing circuits use a Zener diode (in series with an ordinary diode so as not to conduct when the polarizing voltage is connected) across the sensor to limit the breakdown voltage – typically the Zener rating is a few times the supply voltage. What does this do? The answer is that the Zener diode provides the limiting voltage at which the sensor's inductive energy is dissipated. In general, because dI/dt is proportional to the voltage, the higher the voltage rating of the Zener, the faster the discharge and hence the shorter the time taken. Therefore it is desirable to make this voltage as high as is compatible with the rest of the circuit (see Appendix B).

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 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 case of a gradiometer, the problem can be further ameliorated by

⁸ It is possible to calculate this time. A more complete derivation of the circuit behavior is given in Appendix B. There it is shown that the approximate method used here is good enough (for this case). In general, it will be good enough provided that the breakdown voltage is sufficiently high.

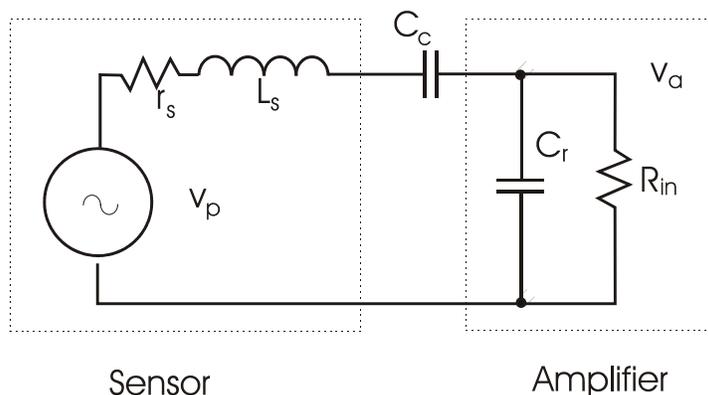
locating the relay roughly equidistant between the two sensors; the offset due to the iron then roughly cancels out. It goes without saying that the relay should be in the unenergized position during the measurement portion of the cycle. To properly protect the amplifier, the relay should be energized **before** the polarizing voltage is turned on in order to disconnect the amplifier from the sensor. It should then be turned off to reconnect the amplifier to the sensor **after** the polarizing current is turned off **and after** the high voltage dissipation part of the cycle is over. This is easily done by using a microprocessor to control the switching of the polarization current and the relay(s).

Chapter 4 Some Other Considerations

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 4 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 below.



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 as 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 about 10^5 Ohms (for the LM394, for example) while the resonating capacitor's reactance at the precession frequency is typically a between 10^3 and 10^4 Ohms.

With these approximations, it is possible to show that the voltage appearing at the amplifier's input, v_a , is related to the precession voltage, v_p , by the following equation:

$$v_a = \frac{\omega L}{r_s} v_p = Q v_p$$

Q is the ratio of inductive reactance of the sensor to the losses in the sensor and 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.

Relay-less polarization

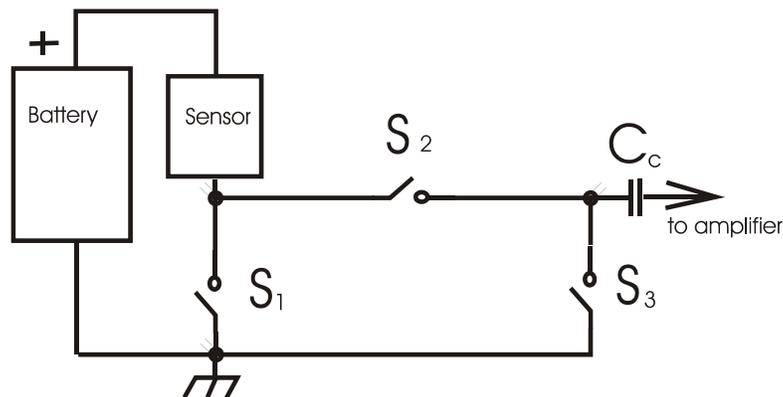
The presence of a relay in the instrument is always a problem. Relays have iron cores and so perturb the values of the local magnetic field. This is not a problem with a shipboard magnetometer towing the sensor at some distance behind the ship but it is a problem with hand-held magnetometers. However, 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 manufacturer's rated lifetime.

There are semiconductor devices which are called HEXFET's which can act as switches and so may be used to replace the relay. They are MOSFET's 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 following circuit diagram. I have shown the HEXFET's as if they were simple switches.

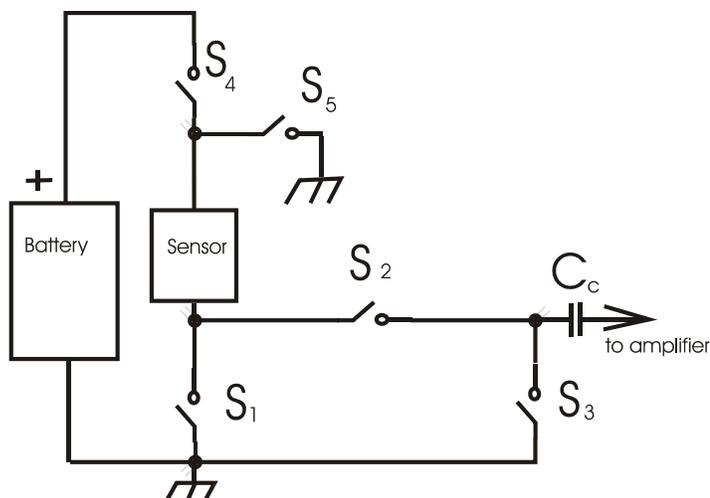
To polarize the sensor, switch S_1 is closed and, to isolate the sensor from the amplifier during the polarization, S_2 is opened and S_3 is closed.



To connect the sensor to the amplifier after polarization, S_1 is opened, S_2 is closed and S_3 is opened.

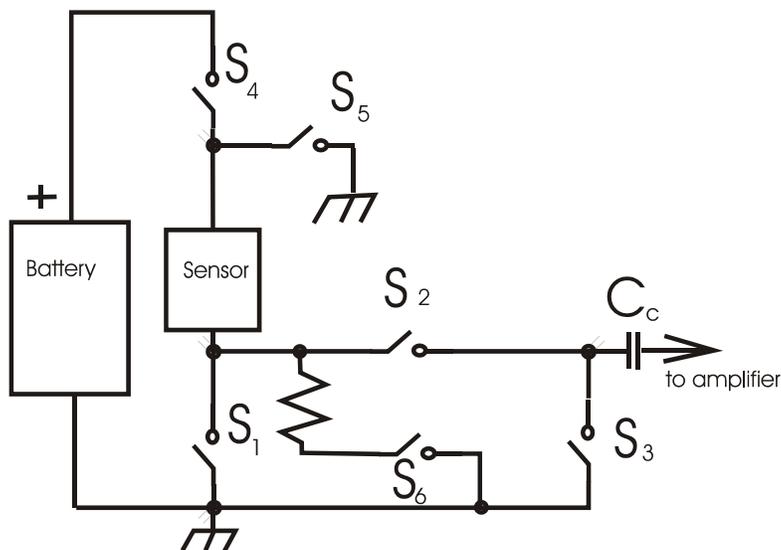
All three of the switches can be replaced by HEXFET's (specifically N-channel) and the opening and closing of the switches is done by applying logic-level voltages to their gates. Since the voltage at the junction of S_1 , S_2 and the sensor will go up to the breakdown voltage of S_1 (which might be $\sim 100V$) when it is opened, it is important that S_2 not break down at this voltage – it must have a high rating than S_1 .

This circuit has one problem and that is that S_1 , 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 the circuit below.



In this circuit, two additional HEXFET's have been added; S_4 and S_5 . During polarization, S_4 is closed and S_5 is open. After polarization, S_4 is open and S_5 is closed. This connects the top end of the sensor to ground and so no leakage current flows through the sensor. S_4 would be a P-channel HEXFET while S_5 would be another N-channel one.

As a final touch, immediately after polarization stops, because the sensor would otherwise 'ring' at the self-resonant frequency of the inductor, it is useful to connect yet another HEXFET switch to the junction of S_1 , S_2 and the sensor to ground through a resistance to act as a squelch of this transient. The final resulting circuit is shown below.



Here, S_6 is a N-channel HEXFET and the resistance is of the order of 10K.

The controller has a more complicated job in switching these HEXFET's but it is not difficult. A typical sequence might be as follows:

1. turn on S_3 and turn off S_2 and S_5 – this isolates the amplifier from the rest of the circuit,
2. wait a few mS,
3. turn on S_1 , S_4 and S_6 – this turns on the polarization current,
4. wait for the polarization time,
5. turn off S_4 and S_1 ,
6. wait about 10 mS to let the ringing die down,
7. turn on S_5 and turn off S_6 ,
8. wait a mS or two and then turn on S_2 and turn off S_3 ,
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 HEXFET's 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 HEXFET's. 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 don't want to switch the gates more rapidly than this time delay; which is the case for the sequence of operations stated above.

The photo below shows a complete microprocessor based PPM using HEXFET switching. This particular board is a first prototype; it is



very approximately 5cm by 8 cm in size and uses mainly surface-mount components. The brown connector on the left side connects to the sensor. The top connector on the right side connects to the battery and the lower one is the RS-232 interface to the outside world. The program allows one to set parameters like the polarization time, number of

cycles of the precession signal to count, etc. The PPM can be operated in either a demand mode where each measurement is initiated by an external command or in an automatic mode where the measurements are made one after another as rapidly as possible. The board was laid out in such a way as to keep the low-level amplifier circuits as far away as possible from the digital and polarization circuitry.

Digital 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 counted. The principal disadvantage of this method is that noise spikes due to ancillary circuits may cause errors in the count. An error of one count will give a very large error in the calculated magnetic field.

In a sense, this method throws away a lot of information. The precession signal is an exponentially decaying sine-wave plus noise and yet the method essentially just looks for and counts zero-crossings. All the amplitude information of the signal is not used.

It is possible to employ other mathematical techniques to determine precession frequency with the necessary precision. Basically, 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. 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 analyze 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 going to be some version of a ‘best fit’ algorithm. This technique 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:

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

It is an equation with five unknowns and so it can, in principle be solved by a least-square best-fit algorithm. Indeed, the situation is not even that bad because you have *a priori* knowledge of some of the parameters already. For example, C is going to be a 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. Calculating the average of the samples will give an approximate value for A.

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 $\frac{1}{B}e^{Ct}$ to give a simplified equation of form:

$$v' = \sin(Dt + E) + noise + average_error$$

where average_error is the DC offset error due to our crude method for determining A. Now we need only solve for D, E and the average_error; a best-fit analysis problem with just three variables. Best-fit algorithms require some initial estimate of the variable before starting. 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 is. There are a number of programs to calculate B at any location on the earth's surface. E is totally unknown and so is the average_error. Nevertheless, most modern lap-top computers are fast enough to do this sort of analysis 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 sample 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.

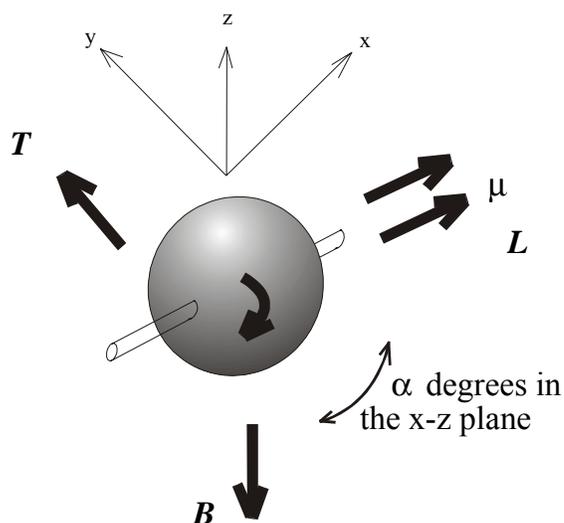
Finally, there are specialized microprocessors called DSP's (digital signal processors) which are specially designed for this type of analysis and they may be built into a complete, self-contained, instrument to measure precession frequency with the required accuracy.

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

Appendix A The Larmor Equation



Consider a proton with magnetic moment, μ and angular momentum, L , oriented in an external magnetic field as shown above. μ and L are in the x-z plane at an angle of α degrees from the $-z$ direction. The external magnetic field, B , is in the minus z-direction so the torque, T , resulting from the interaction between the proton magnetic moment and the external magnetic field is in the plus y-direction since:

$$T = \mu \times B$$

This is a vector cross product and all the quantities are vectors in this equation.

Newton's third law, for angular quantities, is:

$$T = \frac{dL}{dt}$$

In this case, we know that, because the torque is always perpendicular to the angular momentum vector, the additional change in angular momentum is always perpendicular to the existing angular momentum. The resultant is therefore constant in amplitude (and

equal to the original magnitude) but there is a constant change in direction; in other words, the proton will precess at a constant rate. In this geometry, the time-varying angular momentum vector will always lie in the xy-plane. Therefore, we can write the angular momentum vector in two dimensions as:

$$\mathbf{L} = L_{xy} \{ i \cos(\omega t) + j \sin(\omega t) \} + k L_z$$

where i , j and k are unit vectors in the x, y and z-directions respectively. L_{xy} (not bold faced) is the magnitude of \mathbf{L} in the x-y plane and L_z is the magnitude of \mathbf{L} in the z-direction – in this case, L_z happens to be negative but the following analysis is independent of whether it is positive or negative. ω is the angular rate of precession in radians per second.

Since $\boldsymbol{\mu} = \gamma_p \mathbf{L}$, we may write $\boldsymbol{\mu}$ as:

$$\boldsymbol{\mu} = \gamma_p L_{xy} \{ i \cos(\omega t) + j \sin(\omega t) \} + k \gamma_p L_z$$

and then, the torque, \mathbf{T} , becomes:

$$\mathbf{T} = \begin{vmatrix} i & j & k \\ \gamma_p L_{xy} \cos(\omega t) & \gamma_p L_{xy} \sin(\omega t) & \gamma_p L_z \\ 0 & 0 & -B \end{vmatrix}$$

which is:

$$\mathbf{T} = i \{ -\gamma_p L_{xy} B \sin(\omega t) \} + j \{ \gamma_p L_{xy} B \cos(\omega t) \}$$

Now, from the equation for \mathbf{L} , we can differentiate with respect to t to get ($\gamma_p L_z$ is constant and so its differential with respect to t is zero):

$$d\mathbf{L}/dt = i \{ -\omega L_{xy} \sin(\omega t) \} + j \{ \omega L_{xy} \cos(\omega t) \}$$

Therefore, since

$$\mathbf{T} = \frac{d\mathbf{L}}{dt}$$

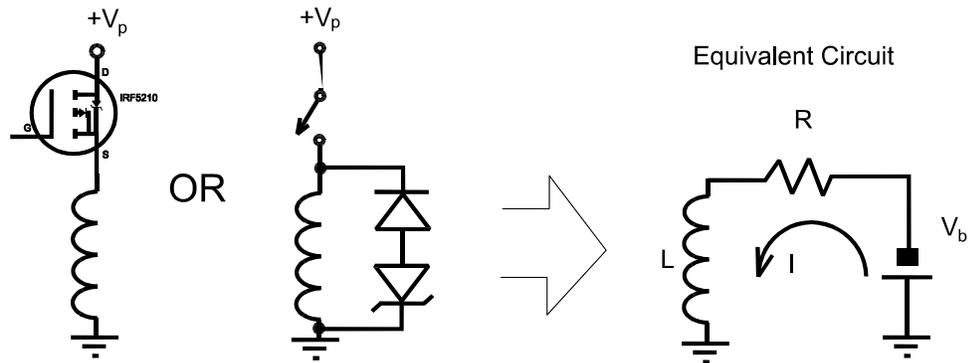
then

$$\omega = \gamma_p B$$

which is the Larmor equation. Note that the result is independent of the orientation of \mathbf{L} with respect to \mathbf{B} provided that L_{xy} is not zero. If L_{xy} is zero (i.e., if \mathbf{L} is parallel to \mathbf{B}), then there is no torque and hence, no precession.

Appendix B The Polarization Cycle Decay

Consider the figure shown below. Here are two possible circuits for polarizing the sensor and the equivalent circuit **after** the HexFET is off or the switch **has been opened** (respectively).



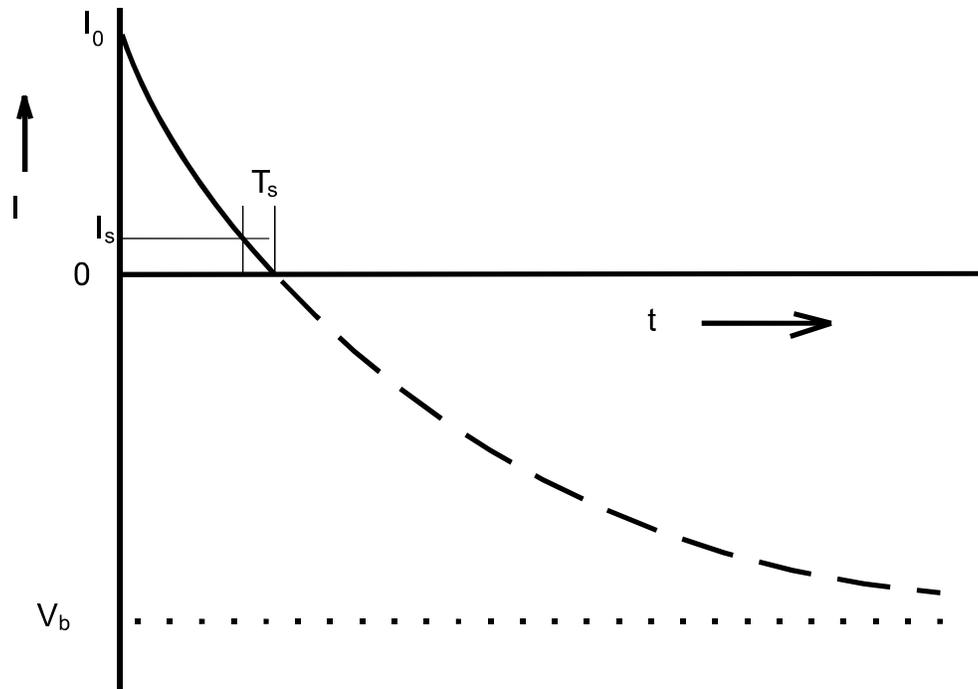
In this figure, V_b is either the breakdown voltage of the HexFET or the voltage of the Zener diode. The equation for I in the equivalent circuit is:

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

where R is the dc resistance of the sensor, and L is its inductance. Given the initial condition that the current is I_0 at $t=0$, just prior to the instant that the switch is opened (or the HexFET is turned off), the solution for I is then:

$$I(t) = \left(I_0 + \frac{V_b}{R} \right) e^{-\frac{R}{L}t} - \frac{V_b}{R}$$

The graph of this equation is shown below:



Here, T_s is the total time for the current to go from I_s to 0. Using the same notation as in Chapter 3, the time at which I goes to zero is t_s . The equation may be solved to find this time (at which I goes to zero) giving:

$$t_s = \frac{L}{R} \ln \left(1 + \frac{I_0 R}{V_b} \right)$$

In the example given in Chapter 3, I_0 was 1.6 A, R was 7.4 Ohms, V_b was 100 V and L was 20 mH. Substituting those numbers into this equation gives a calculated value for t_s of 0.3 mS. The value observed on an oscilloscope was 0.2 mS; in good agreement with the calculated value.

In order to calculate T_s , we need to find the slope of the equation for I as a function of time. This is given by:

$$\frac{dI}{dt} = -\frac{R}{L} \left(I_0 + \frac{V_b}{R} \right) e^{-\frac{R}{L}t}$$

and, using the value for t_s which was obtained in the previous example, we find that dI/dt in the vicinity of where the current goes to zero is about 5,000 A/s. Using this value, T_s turns out to be $\sim 4 \mu\text{S}$ instead of the previously estimated $\sim 2.5 \mu\text{S}$ (in Chapter 3). Obviously, the much simpler analysis in Chapter 3 is sufficiently good **for this sensor and polarization circuit**; we do not need to do the more extensive analysis derived here.

These equations are useful however, for those cases where V_b is very small. For example, some published circuits for PPM's just show a reversed diode across the sensor (sometimes, a few in series). That is, in the second circuit at the beginning of this Appendix, the Zener diode is replaced with a short circuit. In this case, V_b might be as low as just one diode voltage drop ($\sim 1 \text{ V}$ at the typical currents of a ~ 1 to $\sim 10 \text{ A}$) and, t_s may be several mS. Therefore T_s is also very long – perhaps too long!

You may use these equations to do the appropriate calculations for the circuit you use.

Appendix C Relative Sensitivity of some Common Fluids

Substance	Molecular Weight	Number of H atoms per molecule	Density (kg/cubic metre)	Chi relative to water
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

Appendix D Relative Accuracy versus Achieved Signal-to-Noise Ratio

In the table below, it is assumed that the measurement of the frequency takes place over $\frac{1}{2}$ second and that the frequency is close to 2130 Hz (50,000 nT).

S/N	Relative Accuracy	Error in 50,000 nT (nT)
1	9.2E-5	4.6
2	5.1E-5	2.5
5	2.1E-5	1.1
10	1.05E-5	0.5
20	5.28E-6	0.26
50	2.11E-6	0.11
100	1.06E-6	0.05
200	5.28E-7	0.03

Note: In a digital counter, you can expect to get much larger errors for low S/N ratios because noise may cause excess transitions which will be counted as complete cycles.

Bibliography

- 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
- 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
- Burrows, K., *A Rocket Borne Magnetometer*, J. Brit. IRE, **19**, 769-776, 1959
- Cahill, L.J. and Van Allen, J.A., *High Altitude Measurements of the Earth's Magnetic Field with a Proton Precession Magnetometer*, J. Geophys. Res., **61**, 547-558, 1956
- Faini, G., Fuortes, A. and Svelto, O., *A Nuclear Magnetometer*, Energia Nucleare, **7**, 705-716, 1960
- 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
- Haggard, L., *Proton Precession Ferrous Metal Detector*, Practical Electronics, 782-793, October, 1970
- Packard, M.E. and Varian, R.H., *Free Nuclear Induction in the Earth's Magnetic Field*, (abstract only), Phys. Rev., **93**, 941, 1954
- Primdahl, F., *Scalar Magnetometers for Space Applications*, Measurement Techniques in Space Plasmas: Fields, Geophys. Monograph **103**, American Geophysical Union, 1998
- Serson, P.H., *A Simple Proton Precession Magnetometer*, Dominion Observatory, Ottawa, Canada, 1962
- Serson, P.H., U.S. Patent No. 3,070,745, 1962
- Sigurgeirsson, Th., *A Continuously Operating Proton Precession Magnetometer for Geomagnetic Measurements*, Scientia Icelandica, **2**, 64-77, 1970
- Stuart, W.F., *Earth's Field Magnetometry*, Rep. Prog. Phys., **35**, 803, 1972
- Waters, G.S. and Phillips, G., *A New Method of Measuring the Earth's Magnetic Field*, Geophys. Prospecting, **4**, 1-9, 1956
- Waters, G.S. and Francis, P.D., *A Nuclear Magnetometer*, J. Sci. Instrum., **35**, 88-93, 1958

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