FGM-series
Magnetic Field Sensors

SCL

- +5 volt operation
- Three Terminal Devices
- DC to 20 KHz Bandwidth
- Low Temperature Sensitivity
- High Intrinsic Sensitivity
- Built-in feedback coils

Description

The FGM-X series of devices are very high sensitivity magnetic field sensors operating in the ±50 microtesla range (±0.5 oersted). They are simple, essentially three terminal devices, operating from a single +5 volt supply, the connections being ground, +5v and output. The output is a robust 5 volt rectangular pulse whose period is directly proportional to the field strength, (giving a frequency which varies inversely with the field), making it very easy to interface to a computer or micro controller. The typical period swing for the full range of an FGM-3 is from 8.5 μs to 25 μs (~120 KHz to ~50KHz), a clear indication of its high sensitivity.

The FGM-1 is essentially a miniature version of the FGM-3 with a slightly lower sensitivity.

The FGM-2 is a single package containing two FGM-1 sensors at right angles to one another.

Unlike Hall Effect field sensors, which are virtually unusable in this range because of their large temperature sensitivity, the FGM-series has a very low temperature coefficient.

Since the lowest effective Nyquist sampling rate is ~50 KHz, appropriate filtering can provide an AC field bandwidth from DC to ~20kHz.

Maximum Supply Voltage	**7 volts**
Recommended supply Voltage	**5 volts ±0.5 volts**
Typical Supply Current	**12 mA**
Operating Temperature Range	**0 - 50 °C**

Typical Applications

Since the range covers the earth's field magnitude, multiple sensors can easily be arranged to provide compass orientation or full three-dimensional orientation systems, using the local earth's field as a reference, (-gimballed compass or virtual reality helmet devices.)

Other applications include conventional magnetometry, earth field magnetometry, ferrous metal detectors, internal vehicle re-orientation alarm sensors, external vehicle or ship passage sensors, wreck-finders, non-contact current sensing or measurement, conveyor belt sensors or counters and in conjunction with small permanent magnets, movement and proximity sensors and ferrous impurity detectors for non-magnetic alloys.
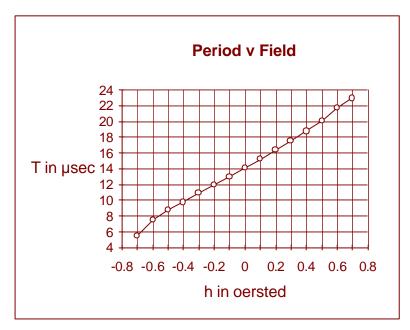
Conventional magnetometry would include magnetic material measurement, B-H loop measurement, Preisach displays and archaeological artifact assessment.

For use with applications requiring a larger range of field strength, external feedback winding techniques can increase both the sensor's maximum range and its linearity. This approach is described in more detail in the separately available "Application Notes" but basically consists of using an overwound solenoidal coil in a negative feedback loop which continuously attempts to zero the field seen by the sensor. By this means the range and the linearity cease to be a function of the sensor characteristics and depend only on the feedback current through the coil.

For added convenience, the miniature versions, FGM-1 and FGM-2, have such a coil built into their normal structure.
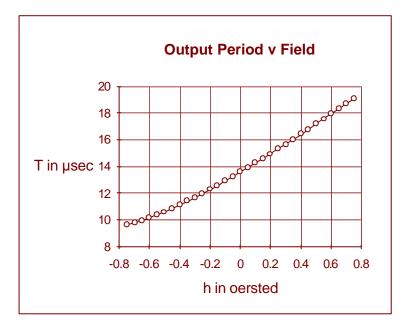
Output Period as a Function of Field Strength

Typical Characteristics

The chart below shows the typical response of the larger size sensor of the range, the FGM-3

Between ±0.5 oersted (±50 µtesla) the non-linearity is about 5.5%.

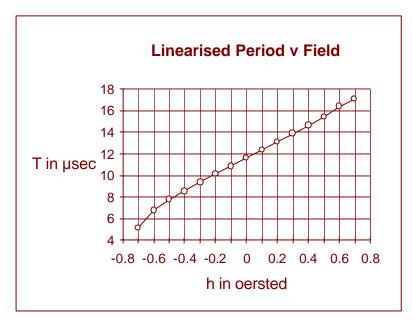A typical response for the miniature types of sensor, the FGM-1 and FGM-2 is given below.

Output Period v Field

T in µsec (vertical axis, values 8 to 20)

h in oersted (horizontal axis, values -0.8 to 0.8)

Between ±0.5 oersted (±50 µtesla) the non-linearity is about 4.8%

This non-linearity varies somewhat between individual sensors, but may normally be expected to be in the region of 5%.

A simple strategy will improve this considerably.

It was stated earlier that the field strength was inversely proportional to the frequency. In practice it will be found that the field strength is more closely inversely proportional to the frequency plus a small constant. If the frequency is measured and a fixed number of kilohertz is added before it is divided into one, to obtain the period, the response curve of period against field will be seen to become much more linear.
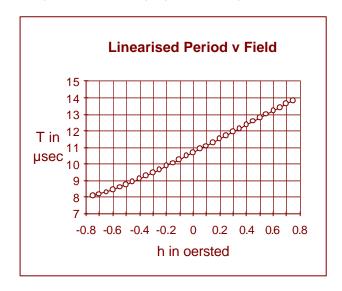
The chart below shows the effect of adding a fixed value of 15 KHz to the incoming frequency of an FGM-3 sensor before inverting to obtain the period. (Plot shows T= 1/(f+c), c=15 KHz)

Linearised Period v Field

T in µsec (vertical axis, values 4 to 18)

h in oersted (horizontal axis, values -0.8 to 0.8)

3

Apart from the droop at -0.7 oersted the linearity has improved considerably. In fact applying the same definition as before the non-linearity between ±0.5 oersted has been reduced from 5.5% to 0.7%
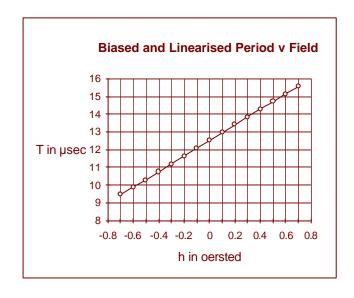
In individual cases varying the 15 KHz figure may produce an even better non-linearity, but in general this value will produce a significant improvement with any production sensor, where individual tailoring is undesirable.

The same technique can be used with the miniature sensors FGM-1 and FGM-2. The chart below shows the effect of adding a constant frequency of 18 KHz to the incoming frequency before obtaining the period by inversion. (Plot shows T=1/(f+c), c=18 KHz)

Linearised Period v Field

T in µsec

h in oersted

For the miniature sensors a further improvement in linearity and also maximum range can be obtained by making use of the built in overwound coil intended for feedback systems. This can also be used as a simple biasing coil to alter the zero field period of a sensor. Inspection of the graph above suggests that moving the h scale to the right might improve the linearity.

This is, in fact, the case and the effect of injecting a current of about minus 1 mA (a 4K7 resistor to a -5 volt supply) into the feedback winding pin, can be seen in the graph below.

Biased and Linearised Period v Field

T in µsec

h in oersted

4

The linear range and linearity can be seen to have improved considerably and applying the same criterion as previously, the non-linearity has been reduced from 4.8% to 0.2% across the increased range of ±0.7 oersted.

If the application is set up in such a way as to measure the period directly, rather than the frequency, followed by an inversion, the linearising technique can still be used.

The method is to use $T/(T+cT)$ as the quantity which is proportional to field strength, where T is in microseconds and c is in kilohertz as before.

For applications requiring greater linearity, the more complex overwound external coil feedback system can provide this together with even more enhanced field strength range if needed.

Supply Voltage Variation

The period (and frequency) of the FGM-series of devices varies with supply voltage, having a coefficient of about 3.5% per volt at the nominal 5 volt supply level. For precise applications good supply regulation is required, but since the transducer's current requirement is low, this is fairly easy to achieve using, for example, single or double regulation with devices from the LM78LXX series.

General Application Notes

Use with Computers and Microcontrollers

The large pulse output gives considerable noise immunity permitting the use of transducers sited at long distances from the main system.

Interfacing is simple in that it requires only one bit of a digital input port per channel of measurement, the technique being to count input pulses for a fixed period to determine the frequency of the incoming signal, from which the field can be calculated. Alternatively, where a faster response is required, the time between successive like edges permits the direct determination of period, from which again the field can be calculated.

With microcontrollers this usually presents no problem, but with systems using many interrupts or extensive multi-tasking it may be necessary to buffer the input signals to deal with the high data rate. However this usually means no more than the addition of a single triple-counter I/O chip even for three-dimensional orientation systems.

For applications such as earth field magnetometry, where readings may only be required at relatively long intervals simple binary division with a single chip 12 or 14 stage divider will reduce the input period to a level where data rate ceases to be a problem to the computer. Alternatively, in such applications where the field variation is extremely small, digital heterodyning with a stable oscillator will also reduce the period but simultaneously maintain the high sensitivity, (in hertz/oersted) to field variations.

For applications which need absolute field magnitude without any orientation sensitivity, it is necessary to use three orthogonal sensors and exploit the fact that the sum of the squares of the three signals is constant regardless of orientation. Provided that the zero offsets, channel sensitivities and linearisation are appropriate to the required absolute sensitivity, this will permit free movement of the sensor head while measuring small changes in absolute field. If the sensor is in constant angular motion, advantage can be taken of this to provide some level of auto-calibration of zero offset and channel sensitivity. (See Application Note - Auto Calibration Algorithm .)
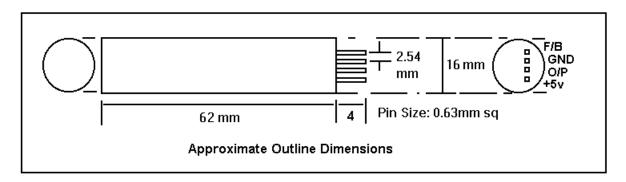
Where the sensor can be permanently fixed, only one sensor is necessary, the zero offset being adjusted to suit the local ambient field strength. This technique is appropriate to fixed ferrous metal detection systems such as conveyor belt counters, vehicle and ship passage detectors and materials magnetometry. A limit to the range of such systems results from the fact that the earth's field itself fluctuates at a low level continuously. The effective range will be a function of the size or likely magnetic moment of the objects being detected, ships generally giving a larger range than vehicles or hand guns. Appropriate filtering of the input frequency variations will enhance range.

Where extremely high sensitivity is required it may be possible to use two sensors in a gradiometer configuration to cancel out the micro-fluctuations of the earth's field. However, this will not always increase range, since the gradiometer sensitivity falls off faster with range than the simple field sensor.

In this context, it should be remembered that the field produced at range by a magnetic moment falls off as the inverse **cube** of the range, so the gradiometer configuration will fall off as the inverse **fourth** power. However such systems may be useful as short range high sensitivity detectors and materials measurement systems. An example might be extremely small magnetic moment inert particles introduced into fluid flow systems for movement detection, such as chemical processing plants or animal internal fluid flow systems in medical research applications.
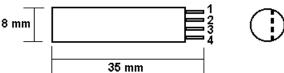
PHYSICAL CHARACTERISTICS

Sensor Outline - FGM-3

Approximate Outline Dimensions

The sensor has been made with a cylindrical form in order to facilitate the overwinding of external feedback coils. Alternatively, it can be readily inserted into a separately fabricated coil on a tubular core.

As a simple guide, for example, a single layer of 0.2mm wire (0.25mm overall say) wound over 60mm of the sensor length will give the equivalent of 4000 turns/metre. Since 1 oersted is approximately 80 ampere-turns/metre, such a winding will produce a field of around 50 oersted/ampere. Thus with up to 100mA flowing it is possible to offset fields of ±5 oersted, increasing the range of the sensor by x10. Negative feedback also brings all the usual benefits of improved linearity and stability, of course.
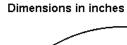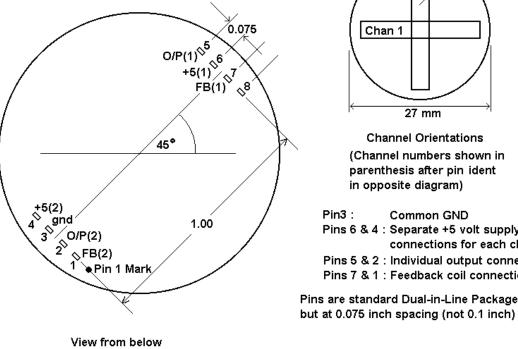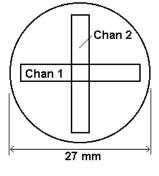
FGM-1 Pin Layout

8 mm

35 mm

Pin 1: feedback
Pin 2: output
Pin 3: gnd
Pin 4: +5 volts

Pins are standard DIL package
types but at 0.075 inch spacing

FGM-2 Pin Layout

Dimensions in inches

0.075

O/P(1) 5
6
+5(1) 7
FB(1) 8

45°

+5(2)
4 gnd
3 O/P(2)
2 FB(2)
1 Pin 1 Mark

1.00

View from below

Chan 2

Chan 1

27 mm

Channel Orientations
(Channel numbers shown in
parenthesis after pin ident
in opposite diagram)

Pin3 :   Common GND
Pins 6 & 4 : Separate +5 volt supply
                 connections for each channel
Pins 5 & 2 : Individual output connections
Pins 7 & 1 : Feedback coil connections

Pins are standard Dual-in-Line Package Types
but at 0.075 inch spacing (not 0.1 inch)

CONVERSION TABLES
Magnetic Flux Density

	gauss	tesla	gamma
1 gauss	1	10^{-4}	10^{5}
1 tesla	10^{4}	1	10^{9}
1 gamma	10^{-5}	10^{-9}	1

Magnetic Field Strength

	amp/metre	oersted
1 amp/metre	1	0.01257
1 oersted	79.58	1

NOTE: Technically, the sensor measures flux density, in gauss, but since in vacuum (and virtually in air) the units of flux density are the same magnitude as those of field strength and since the sensor can only really be used in air, oersted have been used in the text and diagrams as equivalent to gauss.

FGM-series
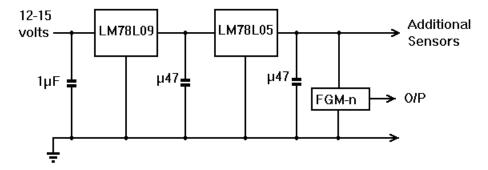**Magnetic Field Sensors**

SCL

Application Notes

General Application Techniques

Decoupling and Power Supply Regulation

If long leads are used with the sensor it is advisable to provide some local decoupling close to the sensor itself where possible. A 10μF electrolytic capacitor is suitable for this purpose.

Since the sensor itself has a sensitivity of a few percent to power supply variations, it is necessary to provide it with some power regulation in most cases. For many applications, such as orientation devices, a single fixed voltage regulator of the LM78L05 (or equivalent) type is adequate. However for applications such as earth field magnetometry or where extremely small field variations are being studied, supply voltage variation needs to be reduced to a level which permits the temperature coefficient of the device to be the limiting performance factor.

Double regulation from 12-15 volts, first down to 9 volts and then to 5 volts, using the LM78L09 and LM78L05 provides a low cost solution, a typical arrangement being shown below.

Double Regulation Supply

Sensor Calibration

For many applications, such as simple field detection or orientation measurement systems, calibration of the sensors is not necessary.

For those applications which do need to measure field strength, a reasonably accurate calibration can be made using simple equipment. A single layer solenoid can easily be made by close-winding enamelled wire on to a tube having an internal diameter large enough for the sensor to be inserted in it. The field inside a long solenoid is given simply by the product of the current flowing in it and the number of turns per metre with which it is wound. Both these items can be measured reasonably accurately, one with a ruler, the other with an ameter.

For most purposes, a winding at least twice as long as the sensor will give a good calibration, consistent with the likely turns/metre measurement accuracy using a ruler, provided its diameter is no greater than necessary.

1

Single axis sensors are the easiest in this respect because they have a small diameter themselves and can be inserted into a small diameter tube. The following tables should be helpful in the the design of calibration coils.

First, the field at the centre of a cylindrical coil of the type suggested is given by:

H=geometry-factor x number of turns per meter x current   (H in amperes/metre)

where the geometry-factor is shown in the following table.

Length/diameter	Geometry-factor
5	0.9806
6	0.9864
7	0.9900
8	0.9923
9	0.9939
10	0.9950

This permits calibration of the coil centre-field. The correction is small, somewhere between 0.5 and 2 percent, but may be worthwhile in appropriate cases.

Away from the centre of the coil the field falls off towards either end and on the assumption that the coil is twice as long as the sensor, the following table gives a factor for this reduction, at either end of the sensor, for various coil geometries. It also shows, in the third column, the percentage by which the field differs from being uniform along the length of the sensor, assuming that the sensor is centrally placed.

Length/diameter	Reduction-factor	Uniformity
5	0.9788	±1.06%
6	0.9847	±0.77%
7	0.9884	±0.58%
8	0.9910	±0.45%
9	0.9938	±0.31%
10	0.9942	±0.29%

The geometry-factor from the first table should be reduced by this percentage to arrive at a mean calibration factor for the coil, for the most accurate results.

The calibration currents required are modest. Since the usable range of the sensor is around ±0.5 oersted or 40 amperes/metre a single layer winding of 0.5mm enamelled wire (which has an overall diameter of ~0.559mm) will only require 23 mA to reach maximum calibration field strength.

In carrying out such a calibration with a solenoid coil, the coil and sensor should be aligned at right angles to the direction of maximum local field as determined by the sensor alone. Where only relative field measurements are needed this can be done by simply aligning the coil and sensor in an east-west direction. If an accurate zero field calibration point is required the sensor will need to be placed in a zero-field location, such as the inside of a small mumetal container, aligned east-west.

Linearity Correction

In practice, the linearity is much better than the data specification suggests (over the recommended operating range) and few applications really require any correction at all. In a plot of field against period a small concave downwards curvature may be detected which is most simply straightened by adding a small proportion of the period squared to the assumed linear relationship. The proportion to add can be estimated by comparing the full scale negative and positive field periods with the zero field period, obtained from a calibration exercise as described earlier. Substituting these three values for T in the data sheet suggested equation, $H=c_0+c_1(T-T_0)+c_2(T-T_{min})^2$, gives three simultaneous equations from which c_0, c_1, and c_2 can be determined.

2

If the periods are normalised by dividing by the zero field period value, $T_0=1$ and the three equations are;

$$-0.5=c_0+c_1(T_{min}-1) \qquad T_{min} \text{ measured at -0.5 oersted}$$

$$0=c_0+c_2(1-T_{min})^2 \qquad T_0=1$$

$$0.5=c_0+c_1(T_{max}-1)+c_2(T_{max}-T_{min})^2 \qquad T_{max} \text{ measured at +0.5 oersted}$$

Physically interpreting these coefficients, c_1 is just the usual slope of a linear relationship, c_2 is the small period squared correction previously discussed and c_0 is to offset the disturbance of the zero field period by the squared term and pull T_0 back to to its normalised value of one.
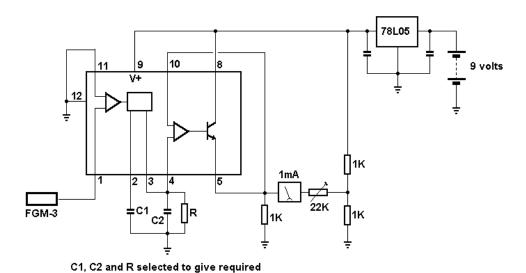
Field Measurement Methods

The simplest method of making field measurements is to use a frequency meter directly on the output of the sensor, set to period measuring mode. In most cases application designers will prefer to dedicate some specific hardware to carry out the conversions automatically and present the results in some acceptable form rather than use the eyeball/pencil and paper method, though this is quite adequate for calibration and familiarisation exercises.

Hardware configurations can vary from minimal, battery powered meter-display detectors through to complex, multiple sensor, computer controlled data collection systems. The following notes describe some useful techniques which may be incorporated into such designs without detailing any complete systems.

Meter or Chart Recorder Outputs

Low cost equipment can be made by using a semiconductor frequency-to-voltage converter such as the LM2907 or equivalent. Many variations of this type of integrated circuit are available in the component catalogues. A suggested circuit for a portable, direct reading instrument is shown below.

C1, C2 and R selected to give required
range (see LM2917 data sheet)

Simple Portable Meter Output Instrument

Since the field strength is inversely proportional to frequency the output is not linear but over the ±0.5 oersted range the non-linearity is modest and provides an acceptably spaced meter scale. At higher sensitivities over a more limited full scale range, the non linearity becomes increasingly neglegible and gives an almost evenly spread scale.
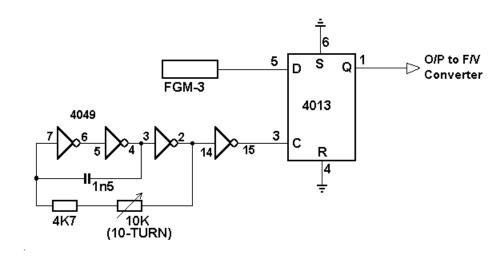
Where the converter device is unable to handle the highest frequency output, simple binary division in a prescaler chip is an appropriate method of bringing the sensor output into an acceptable range.

As an alternative to the diode pump type of converter a phase locked loop can also be configured to provide similar performance.

Digital Heterodyning

Where the full range of the sensor is not required, such as in ferrous material detection or the measurement of small field fluctuations, a technique of digital heterodyning is useful. This is not a true heterodyne process but has close similarities when used over small frequency ranges. It is more akin to the production of aliasses by undersampling and gives very high sensitivity to small signal fluctuations. This makes it very for the remote detection of moving ferrous objects or for the measurement of the earth field fluctuations which accompany magnetic storm activity.

The technique requires a stable but adjustable source of clock pulses of similar frequency to the sensor output. These are used to undersample the sensor output and produce a much lower frequency square wave. One easy way to achieve this is to use the sensor output as the D-input of a D-type bistable and the clock source as the trigger input as shown below.

Digital Heterodyne

The sensor is used in a fixed position and the clock signal is set to a frequency close to the sensor frequency. The output of the bistable is a square wave of frequency equal to the difference between sensor and clock frequencies, similar to a heterodyne mixer.

A small percentage change in the sensor frequency becomes a large percentage change in the bistable frequency. This can be converted to a voltage as before, for meter or chart recorder, but gives a large increase in apparent sensitivity without the need for high gain amplifiers.

The configuration of the CMOS oscillator above is more stable than most and is suitable for such applications as ferrous metal detectors. Some care is needed in decoupling both the oscillator and the sensor to prevent a tendency to frequency lock, if the highest sensitivity is required. However such an arrangement has successfully detected passing vehicles and a measure of its sensitivity can be obtained from the fact that it could pick up a motorcycle in the far lane of a three lane motorway from the grass verge.

The more critical applications of earth field or materials magnetometry need a more stable oscillator, such as a crystal controlled type, but a fixed frequency is normally adequate, since such instruments are not usually mobile.
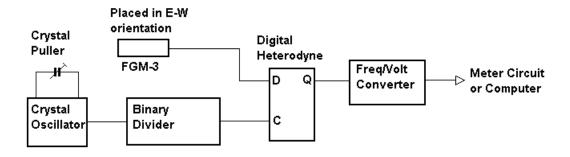
The technique is equally applicable to ferrous object detection or counting of smaller objects passing on a conveyor belt, the sensitivity and range being adjusted to suit the individual system. In this context it is useful to remember that the field produced by a given magnetic moment falls off as the inverse *cube* of the distance, not the inverse square.

All of the above hardware approaches can equally well, if not better, be simulated by software in a computer or microcontroller, often resulting in minimal hardware to achieve sophisticated results.

Earth Field Magnetometer

A block diagram of a modest earth field magnetometer is shown below using the type of circuitry described above.

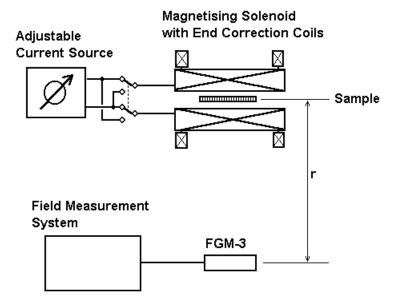## Earth Field Magnetometer Configuration

The sensor should be located in an east-west orientation and its mean frequency is measured. A crystal oscillator and binary divider is then selected to produce a frequency around 500 Hz below the sensor frequency. The sensor signal and the divided clock are fed to a digital heterodyne circuit as described earlier. The sensitivity of the FGM-3 is such that this arrangement will give a swing of about 0 to 1000 Hz for a variation in field of around ±500 gamma. (1 gamma $=10^{-5}$ oersted) This gives enough headroom for most magnetic storms likely to be observed. The exact range can, if required be calibrated as outlined earlier.

The output of the digital heterodyne can be taken to a voltage-to-frequency converter for chart recorder use or to a computer to store or plot the data in whatever form is appropriate. If the sensor is calibrated the results can easily be converted to angular or azimuth variations by dividing by the local horizontal component, (which can be measured by a north-south oriented calibrated sensor.) This gives the variation in radians, readily converted to the more suitable minutes of arc.

A magnetometer of this kind needs to be installed in a location far removed from potential sources of magnetic field interference, such as mains transformers and motor vehicles. Fortunately the inverse cube law mentioned previously helps considerably with this aspect. One exception to this rule, however can be exploited in the initial commissioning of the equipment. If difficulty is experienced in finding an appropriate combination of crystal and binary divider to nearly match the sensor output, in the magnetically quiet location needed, the strategic placement of a small ceramic magnet, at a suitable range, can be used to "pull" the sensor frequency instead.

Materials Magnetometer

A setup with this kind of sensitivity is equally capable of being used as the measurement tool in a classical Gauss-type materials magnetometer. The usual configuration of this instrument is illustrated below.

Materials Magnetometer Configuration

This consists of a controllable magnetising arrangement in the form of an air-cored solenoid as shown in the case of low susceptibility specimens or an electromagnet type yoke to reduce demagnetising effects with high susceptibility materials, in conjunction with a field measuring device. The demagnetising field measured by this device is related to the magnetic induction in the specimen.

Taking the magnetising coil through positive and negative cycles large enough to reach saturation in the specimen produces a measured field which displays the hysteresis loop characteristics of the sample specimen. To avoid shearing the hysteresis loop it is necessary to use a specimen with a small demagnetising coefficient or one with a known demagnetisation coefficient which can be corrected for, such as a cylinder. For straightforward quantitative work the distance r should be large compared to the magnetic length of the specimen if induction is to be measured, requiring high sensitivity in the detector. Comparative work and coercive force measurement can be carried out without satisfying this criterion.

Calibration is often carried out using a standard comparison sample of known demagnetisation coefficient and magnetic properties. The arrangement shown can also deal with the process of anhysteretic magnetisation if an alternating current source is superimposed on the direct current supply. The more common ac-driven type of B-H loop tester has difficulty with this since it cannot measure static fields.
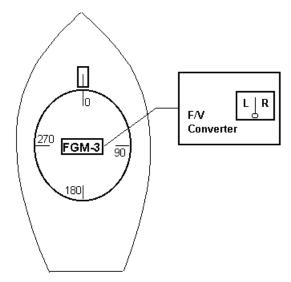
Similar arrangements are used by naval establishments under the name of fixed or portable ranges to determine the efficiency of the various degaussing equipments used to maintain the "magnetic hygiene" of vessels and items taken on board. The other side of this coin is that even more sensitive systems attempt to locate such vessels by detecting the magnetic anomaly caused by their presence.

The latter type of detector, however, needs to be insensitive to its orientation in order to avoid avoid the effects of the earth's field when it is in motion. Provided they can be made orientation insensitive, such detectors can be used as remote wreck-finders by divers

Orientation Measurement Devices

The most well known of these is the common compass, which is used to indicate the direction of the local horizontal component of the earth's field and from this to deduce the heading of the vehicle or vessel in which it is installed.

The simplest is a single axis device known as a steer-on-heading compass or "poor man's autopilot". It consists of a single magnetic sensor mounted on a rotatable disk, marked in degrees around the periphery and fitted with a stationary indicating pointer, as drawn below.

Steer-on-Heading Compass

The output of the sensor is connected to a frequency to voltage converter circuit feeding a centre zero meter display as described previously. In the diagram shown, if the boat veers to the left the meter needle swings to the right, indicating the need to steer in that direction to correct the course. If the rotating disk is turned to a new heading the needle will show the shortest direction to steer in until the new heading is reached when it will return to the centre position. The size of the deflection gives an indication of the amount of correction needed at any time. This type of steering system is said to be easier on the helmsman than having to remember and follow a degree bearing.

The implications of tilt in the sensor will be discussed in detail later but, for this simple system to be practical, the sensor must be gimballed and weighted so as to keep its axis level at all times. Since it is only a single axis device it only needs a single gimbal, provided that the gimbal rotates with the heading disk.

The next level of complexity is a two axis compass and for this it is best to replace the frequency to voltage converter with a microcontroller of some sort as a number of more complex operations need to be carried out especially if a readout display is wanted. Many varieties exist all capable of dealing with the requirements of a compass, but because the sensors have their own analogue-to-digital feature, microcontrollers which have frequency or period determining features built in are the obvious choice in this instance.
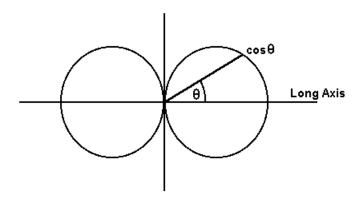
Microcontroller selection has been dealt with better elsewhere and these notes will restrict the subject matter to the compass design principles using the FGM type sensor.

Angular Sensitivity

It is useful to look first at the angular response of an individual sensor. Because of its structure it "sees" the full magnitude of a field which is aligned along its long axis. For any field at right angles to this axis it gives zero output in the sense that its period corresponds to that of a zero field condition. For a field aligned at an angle between these two extremes the response is proportional to the projection of the field on to the long axis of the sensor, therefore to the cosine of the angle between field and sensor.

This gives rise to the classic figure-of-eight polar diagram, consisting of two contacting circles or, in the three dimensional case, two contacting spheres.

Polar Response of Sensor

If sensors are aligned along the axes of any two or three axis coordinate system the sensor outputs represent the direction cosines of the field vector with respect to that coordinate system. For convenience the chosen system is usually cartesian but this is not a requirement.

As described under the calibration and linearising techniques it is convenient to normalise the sensor readings by dividing through by the zero-field period. In orientation type devices it is also convenient to then subtract one from these normalised values to yield equal positive and negative ranges about zero. These adjusted values are then proportional, but not yet equal, to the direction cosines of the field vector.

The reason is that no two sensors are exactly alike in absolute sensitivity and must now be calibrated so as to achieve a standard sensitivity. This can be done by the calibration coil method described earlier, after which proportionality constants can be assigned as multipliers to equalise the sensitivities. Alternatively it can be done by aligning the individual sensors in turn along the local earth field vector in the two possible directions, 180° apart and determining the corresponding maxima and minima for each sensor. Proportionality constants are again assigned to equalise the sensitivities.

Two Axis Orientation Sensing

The two axis compass uses twin sensors superimposed at right angles to one another in the same location and both constrained to lie in the horizontal plane. The sensitivity equalising process in this case can be semi automated by rotating the sensors through a full 360° and allowing the software to determine the maxima and minima for both axes.

If the two now standardised values are regarded as the x and y components of the local field vector, $\mathbf{h}$, having a modulus equal to $\sqrt{(x^2+y^2)}$, the final normalisation can be realised by dividing each component by this modulus. This gives the true direction cosines of the field vector which together define the unit vector $\mathbf{i}$, having the same direction as the field vector, $\mathbf{h}$.

This process eliminates the effect of any variation of the absolute magnitude of the measured field, since the sum of the squares of the direction cosines always equals one. Earth field variations are insignificant in this context, but supply or ambient temperature changes are neutralised provided all sensors are equally effected.
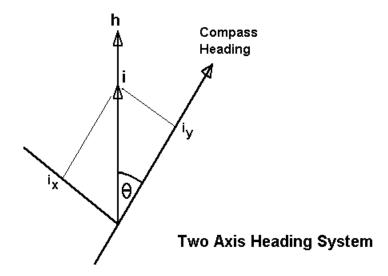
The direction cosines can be readily converted to a more customary representation such as angular heading as follows.

Assume that the compass heading indication is aligned with the y-axis and label the components of the unit vector, i_x and i_y. Then it can be seen from the following diagram that if θ is the conventional heading angle,

$$\tan\theta = -i_x/i_y \qquad \text{and} \qquad \theta = \tan^{-1}(-i_x/i_y)$$

and the compass heading is simply the arctangent of the ratio of the x and y components of the unit vector in the earth's field direction.

Two Axis Heading System

For a three dimensional coordinate axis system with the z-axis at right angles to the other two, there is no conflict with anything that has been said so far, provided that the z-axis remains vertical. In fact this becomes the necessary condition for the successful operation of this type of compass, which needs to be gimballed in two directions and appropriately weighted.

It will be evident that some attention to signs and the possible divisions by zero will be required in considering the full circle of 360°. While this arctangent solution may be possible for a computer with trigonometric functions in a high level language, it is not really appropriate for a lower level of implementation such as a microcontroller, though the underlying principle remains the same in alternative approaches.

The full circle in which the heading vector lies may be segmented into eight 45° octants and the octant occupied by the field vector can be identified by simple non trigonometric tests, easily applied in software.

The rules which do this involve the signs of the i_x and i_y components and the comparative magnitudes of these components taken as an ordered set. For example if $i_x<0$ and $i_y>0$ the heading must lie in the first quadrant. If, in addition, $|i_x| < |i_y|$ it must lie in the first octant between 0° and 45° as in the previous diagram. Other combinations uniquely identify the remaining octants in the fashion illustrated in the chart below.

sign of i_x	sign of i_y	$|ix|>|i_y|$or$|i_x|<|i_y|$	octant no
negative	positive	less	1
negative	positive	greater	2
negative	negative	greater	3

negative	negative	less	4
positive	negative	less	5
positive	negative	greater	6
positive	positive	greater	7
positive	positive	less	8

The implementation of these rules on their own provide an eight point compass with a ±22.5° accuracy, which while not very precise may be adequate for some undemanding applications. There are other benefits in more sophisticated versions.

The first advantage of this technique is the Gray code like way in which the octant rules work. At each octant boundary only one of the rule parameters changes. For example at 45° no sign changes occur but the inequality between i_x and i_y changes direction. At 90° no inequality changes occur and the sign of only i_y changes. This property prevents large scale jitter and confusion which might otherwise occur at the octant transitions if the changes were not totally synchronous.

A second advantage is that in each of the octants, a linear function of either i_x or i_y can be identified which is virtually equal to the desired heading angle, to within a small error. In the 0° to 45° range if ki_x is interpreted as a radian angle it is in fact very little different from the appropriate arctangent for that octant. If $k=1.08$ the error in doing this is nowhere greater than about 1.25°. If k is ignored and the unit vector x component alone is interpreted as radians the error is never worse than 4.5°, permitting the implementation of a 5° precision compass very easily. Note that for this purpose the modulus of the x component is used, eliminating the need to consider signs. Note also that $|i_x|$ is less than $|i_y|$.

In the next octant, between 45° and 90°, ki_y interpreted as an angle and subtracted from 90° is very close to the correct heading. This pattern repeats around the full circle and leads to the following rule.

Whichever of the direction cosines is the smaller is interpreted as an angle and in odd octants is added to the nearest quadrant boundary, but in even octants is subtracted from the nearest quadrant boundary to obtain the heading. (If as is likely in a software implementation the octants are numbered 0 to 7 rather than 1 to 8, the odd and even should be reversed in the previous statement of the rule.) In conjunction with using $k=1.08$ this rule will provide almost ±1° precision in a software implementation requiring no trigonometric functions.
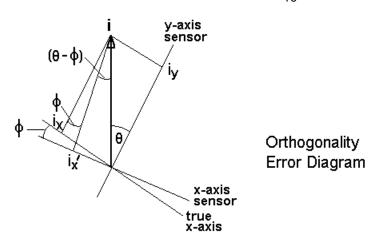
Alternatively, since the error is small, a very short lookup table of adjustments to be added to the heading obtained with $k=1.0$ will improve the precision to a level of around ±0.5°.

measured angle	added correction
0° - 20°	0°
20° - 29°	0.9°
29° - 33°	1.8°
33° - 37°	2.7°
37° - 39°	3.7°
39° - 41°	4.3°

It should not be assumed that this technique will enable a compass of this accuracy, only that the contribution to the total error budget from this source will be minimised to the extent indicated. Other sources may contribute larger errors in a final design if they are not suitably addressed.

One important potential error is lack of orthogonality in the axes of the two sensors. This can cause a smoothly varying error around the whole compass circle which can be much larger than those discussed above. Fortunately there is a relatively simple correction technique for this as can be seen from the following analysis.

In the diagram below **i** is the unit vector in the field direction, θ is the heading angle and ϕ is the small angular error by which the x-axis sensor departs from the correct right angled position. Also i_x is the true x component of the unit vector, i_y is the true y component and i_x' is the apparent (measured) x component of the sensor in error.

Orthogonality Error Diagram

It can be seen from the geometry of the figure that

$$i_y = \mathbf{i} \cos\theta$$
$$i_x = \mathbf{i} \sin\theta$$
$$i_x' = \mathbf{i} \sin(\theta - \phi)$$

Expanding the last relation

$$i_x' = \mathbf{i} \sin\theta \cos\phi - \mathbf{i} \cos\theta \sin\phi$$

Since ϕ is small $\cos\phi$ may be taken as one and $\sin\phi$ as just equal to ϕ giving

$$i_x' = \mathbf{i} \sin\theta - \phi \mathbf{i} \cos\theta = i_x - \phi i_y$$

Hence

$$i_x = i_x' + \phi i_y$$

It can be seen from this that the desired x component of the unit vector can be obtained from the apparent measured component, for all angles, by adding a small fixed portion of the y component. The proportion to be added is equal to the orthogonality error in radians.

The value of ϕ can be found, for a standardised and normalised sensor set by rotating the configuration in the earth's field and measuring the angle between the zero-field positions of each sensor. Alternatively the algorithm can be added retrospectively to an otherwise completed compass by checking the error during a full 360° rotation. The value of ϕ can be taken to be the average of the errors at 90° and 270° shown by the digital display. Such a determination needs only to be made once. If the microcontroller has no convenient way of memorising the correction it could alternatively be read from a trimpot value on power-up using RC timing or some other relatively crude analogue input method. The orthogonality then becomes one of the possible adjustments available to the user during the compass "boxing" exercise.

Probably the largest of the final observed errors will arise from failure to constrain the axes of the sensors to the horizontal plane. The errors depend on the direction of tilt and the heading, and on some headings small angular tilts will multiply up to much larger heading errors. For example on a north heading a 1° north-south tilt will produce no error, but a 1° east-west tilt will give rise to almost 2.5° of heading error. There is no simple cure for this other than effective double gimballing, suitably weighted, though short term averaging of multiple readings can improve the stability of the displayed output.

Another more complex alternative is to use a gravity sensor to determine the direction of the gravity vector and use trigonometric calculation to correct for the effects of tilt.

A final aspect of overall accuracy concerns the required precision of sensor readings. Interestingly, this is surprisingly lower than might be thought. Using the type of algorithm described earlier, a full 360° of 1° precision requires only that the measured components be slotted into one of forty-five almost evenly spaced bands. A relatively low six-bit binary measure will cover this. For a 5° precision a miserly four bits is adequate.

In conclusion, for those who may design, build and use a compass, in anger, the illusion of precision created by a 360° digital display may hide a lack of precision which is real. The cautious navigator rarely places total faith in compass accuracy and never trusts it as his sole instrument of navigation.

Three Sensor Systems

The use of three orthogonal sensors permits a three dimensional determination of both the magnitude and direction of the local field vector. This determination, however, is only made with respect to the axis system of the sensor configuration and not in any absolute space. Nevertheless it can provide the basis of many interesting applications other than the compass.

The compass is not the only device which requires absolute referencing. The extension to three dimensions permits in principle, the exploitation of the earth's field in "virtual reality" simulations, with the possible advantage of a "free roving" capability.

The potential to free rove in a large space is a consequence of the fact that the field behaves as a fixed orientation vector everywhere in the space. It can be converted to the forward looking vector of a virtual reality helmet, provided that it can be referenced to some absolute space.

The sensor configuration alone is not adequate for the following reason. For each angle the field can take with reference to the sensor axis system it is possible to rotate the axis system a full 360° around the field vector without any change in the sensor outputs. This ambiguity must be resolved to obtain the desired absolute reference and requires one more fixed orientation vector. The obvious one is the earth's gravity vector which will always provide a local vertical.

While a compass design can make use of a slow response device such as a mercury pool on resistive quadrants or a dielectric bubble on capacitive quandrants, these are useless for virtual reality applications. They usually do not have the angular range and the certainly do not have the speed required to follow rapid head movements.

The minimum requirement would a be speed compatible with a flicker free video image refreshing system, say 70 Hz, though some systems specify a response rate of 250 Hz. To satisfy this kind of requirement calls for something like an accelerometer configuration with a flat bandwidth of this order, which also extends down to DC. Such devices have been recently developed, spurred on by the automobile airbag market, but low-g versions are still very expensive in small quantities. They also have relatively poor signal-to-noise ratios at wide bandwidths. This is not too severe a problem, however, since like the compass algorithms already described, high angular accuracy can be obtained with low binary spans.

This is a fairly complex subject and details are deferred to a later application note.

Pseudo Three Axis Systems

There is a class of systems which use a three axis sensor sytem, but eliminate the need for the gravity vector by an additional constraint on one axis. They have the superficial appearance of three dimensional systems but do not exploit all the possible degrees of freedom. The searchlight is a classic example. It rotates in azimuth around 360° and could rotate in elevation through 180°, but does not have any mechanism for rolling around the remaining axis, since it would be entirely pointless.

If only the human head was satisfied by the same mechanism! Virtual reality would be much easier to implement.

12

The reason that this works is that as soon as the roll axis is constrained to remain horizontal, the rotational ambiguity around the field vector, mentioned previously, disappears. The trigonometry of the unit vector components is soluble and yields not only the azimuth angles, like a compass, but also the elevation angles.

Gun platforms fall into this category, as do steerable satellite type aerials, some robot mechanisms and any device which needs to point to a direction in space from a horizontal platform. Complex devices of this nature are probably well served by the expensive mechanisms they already employ, but there may be many simpler applications which could benefit from a low cost magnetic sensor configuration and a microchip solution, previously not economic.

One interesting idea may be exploitable by the economy end of the flying sport. Aircraft magnetic compasses are notoriously impossible objects, since even the addition of a gravity vector sensor solves nothing when it indiscriminately combines gravity with the accelerations of manoeuvering. In level flight a gimballed fluxgate compass works well but is useless in turns. Nevertheless it remains a reasonable tool to a power pilot in transit. Since a glider pilot spends a great deal of time in spiral turns, chasing thermals, it is not very appropriate most of the time.

If one axis could be reasonably constrained most of the time, a usable compromise might be achievable. Since the full horizontal rotation of 360° is required and roll angle can be large, the only restriction possible is in the pitch axis. Aircraft do not generally spend very long periods in pitching manoeuvers except during aerobatic activity. They may, however, alter pitch modestly during climbing, descending or turning. During any steady state version of these activities acceleration or deceleration along the line of the fuselage is small or nil.

If a three dimensional sensor configuration were gimballed transversely and suitably weighted, it could perhaps maintain the pitch axis of the sensor set sufficiently horizontal to allow the strategy under discussion to generate a heading and additionally a bank angle of acceptable precision.
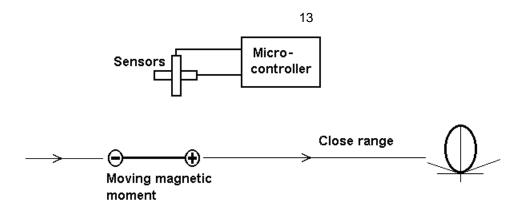
Whatever the precision, it would represent a vast improvement on the conventional fully gimballed compass and add half of an artificial horizon into the bargain, at lower weight and cost than any gyroscopic equivalent.
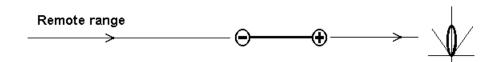
Three and Two Dimensional Ferrous Detectors

It is possible to elaborate the design of fixed single sensor vehicle detectors described earlier, with advantage, by using a two sensor version. Even when restricted to the horizontal plane, an orthogonal sensor set can provide more information, in the sense that it can provide both angular and magnitude signals for the anomaly caused by the vehicle passage.

An object with a magnetic moment possesses an external pattern of lines of force similar to that of a permanent magnet. This line of force pattern combines additively with the earth's field lines of force which consist locally of straight parallel lines. If the disturbing magnetic moment passes very close to the sensors it produces not only a variation in field magnitude but also large swings in the angular orientation of the detected field. If the passage is more remote from the sensors, not only is the magnitude of the signal reduced, but so also is the total angular swing.

While the time variation of these parameters gives some indication of the speed of the passage, if the magnitude of the signal is plotted against the angle in a polar diagram, what results is a time invariant "signature" of the object. In some sense this signature contains information about the range, since for a close passage it will have a large angle polar diagram and for a remote passage a small angle diagram. This range is not absolute as it will also depend on the equivalent magnetic length of the magnetic moment being observed, which is roughly correlated with the size of the vehicle most of the time. The fall off in field strength is proportional to the inverse cube of the ratio of the range to the magnetic length, so the field from large objects falls off more slowly than that from small ones.

13

Two Axis Ferrous Detector

The actual magnitude and angle variations will be quite small but can be increased to usable size by the digital heterodyne method or, in this case its software equivalent. The polar diagrams shown are oversimplified guesswork and not based on any tests.

However it seems that this is an area worthy of a little more serious research on practical real life situations, since it may resolve the problems of lane separation and vehicle classification in multiple vehicle studies.

Orientation Sensitivity Elimination

The ferrous detection systems discussed so far have been static ones and the fixed large signals produced by the earth's field can be relatively easily eliminated from the desired indications. In situations where the sensor configurations will inevitably be subject to unpredictable movement, the high orientation sensitivity becomes a serious disadvantage in the search for very small signals.

However if we consider a perfectly standardised and normalised, perfectly linearised and perfectly orthogonal sensor set, the problem is easy to deal with since the sum of the squares of the three outputs must always be equal to the field vector modulus squared, a scalar quantity without any orientation. The success of any real implementation will clearly be only a function of how close to perfection the above requirements come.

The mathematics are simple and readily implemented on computer or microcontroller.

The basic sensitivity of the sensors is adequate, matching of the calibrations is more constructive than absolute accuracy, orthogonality correction can be carried out to a fairly high degree, but non-linearity may give the most troublesome source of error.
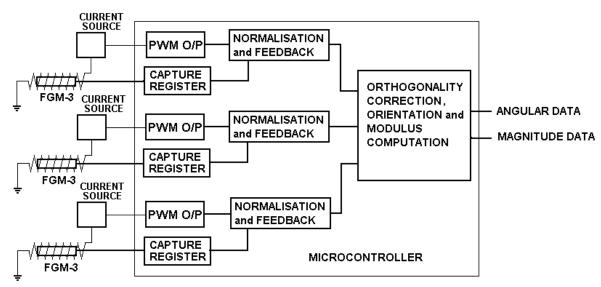
There is a technique which can be very helpful in these circumstances and this is the use of some sort of negative feedback, well known for its ability to improve linearity and stability.

The method consists of overwinding the sensor with a solenoidal coil in which a controlled field can be produced and automatically adjusting this field to cancel out to zero, the local field which the sensor would otherwise experience. The solenoid current giving rise to this cancelling field must be proportional to the local field being cancelled. Since the sensor only ever sees a zero field, its own non-linearity is no longer of consequence and the cancelling current is a direct and linear measure of the local field magnitude.

<div align="center">14</div>

This approach obviously calls for a digital-to-analogue converter to control the current in the cancellation coil, but with a microcontroller, this could be a pulse width modulated, single-bit, output and low pass filter arrangement, as used so successfully in many current low-cost digital audio devices. The software complexity increases but the hardware cost is still held low, probably calling only for a linear current generator of modest current capability.

In any case, total 360° orientation de-sensitising is not always needed and reductions in the angular variation achieved by other means will often considerably improve performance, as for instance in the case of a detector carried in a normally level vehicle or a neutral buoyancy weighted float, trailed just submerged. An error may exist in the output but it remains passably constant.

THREE DIMENSIONAL SENSOR SYSTEM

This type of system could find uses as a detector of seabed wrecks in modest depths or as a search tool in archaeological studies. Constructed with sufficient care, it provides a low cost and compact alternative to nuclear magnetic resonance devices in some applications.

<div align="center">15</div>