A MAGNETIC HEADING REFERENCE FOR THE ELECTRO/FLUIDIC AUTOPILOT

PART I

By Doug Garner
(EAA 15611)
MS494 NASA
Langley
Hampton,
VA 23665
804/827-3321

IN LAST SUMMER'S article on fine tuning the electro-fluidic wing leveler, I promised to put together a paper on the construction of a magnetic heading reference that would make an honest autopilot out of the wing leveler. Well, after much labor and sorrow, this is it. I hope you like it.

Magnetometers have had a nagging fascination for me for many years. The work on the original electro-fluidic autopilot, here at NASA, gave me an excuse for investigating these instruments in some detail. The magnetometer produced for this project was fairly simple and worked reasonably well, but it contained some components and materials that were hard to obtain. Since that time a lot of sparetime research has gone into the development of a magnetometer, suitable for autopilot applications, that was inexpensive and easy to build and contained a minimum number of exotic components.

The resulting instrument has been very gratifying. The only hard-to-get item is the iron alloy core. This is a standard transformer core which can be obtained from the manufacturer with only a minor degree of hassle. In fact, both core and windings have recently been made available to EAA members through a commercial coil-winding specialist.

The associated electronics consist of two readily obtainable integrated circuit chips, a transistor and a few resistors and capacitors. There's nothing very tricky about the construction and adjustment, and the output is a good, healthy DC signal that can be fed directly into the existing wing-leveler circuitry.

To the best of my knowledge, the following article will be the first do-it-yourself article on navigational magnetometers ever to be published anywhere.

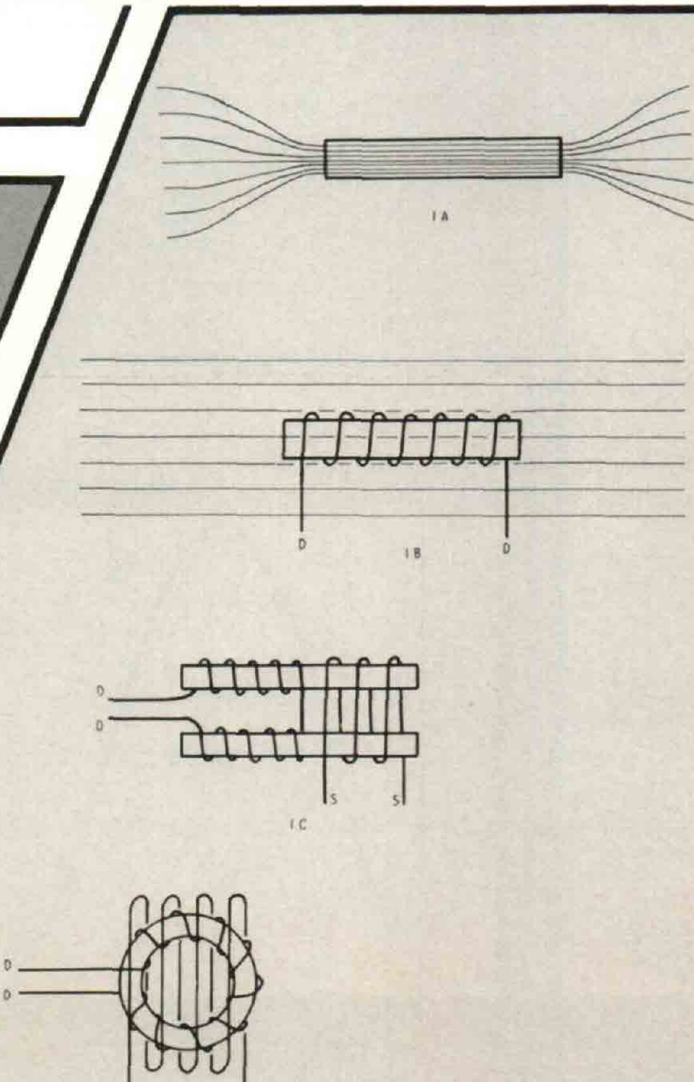

MAGNETOMETER GEOMETRY

FIGURE 1

Compasses and Gyroscopes

A simple wing leveler is a worthwhile contribution to the safety and convenience of flying, but it will not hold a constant heading for any length of time without the addition of some sort of heading reference. Commercial wing leveler autopilots have traditionally used the directional gyroscope for this function.

The directional gyroscope was developed in the late twenties as the result of the joint effort of Jimmy Doolittle and Elmer Sperry, Jr. to overcome a couple of basic faults in the magnetic compass, which limited its usefulness as a "blind flying" instrument (Ref. 1). The magnetic compass and the DG have been an inseparable team ever since, making up for each others deficiencies.

The free gyro has no idea which way is north, but has excellent short-term stability and will faithfully point in the direction you last told it was north for short periods of time throughout all but the most violent maneuvers. It must be reset, periodically, to the magnetic compass reading, either manually or automatically.

The magnetic compass, on the other hand, knows which way is north most of the time but gets flustered in rough air because of the extreme difficulty in applying adequate damping to its mechanical element. It is not much help in turns because of an effect called "northerly turning error", which we shall discuss later in wearisome detail.

Simplicity and reliability were major criteria in our original autopilot development (Ref. 2), so we wanted to stay away from gyros completely. We decided to take a good look at the compass as a heading reference and see if we could resolve its problems in some other way than by backing it up with a free gyro.

The magnetic compass is one member of a class of instruments called "magnetometers", whose function is to detect magnetic fields and to measure their magnitude and/or direction. Some of the more modern versions of the magnetometer contain no moving parts and thus avoid the damping problem. They also produce an electrical signal which can be fed directly into the wing leveler.

The two most promising types for this application appear to be the "Hall effect" magnetometer, which depends upon an obscure effect of magnetic fields on the conductivity of various metals and semiconductors and the "fluxgate" magnetometer, utilizing the phenomenon of magnetic saturation of an iron alloy core.

After some lab work, I concluded that the Hall-effect devices which were readily available at that time (1973) were just not sensitive enough for the job. About that time I got a fluxgate device working well, so that's the way we went.

You will note that, although this was a government research project, I did not take the easy way out and just go out and buy a commercial magnetometer. We were shooting for simplicity and low cost and I figured that if I could build it myself, it had to be simple, and if the materials didn't cost more than a few bucks, it couldn't be but so expensive to manufacture.

How Does A Magnetometer Work?

This section gets a little sticky in spots and if you don't go for this sort of thing, you can skip it and go on to the how-to-build-it part. It generally helps to know what's going on, but experience has shown that you can build a pretty good autopilot without really understanding it. Just hang in there.

The operating principles of the fluxgate magnetometer seem to be one of the best-kept secrets of the scientific community — at least I couldn't make much out of the usual textbook explanations. The following story, patched together from a number of different sources, would probably offend a magnetics expert, but it makes sense to me (I think).

In order to talk lucidly about magnetic devices, we must employ a gentle fiction called "lines of flux". These lines are used to illustrate the direction and intensity of a magnetic field in the same way that streamlines are used to illustrate the flow of air around an airfoil.

First, we need some iron alloy that is highly "permeable" and has a very sharp "saturation characteristic". This means that it has a very low "resistance" to magnetic flux, but that, once a certain density of magnetic flux is flowing through it, it will "saturate" and will then have a very high resistance to any more flux.

If we place a strip of this alloy parallel to the earth's magnetic field, as in Figure 1A, some of the lines of flux, due to the earth's field, will take a short cut through the alloy strip, since it offers less resistance to their flow than does the surrounding air. Now, if we place a coil of wire around the alloy strip, as in Figure 1B, and pass enough electrical current through the coil to "saturate" the strip, the lines of flux due to the earth's field will no longer want to flow through the strip, since its permeability has been greatly reduced.

Thus, the strip of iron alloy acts as a "flux gate" to the lines of flux of the earth's magnetic field. When the strip is

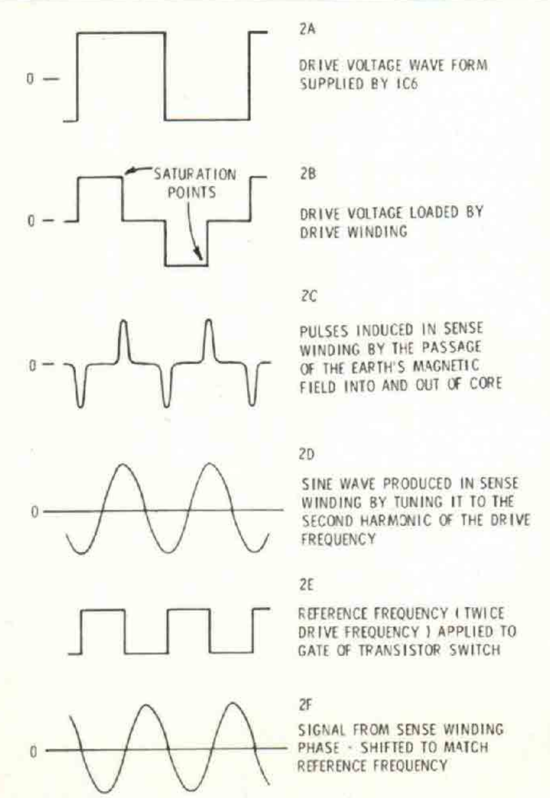

MAGNETOMETER WAVE FORMS

FIGURE 2

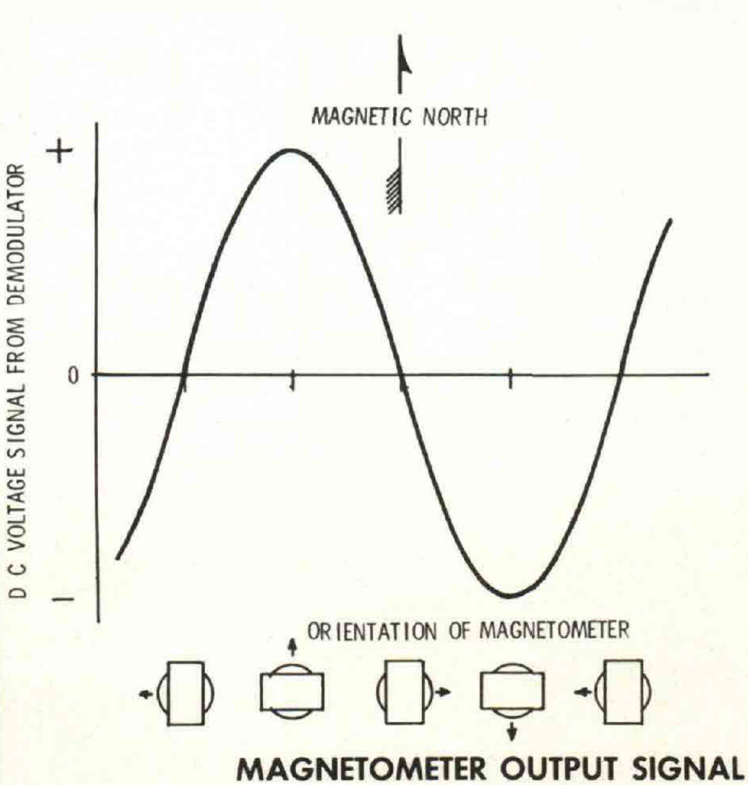

MAGNETOMETER OUTPUT SIGNAL

FIGURE 3

not saturated, the "gate is open" and the surrounding lines of flux bunch together and flow through the strip, but when we saturate the strip by passing a current of electricity through a coil wound on it, the "gate is closed" and the lines of flux pop out and resume their original paths.

Now the basic laws of electricity tell us that, when a line of magnetic flux "cuts", or passes through, an electrical conductor, it induces a voltage in that conductor. There must be relative motion between the line of flux and the conductor for this to happen. Well, if we apply an alternating current to the drive winding, D-D, of Figure 1B, we will be opening and closing the "flux gate" at twice the frequency of the alternating current, and we will have lines of flux from the earth's field moving in and out of the alloy strip at a great rate. If we can arrange to have these lines of flux pass through an electrical conductor (we'll call it the "sense winding") each time they pop into or out of the alloy strip, we will have a voltage induced in this conductor which is proportional to the number of lines of flux cutting it, and thus proportional to the intensity of that component of the earth's magnetic field which lies parallel to the alloy strip.

Right here, though, things get a bit sticky. We can't very well saturate the alloy strip without creating a lot of other lines of flux (conveniently omitted from Figure 1B), and we must sort these out from the lines of flux due to the earth's field, to get a meaningful signal. A number of sneaky schemes have been devised to avoid this problem (Ref. 3, 4). The one I used is shown in Figure 1C. Two identical alloy strips are used, and the saturation, or "drive", windings are arranged so that a closed magnetic circuit is formed. The lines of flux from the earth's field are still drawn into and expelled from both alloy strips each time the strips change from the saturated to the unsaturated state and back, and, if we put a "sense winding" (S,S) around the whole kluge, as in Figure 1C, this winding will be cut at each passage of the lines of external flux, and will produce just the signal we want. Now, however, the lines of flux resulting from the "drive" windings can build up and collapse without cutting the sense winding, and we are home free.

This is the kind of magnetometer we used in the original electro-fluidic autopilot (Ref. 2), using strips of Mu Metal as the core materials. Unfortunately, you can't just run down to the local Radio Shack store and pick up a couple of strips of Mu Metal. The stuff is awkward to get in the first place and, after it is cut to size, it must undergo an exotic heat treatment in an atmosphere of hydrogen to restore its magnetic properties.

In looking around for a handy source of core material for you homebuilders, I found that the resourceful William A. Geyger had anticipated this problem (Ref. 3, 5, 6). It seems that a fairly standard series of toroidal transformer cores are produced by several manufacturers in a wide variety of magnetic materials, including several that will serve well as magnetometer cores.

In Figure 1D we see that a toroidal core can serve the same function as the two alloy strips. With no air gaps at the ends, it is even a bit more efficient magnetically. It is considerably more difficult to wind, but it can be done, with sufficient patience and fortitude. These cores were never intended for magnetometers and are not very uniform magnetically, but with some attention to final adjustments, they will work fine in this application.

Figure 2 shows the details of what goes on in various parts of the magnetometer circuitry. The drive winding is excited by a square wave of suitable frequency and amplitude (Fig. 2A), so that the core is saturated halfway through each half cycle. When the core saturates, the impedance of the drive winding is reduced to a very low value, and virtually shorts out the amplifiers supplying the drive voltage so that the drive voltage is reduced to nearly zero for the remainder of the half-cycle (Fig. 2B). As the polarity of the drive voltage reverses, at the end of the first half-cycle, the core unsaturates and allows the drive voltage to reach full amplitude until about the center of the next half-cycle, when saturation again occurs, and the drive voltage is again reduced to near zero. (All this sounds kind of brutal as far as the driving amplifiers are con-

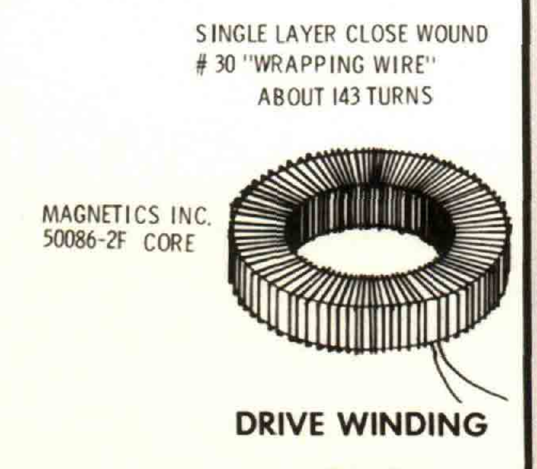

DRIVE WINDING

FIGURE 4

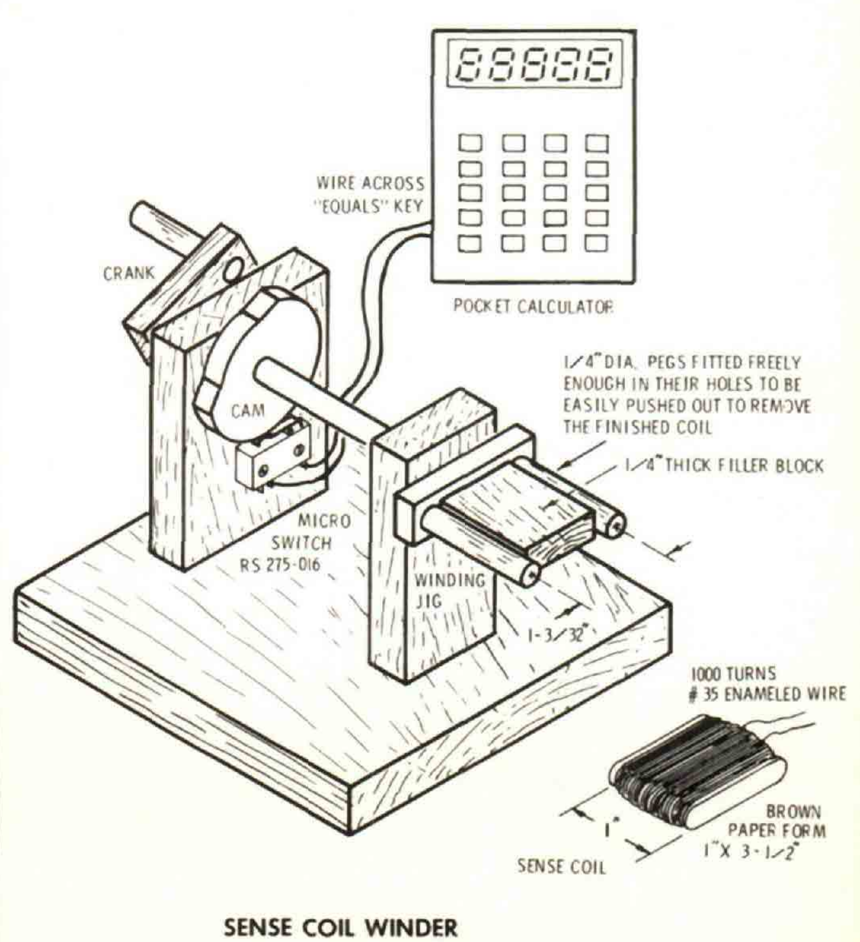

SENSE COIL WINDER

FIGURE 5

cerned, but they are designed to accept this sort of abuse without any problems.)

Now, as we have already seen, any external magnetic field will be drawn into the core when the core is unsaturated, and will be expelled when it becomes saturated. Each time these external lines of flux are drawn into the core, they pass through the sense winding and generate a voltage pulse whose amplitude is proportional to the intensity of that component of the external field which is parallel to the centerline of the sense winding. The polarity, or direction, of this pulse will be determined by the polarity of the external magnetic field with respect to the sense winding. When these lines of flux are expelled from the core, they cut the sense winding in the opposite direction and generate another voltage pulse of the same amplitude but of opposite polarity.

Figure 2C shows how these pulses look (since this particular magnetometer design results in very loose magnetic coupling between the core and the sense winding, a great deal of "ringing" will result from each pulse, and you will not be able to observe these pulses as they are shown in Figure 2C; but they are in there doing their thing, and I have shown them this way to make this explanation a little clearer).

Since these pulses represent the raw information that the magnetometer gives us about the external magnetic field, it is important to understand exactly what they are saying to us.

First, it should be noted that this pulse pattern is repeated twice for each cycle of the driving frequency (2A). This means that the information is coming out of the magnetometer at twice the frequency of the driving voltage, and this leads to the designation: "Second harmonic flux gate magnetometer."

As we have seen, the amplitude of the pulses is proportional to that component of the external magnetic field which is parallel to the centerline of the sense coil and it should be obvious that the direction, or polarity, of each pulse is determined by the direction of the external magnetic field with respect to the sense winding. Thus, these pulses give us information about both the amplitude and direction of the earth's magnetic field with respect to the sense winding. This is the kind of information we need for our direction reference but we still have to convert the pulses into a DC signal compatible with our wing leveler.

The first step is to tune the sense winding to a frequency of twice the drive frequency. This will convert our series of pulses into a sine wave (Fig. 2D), whose amplitude is proportional to the amplitude of the pulses and whose phase will reverse when that of the pulses does. (At this point you can see why we want the core to saturate halfway through each drive cycle, as in Fig. 2B. This leads to an even spacing of the positive and negative signal pulses, Fig. 2C, and allows these pulses to be efficiently converted into a sine wave by the tuned sense coil.) In order to convert this sine-wave signal into a DC signal, we must pass it through a "phase-sensitive demodulator". We'll get into the mechanics of this later, but for now we will just worry about what it does.

First, we need a reference voltage which consists of a square wave of twice the frequency of the drive voltage (Fig. 2E). This reference voltage is generated by the same oscillator that gives us the drive voltage and it always retains the same phase relation with it.

The phase of the sine wave signal from the sense winding must now be shifted a bit to get it exactly in phase (or 180 degrees out of phase) with the reference voltage as in Fig. 2F. Now the job of the phase-sensitive demondulator is to invert the polarity of the signal from the sense winding every time the reference voltage goes positive. This means that, for the conditions shown in Fig. 2E and 2F, the negative-going half of the sine wave will be flipped over and made positive and the positive-going half will be left as it is, giving the wave form of Fig. 2G. If this wave form is now passed through a low-pass filter, to strain out the lumps, we end up with a positive DC signal, whose amplitude will be proportional to that of the original sine wave signal.

If the direction of the magnetic field is now reversed with respect to the magnetometer, the phase of the signal, 2F, will change by 180 degrees with respect to the reference voltage, 2E, and the positive half-cycles of the signal voltage will be flipped over, giving us a negative DC signal out.

The net result of all this is shown in Fig. 3 where the sensitive axis of the magnetometer (the centerline of the sense winding) is kept horizontal and rotated through 360° with respect to the earth's magnetic field.

This, finally, is the kind of signal we need to serve as a heading reference for our wing leveler. We use one of the two zero crossings as our on-course reference. If the aircraft deviates in one direction, we get a positive error signal; if it deviates in the other direction, we get a negative error signal. The slope of the curve is opposite for the two zero crossings, so only one will be stable.

The simplest way to select the direction in which the airplane is to fly is to rotate the magnetometer with respect to the airplane, so that the airplane is headed in the direction we wish it to go when the magnetometer output is at the stable zero crossing.

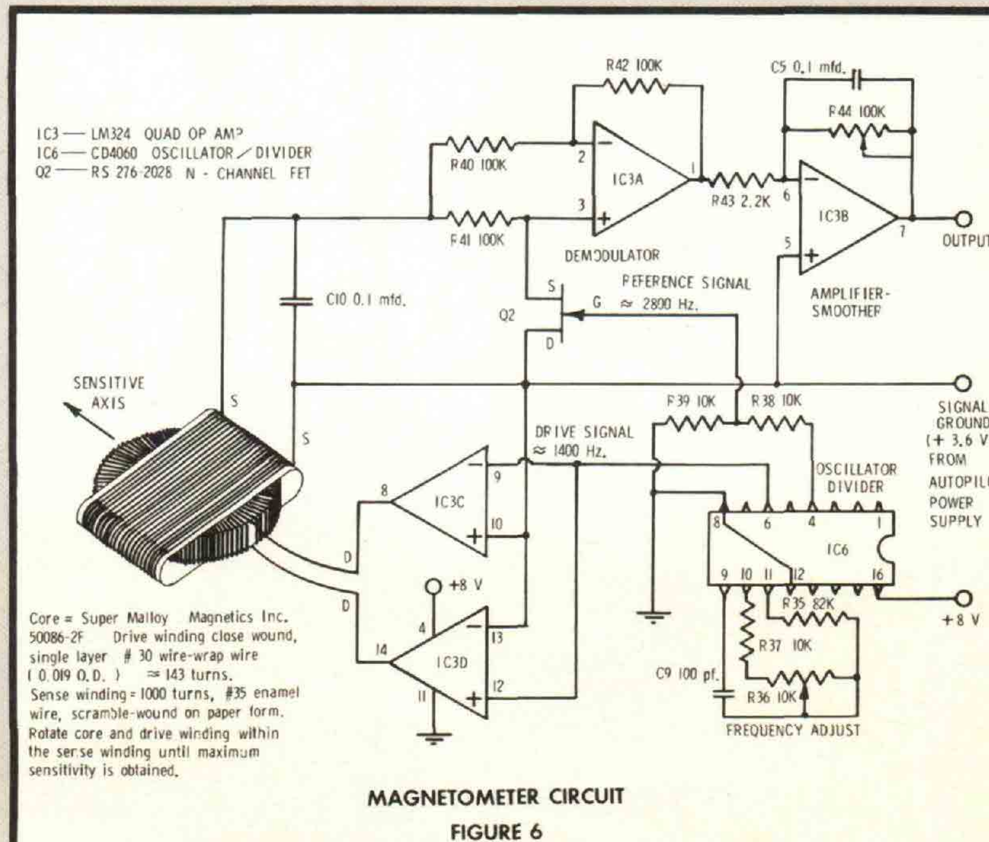

MAGNETOMETER CIRCUIT

FIGURE 6

The particular core I chose was Magnetics, Inc., 50086-2F, made from 0.002 inch thick Supermalloy tape (Ref. 7). There were a number of parameters involved in this selection and I tried to go for the best combination of performance and ease of winding. Other sizes and shapes will work, of course, but the correct combination of winding turns, drive frequency and voltage will have to be worked out for each. The material is the critical item. Two other materials are also suitable, 4-79 Mo-Permalloy and Square Permalloy, but these do not appear to offer any advantages, so if you are not a compulsive pioneer stick with the specified core.

As noted before, this stuff is sensitive to deformation and any shock, like dropping it on a hard floor or even a hard squeeze between your fingers, can disturb its magnetic characteristics. So treat the core gently and do your winding in a room with a soft carpet.

Applying the drive winding is kind of like knitting a sock; it takes some patience, so pick a time and a place where you won't be disturbed for awhile. I have tried various routines for applying this winding and I wrote up what I thought was the easiest for the Oshkosh and Lakeland forum notes, but last year at our Oshkosh autopilot workshop, Tom Kuffel of Seattle, Washington came up with a procedure that takes much of the heartache out of the job, so I've been doing it his way ever since:

Take nine feet of Radio Shack 30 gauge "wrapping wire" (278-503) and secure one end to something solid, like your bench vise. Now slip the core over the wire to about the center and start winding the free end on the core in a single layer, holding a bit of tension on the other end to keep the winding in place. The free end of the wire must be passed through the center of the core for each turn. Keep the turns tightly together on the inside of the core and evenly spaced at the outside. Avoid kinks and laps.

When you get the core about half filled, leaving a few inches of wire for connections, secure this end of the winding with epoxy, Scotch tape, bubble gum or something, and wind on the rest of the wire from the other end. You should end up with a single layer of winding filling the entire circumference of the core (about 143 turns) with both ends of the winding coming out at the same place (Fig. 4). When everything looks good, secure the ends with epoxy, so the winding won't come loose, and cut the ends off evenly to about three inches and strip about an eighth of an inch of insulation off each.

The "wrapping wire" has a plastic insulation which is much easier to strip than the enameled wire originally used and the overall diameter (0.019") gives about the right number of turns, when close wound, so you don't have to count turns as you go.

The sense winding is wound on a jig, as shown in Figure 5, so it will fit over the core and drive winding properly. Cut a piece of brown paper, like from a grocery sack, one inch wide and three and one-half inches long, and wrap this around the jig, lapping the ends on the flat side and holding them together with a bit of Scotch tape.

Put on one thousand turns of number 35 enameled magnet wire. I use a small lathe with a mechanical counter attached to the headstock for this operation. You can make up a coil winder, something like that shown in Figure 5, and try to keep count of the number of turns mentally, or you can add some kind of a counter. Most (but not all) pocket calculators can be converted into turn counters without much effort. If you can locate the connections to the "equals" key, run a couple of wires out to a microswitch which is actuated by a cam on your coil-winder shaft once per revolution. Key in "zero plus one" and then each time the "equals" contact is closed it will add one count to the display.

It is virtually impossible to lay this small wire on in even layers, but try for a fairly even distribution, allowing it to build up a bit at the center, and stay a sixteenth of an inch or so from each edge of the paper form. Every couple of hundred turns, swab on some epoxy or coil dope to hold things together but try not to glue the coil onto the winding jig.

Again, the wire size and number of turns are not all that critical but stick-

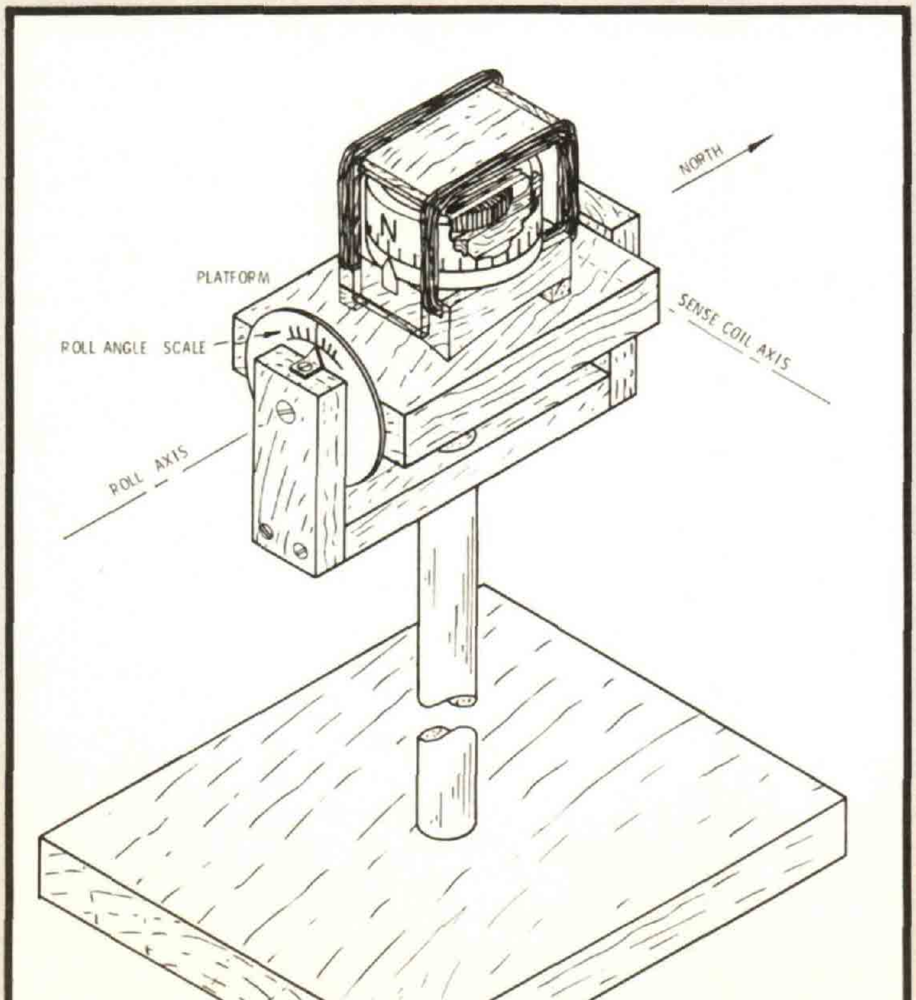

FIGURE 11 NTE CALIBRATION FIXTURE

ing closely to specifications will save some hassle when you get to the final adjustments. Clean the enamel off the ends of the sense windings (very carefully) and solder on some lead wires. The best thing I've found for this is the "multi-colored ribbon cable" available from many supply houses. Use a four-strand width, several feet long (get the most flexible stuff you can find, not the solid-conductor type stocked by Radio Shack).

After the leads are soldered to the coil windings, use epoxy to insulate the splices and to glue the sense coil lead wires to the coil so that any strain is taken by the lead wires and not by the fine winding wire. The drive winding can be soldered, temporarily, to the other two leads of the ribbon. The core and drive windings can be inserted into the sense coil at this time, but don't epoxy them in place yet.

Precision Winding, Inc. (Ref. 8) can supply you with the core, sense and drive windings for $18.70 if you want to save some work. These windings are very neat, and work just fine. These folks will also supply the bare cores for you masochists who want to roll your own.

The Electronics

At this point we'll take a look at the electronics that supply the drive voltage to the magnetometer windings and process its output signal. Figure 6 shows the basic circuit I worked out for this magnetometer. I took advantage of some modern electronic components and techniques to avoid a few of the problems that made the traditional flux-gate magnetometers so difficult to build.

The flux-gate magnetometers used in World War II aircraft were driven by 400 cycle power, since that was the only AC power available in the aircraft. The 800 cycle reference frequency needed for the demodulator was generated by a "frequency doubler" from the original 400 cycle AC, and this operation required an exotic saturating transformer that was pretty hopeless for the homebuilder.

Well, modern integrated circuit chips now allow us to generate most any frequency we want, in a very simple way. Frequency doubling is still a bit of a problem but "frequency halfing" is a snap with a counter/divider, so we just start with the reference frequency we want and divide it by two for the drive frequency. In fact, for about a buck-and-a-half we can get all this on one chip. The CD4060 gives us an oscillator and a fourteen-stage counter/divider all in one hunk.

Since we have all these counter stages available, we can reduce the size of the oscillator tuning capacitor by setting up the oscillator to run at a much higher frequency (179.2 kHz) and tapping off the reference frequency (2800 Hz) at the divide-by-sixty-four pin (pin 4), and the drive frequency (1400 Hz) at the divide-by-one-hundred-and-twenty-eight pin (pin 6).

The drive frequency is applied to the drive winding of the magnetometer by way of two op amps, IC3C and IC3D, Figure 6. This arrangement supplies sufficient power to drive the core and neatly avoids the DC bias problems inherent in most of the other drive circuits I have tried.

Saturation of the core is a function of drive voltage, number of turns and time. Given the particular core size and material, the number of turns in the drive winding and the drive voltage, the time for saturation turns out to be about 0.00018 seconds. Since, for reasons already discussed, we want the core to saturate about halfway through each half-cycle (Figure 2B), this means that one drive cycle wants to be four times this, or 0.00072 seconds long, giving the drive frequency of about 1400 Hz as previously noted.

Traditional phase sensitive demodulators used a lot of special transformers and diodes and stuff, and were generally pretty traumatic to get set up and working right. Even the more modern magnetometer circuits seem to be still hung up on this ancient technology. Some years ago I stumbled on a really neat demodulator, based on one of the quirks of op amps, and filed it away for future use. It turned out to be just the thing for this magnetometer. I can't recall the original source but it turned up later in one of Don Lancaster's excellent "Cookbooks" (Ref. 9). He calls it a "plus-one-minus-one amplifier."

IC3A and Q2, Figure 6, make up this demodulator. When Q2 **is not** conducting, the gain of the amplifier is **plus one**; that is, it has a gain of one and the polarity of the signal **is not** changed. This is the net result of a gain of **minus one**, due to the negative input, and a gain of **plus two** due to the positive input. (In this configuration, the numerical gain through the positive input of the op amp will be greater, by one, than the gain through the negative input — the quirk I referred to previously.) When Q2 **is** conducting, which happens every positive half-cycle of the reference voltage, the positive input is shorted out, leaving a net gain of **minus one**.

So, every half-cycle of the reference voltage, the signal voltage from the magnetometer is inverted, turning the AC signal into a lumpy DC signal whose polarity is determined by its phase relation to the reference voltage, as shown in Figure 2E, F, G. As noted before, the reference voltage comes from the oscillator/counter (IC6) pin 4. The voltage divider R38, R39 is used to avoid overdriving the gate of Q2 and stirring up a lot of unnecessary trouble.

IC3B smooths the resulting DC signal and raises it to a level suitable for use in the autopilot.

Adjustments

If you did a good job of duplicating the magnetometer windings and have led a good, clean life and belong to the right political party, you may be able to complete the final adjustments without the aid of an oscilloscope, but a scope will save you a lot of confusion and a dual-trace scope will be a downright luxury.

If you worked your way through the "how-does-it-work?" section, you will see that there are several things about the various wave forms that appear in the circuit that have to be just right, if you expect to get a usable signal from the output.

First, as we have seen, the core must saturate about halfway through each half-cycle of the drive voltage. Actually, this is not too critical, and we will be tweaking the drive frequency to get the other parameters in line, but it doesn't hurt to take a look at this wave form first. A scope connected across the drive winding should show a reasonable approximation to Figure 2B. Since both ends of this winding are being driven, don't ground the scope to any other part of the circuit. You should be able to center the saturation points by adjusting the frequency control pot, R36, in Figure 6.

For the next part of this operation, you need to locate the magnetometer well away from any large pieces of magnetic material. You should have it on one end of a length of that ribbon cable, three or four feet long. Place the magnetometer on a non-magnetic stand of some sort that can be rotated about a vertical axis. The core should lie in a horizontal plane. The calibration rig that you will need to make to check out the northerly-turning-error compensation (Fig. 11) can be used here.

The next two adjustments are rather tricky, as they have to be made pretty much at the same time. What you will, primarily, be trying to do, is to get the sense winding tuned to the second harmonic of the drive frequency. The trouble is that there will probably be only a couple of positions of the core within the sense winding that will give a reasonable approximation of a sine-wave shape to the signal.

If you fabricated the sense winding to exact specifications, it should tune to the right frequency (about 2800 Hz) with the specified tuning capacitor (C10, Fig. 6), and this frequency should fall within the range of the oscillator/divider (IC6, Fig. 6). Hook your scope across the sense winding with the scope ground attached to "signal ground" of the circuit. Connect the "external sync" terminal of the scope to pin 6, IC6.

Position the magnetometer so that the centerline of the sense coil is pointing approximately north-south and adjust the frequency of the oscillator, using R36, until you get a waveform of maximum amplitude. This waveform may be rather distorted and hacked-up looking because of magnetic irregularities in the core. Hold the sense coil fixed in its north-south orientation and rotate the core and drive winding within it. At some position you should observe a fair approximation of a sine wave, with an amplitude in the order of 0.8 volts peak-to-peak. Readjust the frequency again for maximum amplitude. To be sure you have the thing tuned to the right harmonic, rotate the whole coil assembly about a vertical axis while you watch the wave form coming out of the sense winding. The amplitude of your sine wave should diminish to a minimum as the magnetometer is rotated to an east-west orientation (there will be some weird looking noise left, resulting from higher harmonics). As you continue to rotate it, the sine wave will build up again to a maximum at 180 degrees of rotation, but it will now be 180 degrees out of phase with the original signal if your scope is synchronized to the drive frequency. When you have the best looking sine wave you can get, with the core centered in the sense winding, epoxy it securely in place.

If you have a good-looking sine wave, but little or no change in amplitude when you rotate the magnetometer in the earth's magnetic field, you are tuned in to the wrong harmonic. Try again — only the second harmonic will do.

If you have everything wired up right, and still don't get a good second harmonic signal from the sense coil, the resonant frequency of your sense coil and C10 is probably not in the frequency range that can be covered by the oscillator by adjusting R36. This could be caused by getting the wrong number of turns on the sense winding or by getting the wrong values for R35, R36, R37, C9 or C10.

It's probably easier to wind a new sense coil than to try to tune one that has the wrong number of turns but if you have a capacitor decade box, you can change the value of C10 a little at a time and keep swinging R36 back and forth until you find the proper tuning.

A scope connected to the output of the demodulator (IC3, pin 1) should give you a signal that looks something like Figure 2G (or an inverted image of 2G) with the sense coil axis lined up north and south. If the magnetometer signal is not phased up properly with the reference signal, you will get a hacked-up looking version of Figure 2G with part of the wave form on each side of the zero line. Small changes in drive frequency will give rather large shifts in the phase of the sense winding output signal, so, by tweaking R36 a bit, you should be able to bring in a good approximation of Figure 2G (or its inverse). The lumps will go down to zero and then reappear on the other side of the zero line as the magnetometer is rotated 180°.

A voltmeter connected from the output of the amplifier/smoother (IC3, pin 7) to signal ground will swing from plus to minus about 3.3 volts as the magnetom-

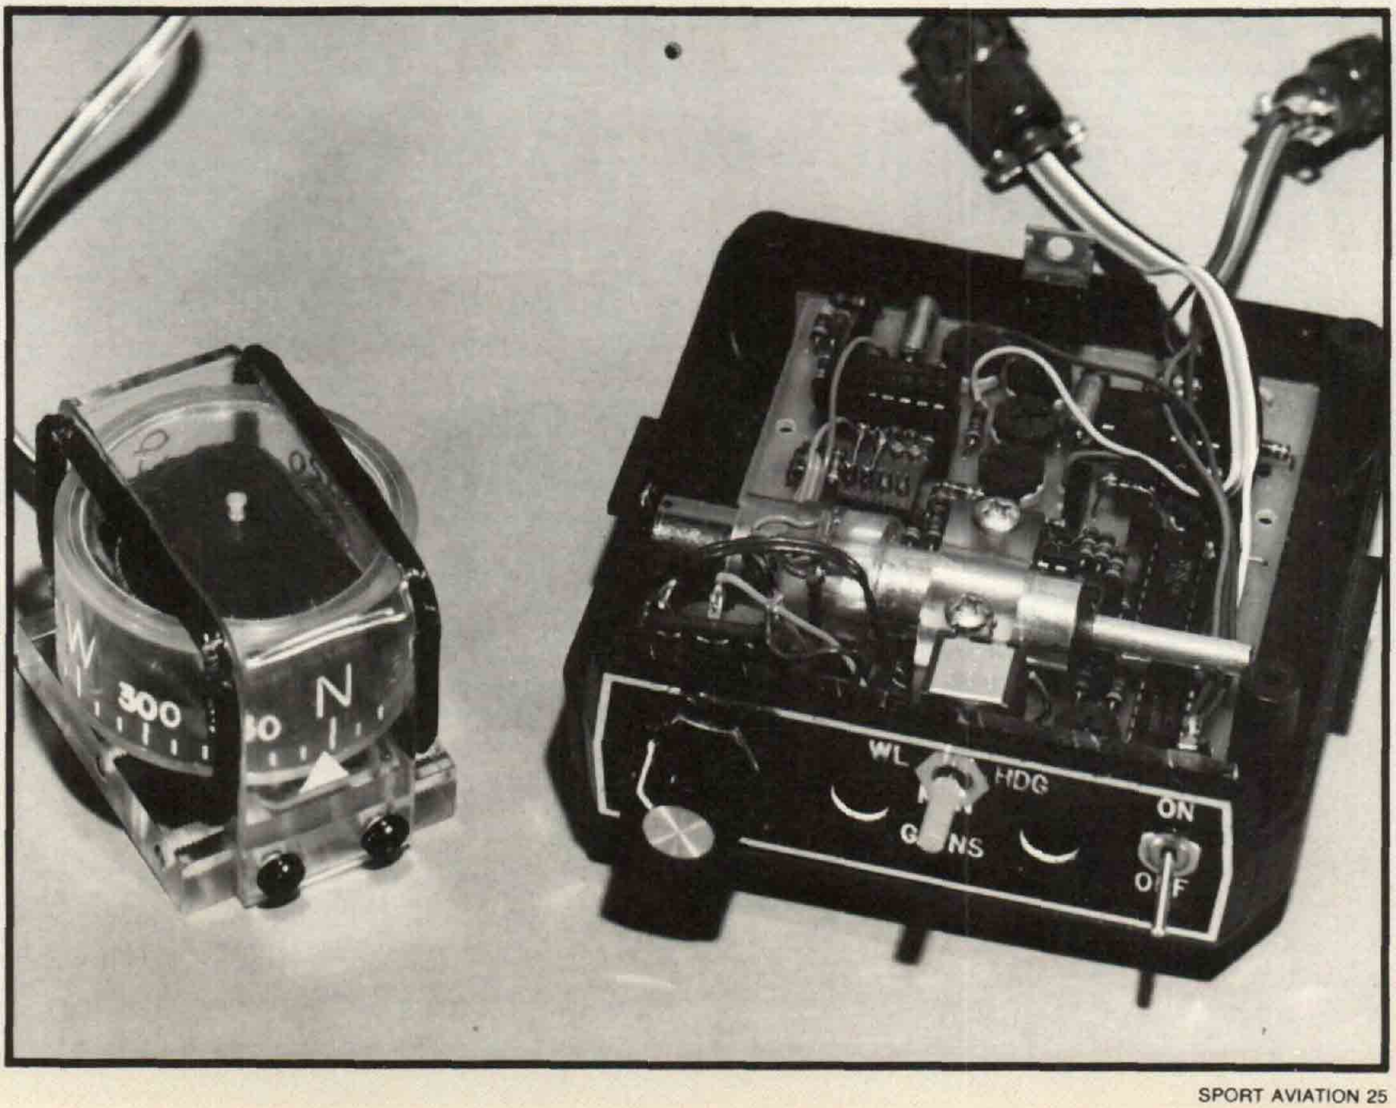

eter is rotated through 180 degrees. R44 controls the gain of this amplifier (the higher the resistance, the greater the gain), and will be used later to set the gain of the heading-hold function of the autopilot for optimum performance. For now, set R44 so the output of the amplifier/smoother is about two volts when the magnetometer is lined up north and south, and again tweak the oscillator frequency (R36) to get a maximum voltage reading. You can use this same ploy for setting the oscillator frequency and positioning the core in the sense coil, if you don't have a scope. Just line the sense coil axis up approximately north-south, and adjust the frequency (R36) for maximum voltage output from the amplifier/smoother (reduce the gain with R44 to avoid saturation, if necessary). Now, holding the sense coil still, rotate the core inside it, again looking for a maximum voltage reading.

Not all cores have the same sensitivity and you may have to reduce the value of R43 to get a usable output signal from one that has been dropped, or stepped on, or left out in the rain.

Construction Summary

The preceding discourse has been rather long and tedious because I have this hang-up about believing that people ought to know exactly what they are trying to do when they build something. For those who do not share my convictions, I have included the following summary on the construction and adjustment of the magnetometer to show you that it really isn't all that complicated:

1. Get yourself the right kind of core and wind the drive coil on it as shown in Figure 4.
2. Make a jig and wind the sense coil on it, as shown in Figure 5.
3. Build up the electronics as shown in Figure 6.
4. Adjust both the oscillator frequency (R36) and the position of the core within the sense winding for maximum output voltage swing from the amplifier/smoother when the magnetometer is rotated in the earth's magnetic field.

— Continued Next Month —

References

1. Dwiggins, Don: Famous Flyers and the Ships They Flew, Grosset and Dunlap, Inc., 1969.
2. Garner, H. Douglas and Poole, Harold E.: Development and Flight Tests of a Gyro-less Wing Leveler and Directional Autopilot. NASA TN D-7460, 1974.
3. Geyger, William A.: Nonlinear-Magnetic Control Devices, McGraw-Hill, 1964.
4. Geyger, William A.: Magnetic-Amplifier Circuits, McGraw-Hill, 1954.
5. Geyger, William A.: Flux-Gate Magnetometer Uses Toroidal Core, Electronics, June 1962.
6. Geyger, William A.: The Ring-Core Magnetometer, A New Type of Second Harmonic Flux Gate Magnetometer, AIEE Transactions 81 Pt. I, March 1962.
7. Magnetics, Inc., Butler, PA 16001, 412/282-8282.
8. Precision Winding, Inc., 3407 W. Douglas, Wichita, KS 67213. 316/942-2811.
9. Lancaster, Don: CMOS Cookbook, Howard W. Sams and Co., Inc., 1977.
10. Roskam, Jan: Airplane Flight Dynamics and Automatic Flight Control, Roskam Aviation and Engineering Corp., 519 Boulder, Lawrence, KS, 1979.
11. Garner, Doug: Fine Tuning the Electro-Fluidic Autopilot, Sport Aviation, August 1980.
12. Garner, H. D.: Magnetic Heading Reference, U.S. Patent 3,943,763, 1976.

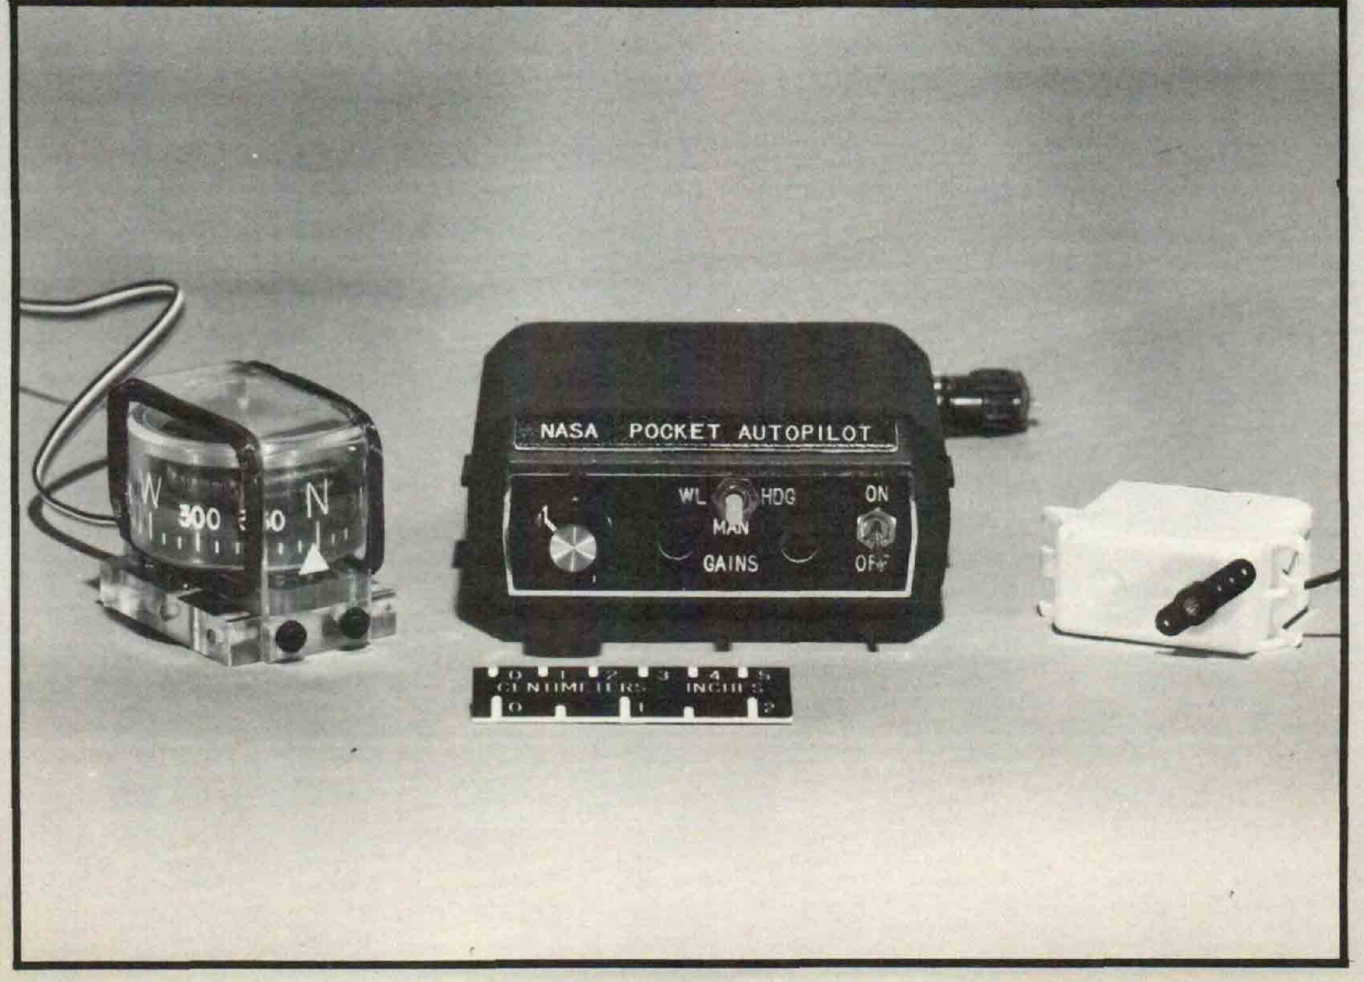

A MAGNETIC HEADING REFERENCE FOR THE ELECTRO/FLUIDIC AUTOPILOT

Part II

By Doug Garner
(EAA 15611)
MS494 NASA
Langley
Hampton,
VA 23665
804/827-3321

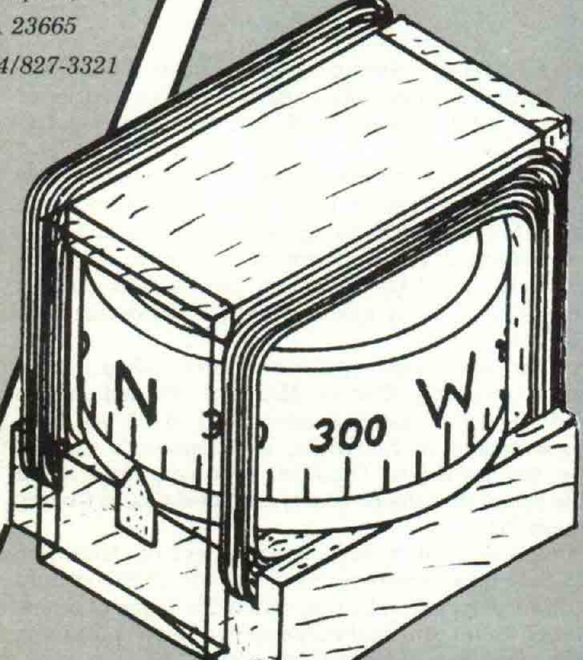

Application

MATCHING THE DYNAMICS of an autopilot to the dynamics of a particular aircraft for optimum performance is a complex operation, and is far beyond the scope of this article. If you'd like to dig into this fascinating subject, Reference 10 covers it better than any other text I have seen (I've just finished taking a course in airplane dynamics given by the author, so I could be prejudiced).

For our simple autopilot, we can just add the output signal from the magnetometer to that of the rate sensor in the wing leveler (Ref. 11), using the wing leveler for stability and the magnetometer as a sense of direction. This works out pretty well and just leaves us with the gains of the magnetometer and the wing leveler to adjust by cut-and-try for best performance. You could do better, but not without a lot of hassle.

Northerly-turning-error compensation is a nice, but not entirely necessary, feature. Without northerly-turning-error compensation, you avoid some construction and calibration problems but you must set both the aircraft and the magnetometer on the desired course before turning things over to the autopilot. With NTE compensation you can dial in any course you want on the magnetometer and the airplane will turn around and go that way. This is not exactly total automation, since you will have to hold the airplane's nose up while the autopilot is making the turn. We will look into NTE compensation later in this article. You can always add it after you get all your other problems under control.

Figure 7 shows one simple way to mount the magnetometer in the airplane. In this case, the four-wire ribbon cable connecting the magnetometer to the autopilot is threaded through the pivot tube and allowed to hang in a loop behind the instrument panel where it will periodically become wound up in a kinky wad and have to be unwound. A more elegant solution for the electrical connections would be a set of flat slip rings, fabricated from printed circuit board, but these will pose their own problems. Some kind of spring loading, to increase the load on the friction washer, will probably be needed to keep the thing from creeping around under vibration. Do not attach the scale permanently to the magnetometer assembly until you get to the calibration procedure.

Be sure to avoid any magnetic materials in this construction and do not use any large pieces of non-magnetic metals in close proximity to the magnetometer coils as the eddy-currents induced in them will reduce the sensitivity of the device.

Try to stay away from other magnetic devices in the airplane. Mount the magnetometer at least one foot away from your wet compass (two feet from a Hamilton vertical-card compass). An overhead mounting position would be best for some aircraft but watch out for those radio speakers they like to install in commercial airplanes; they create a magnetic field you wouldn't believe. Of course, the best place for the magnetometer is out in a wing tip or back in the tailcone, but you would need some way to rotate the thing from the cockpit to set your

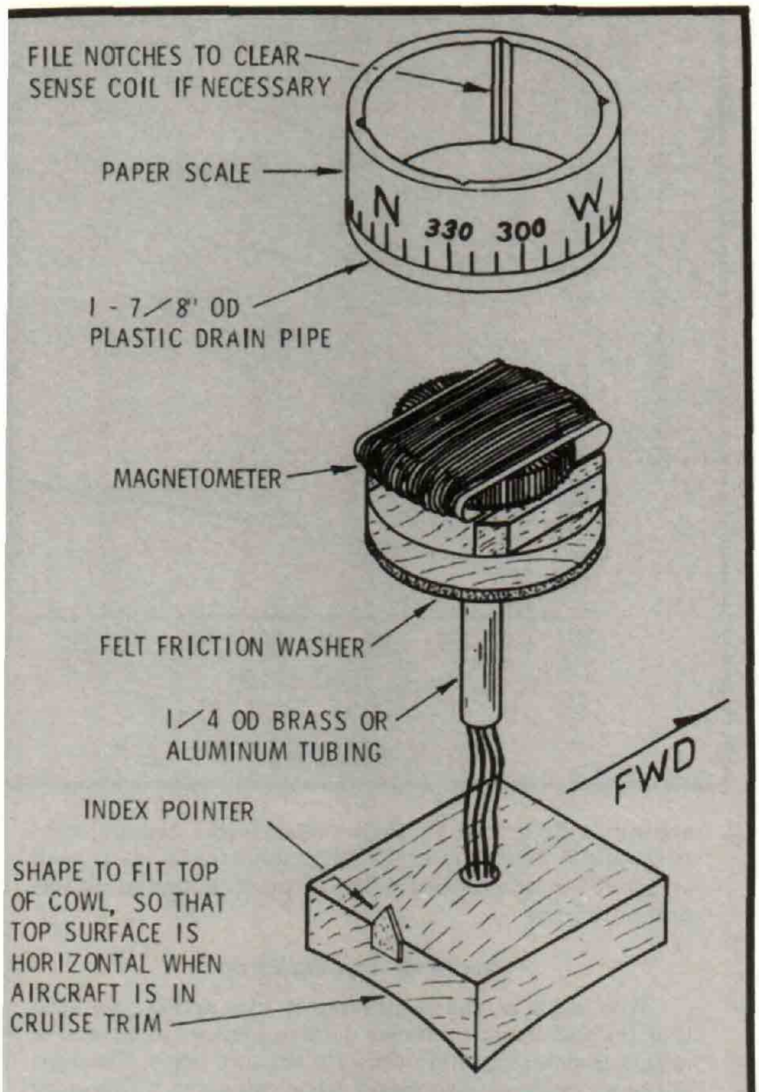

MAGNETOMETER ASSEMBLY

FIGURE 7

PARTS LIST PA-4 AUTOPILOT

ALL FIXED RESISTORS 1/4 WATT EXCEPT R50

R1, R2	10,000 ohms matched pair
R3, R4	4,700 ohms matched pair
R5, R6	220 ohms matched pair
R7, R8	1.0 megohm matched pair
R9, R10, R31	1,000 ohms
R11, R48	1,000 ohm trim pot - Jameco 63P102
R12, R22, R37, R38, R39, R45	10,000 ohms
R13	5,000 ohm trim pot Jameco 63P502
R14, R28	3,300 ohms
R15	27,000 ohms
R16	1.0 megohm
R17, R19, R27	5,600 ohms
R18, R25, R40, R41, R42	100,000 ohms
R20	10,000 ohm trim pot RS 271-218
R21, R46	1,200 ohms
R23	10,000 ohm volume control pot with knob - linear taper
R24	120,000 ohms
R26, R33, R34, R36	10,000 ohm trim pot Jameco 63P103
R29	270,000 ohms
R30	270 ohms
R32	47,000 ohms
R35	82,000 ohms
R43	2,200 ohms
R44	100,000 ohm trim pot RS 271-220
R47	1,500 ohms
R49	330 ohms
R50	0 to 3 ohms - 1 watt
C1, C2, C3, C5, C7, C10	0.1 Mfd. Mylar or Metalized Film
C4	1.0 Mfd. Metalized Film or Tantilum
C6, C8	0.01 Mfd. Mylar or Metalized Film
C9	100 pf.
C11, C12, C13, C14	10 Mfd. 25 WV Electrolytic
IC1, IC2, IC3	LM324 (RS 276-1711) Quad Op Amp Pin 11 = Grd.; Pin 4, IC1 = +12 V Pin 4, IC2 and IC3 = +8V
IC4	LM 334 (RS 276-1734) Current Source
IC5	LM556 Dual Timer (RS 276-1728)
IC6	CD 4060 Oscillator/Divider (Jameco)
VR1	LM 340-8 Voltage Regulator 8 V "
VR2	LM 340-5 Voltage Regulator 5 V "
Q1	2N4401 (or equiv.)
Q2	RS 276-2028 N channel FET
D1	IN914 (RS 276-1122) Silicon Diode
T1, T2	Thermistors Fenwal GB 32L1 matched pair
SW1	SPST toggle switch (RS 275-612)
SW2	DPDT Center-Off toggle switch (RS 275-620)
Blower Motor	Micro Mo 1516-G-012 S
Instrument Case	RS 270-285

SCHEMATIC PA-4 AUTOPILOT

FIGURE 8

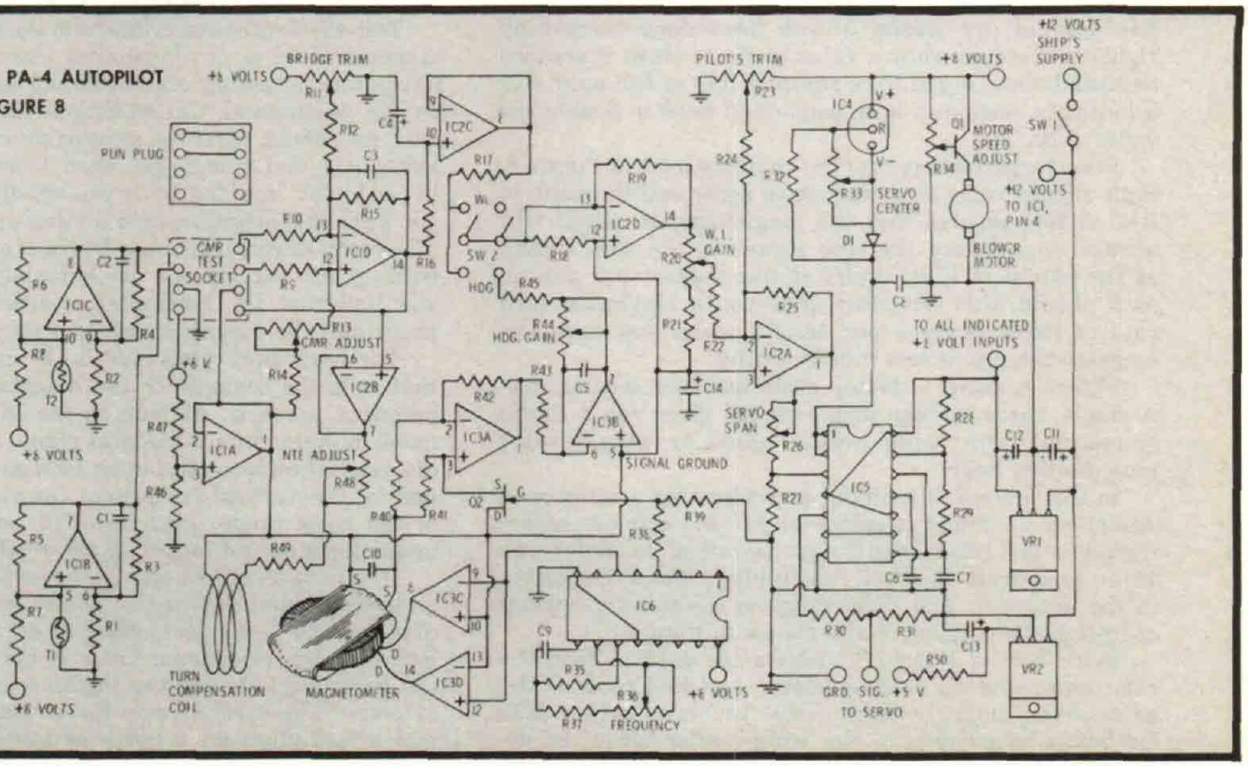

course (like in the Spirit of St. Louis). Two remotely-mounted magnetometers, at right angles to one another, with a resolver on the instrument panel to sort out their signals and select the desired heading, would be the most sophisticated solution but we will leave that as a homework problem.

The output signal from the magnetometer is fed into the wing leveler at the point marked "auxiliary input" in Figure 3, Reference 11. There are a couple of minor problems involved in using this circuit with the magnetometer, and, if you haven't already built up the wing leveler of Reference 11, you will find the updated version, Figure 8 of this article, a little easier to deal with.

If you use the wing-leveler circuit from Reference 11 you will need a switch between the magnetometer output and Ra to switch the heading-hold feature on or off, and you should also substitute a double-pole-double-throw switch for SW2, wiring up the added pole in series with the magnetometer on-off switch so that the magnetometer can never be left connected when the wing leveler is switched to the "manual trim" mode. This is an unstable combination and could get pretty exciting.

The other problem involves the relative limiting of the magnetometer and the rate signals. We'd like the magnetometer output to limit at a signal voltage somewhat lower than that of the rate signal so that the rate sensor always maintains control of the airplane. Specifically, if you are going to opt for the northerly-turning-error-compensation bit, and make course changes just by resetting the magnetometer, you want the output from the magnetometer to limit at a voltage equal to that produced by the rate sensor at a three-degree-per-second turn, so as to limit large course changes to nice, sedate, two-minute turn rates and avoid vertical-bank maneuvers.

In the circuit from Reference 11 the auxiliary input is downstream from the rate gain potentiometer R18, which means that the ideal value of Ra will depend upon the setting of R18 required for optimum performance of the wing leveler for your particular airplane and control surface set-up.

If you go this route, you should first fly your autopilot in the wing-leveler mode long enough to establish a good setting for R18; then choose a value for Ra for proper limiting. Since the rate signal from IC2C limits at about five degrees per second (if you have done everything right), you would want a value of Ra of about seventeen thousand ohms if you were running R18 at full gain. For a half-gain setting of R18, you would need to double the value of Ra, etc.

I rearranged this part of the circuitry a bit in Figure 8. Both signals come in at the same point and the ratio of R17 to R45 insures that the magnetometer signal will always limit before the rate signal. If the rate signal, at the output of IC2C limits at five degrees per second, as it should, this setup will give you a maximum turn rate of three degrees per second when you turn the magnetometer to a new course setting.

Figure 8 shows a better mode-selection scheme, too. A single, three-position toggle switch gives you a choice of manual trim, wing leveler alone or wing leveler plus heading hold.

In the process of building and adjusting a number of autopilots we found it easier to put the common-mode-rejection test components (as described in Reference 11) in a separate "test set" which plugs into a test socket in the autopilot. The "Run Plug" is inserted for normal operation. This option is also shown in Figure 8.

Since R20 of Figure 8 controls the gain of both the rate sensor and the magnetometer, and R44 controls the gain of the magnetometer alone, you must adjust R20 for best performance in the wing-leveler mode, as described in Reference 11; then switch to the heading-hold mode and set the gain of the magnetometer (R44) as high as possible without causing the airplane/autopilot combination to hunt.

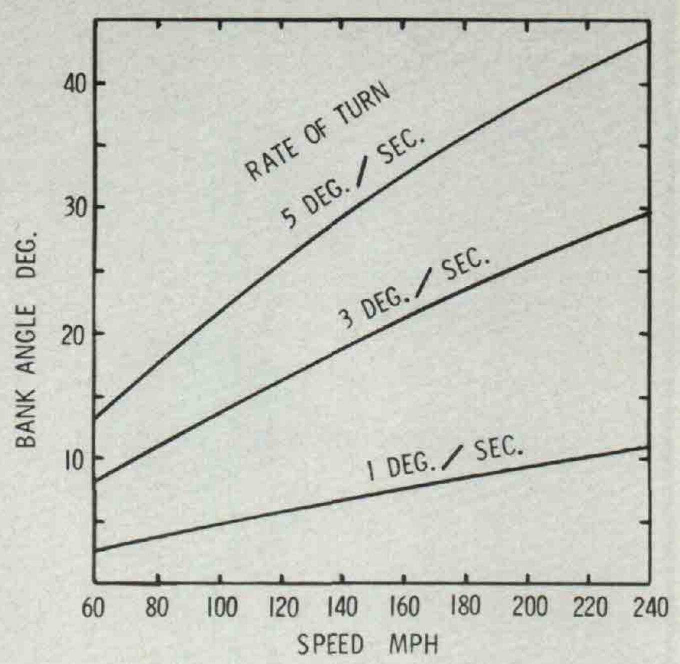

FIGURE 9

Northerly Turning Error

'Way back at the beginning of this article, we said that the two major problems with compasses were lack of adequate damping and northerly turning error. The first of these problems disappears when we go to a flux-gate magnetometer, as we no longer have any moving parts to damp. It is time now to examine the second problem, northerly turning error, which doesn't really bother us much until the airplane banks, either as a result of a turn or of severe turbulence.

This whole problem arises from the fact that the earth's magnetic field is not horizontal (except for those lucky navigators operating around Brazil and Central Africa). In the continental United States the lines of magnetic flux are about sixty to seventy-five degrees from the horizontal and things get even worse in Alaska. This is called the "dip angle." If you are flying over Virginia, you have this magnetic field coming up at you at an angle of seventy degrees from the horizontal. If you break this field up into two vectors, one horizontal and one vertical, you find that the horizontal component of the field is about half the magnitude of the vertical component.

The horizontal component is the only part of the field that tells the compass or the magnetometer which way magnetic north is. As long as the compass or magnetometer remains in a horizontal plane, it does not see that big vertical component, but as soon as the airplane banks part of the vertical component (proportional to the sine of the bank angle) gets into the sensitive axis of the magnetometer and louses up the reading.

If you went through a three-dimensional vector analysis of what the magnetometer sees in a banked turn (don't worry — I'm not about to do it here), you would find that the directional error is not only a function of the bank angle but of the instantaneous heading of the airplane. The error is zero for an east or west heading and a maximum for a north or south heading. Because of the way things add up, the compass error related to a north heading is more confusing to the pilot than that for a south heading, so that's what inspired the name "northerly turning error."

Anyhow, the result of all this is that any time the compass or the magnetometer gets out of a horizontal plane it starts to see enough of that insidious vertical component of the Earth's magnetic field to make it useless as a heading reference.

NTE Compensation

Several ways have been devised to avoid this problem. In some navigation systems the magnetometer is simply maintained in a horizontal plane by means of a free gyroscope, no matter what the airplane may be doing.

The more usual solution is to set a directional gyro to the compass heading when everything is straight and level, and then use the DG reading to get you through the turns and the turbulence.

Both these solutions involve the gyro syndrome — a feature I wanted to avoid in the interests of reliability and cost. This reluctance to put my trust in gyroscopes led to the notion that if you could compute the magnitude of the error, you could feed an equal and opposite signal into the system and just cancel out the error. This correction can be in the form of a magnetic field applied to the environment of the magnetometer itself, or an electrical signal added to the output from the magnetometer (Ref. 12).

The magnetic-field version of this correction signal is a bit easier to implement and will be the only one discussed here. (This same magnetic field correction can be applied to a compass. Used with the Hamilton vertical-card compass, with its excellent damping you get a reading during a turn that is almost identical with that of the directional gyro.)

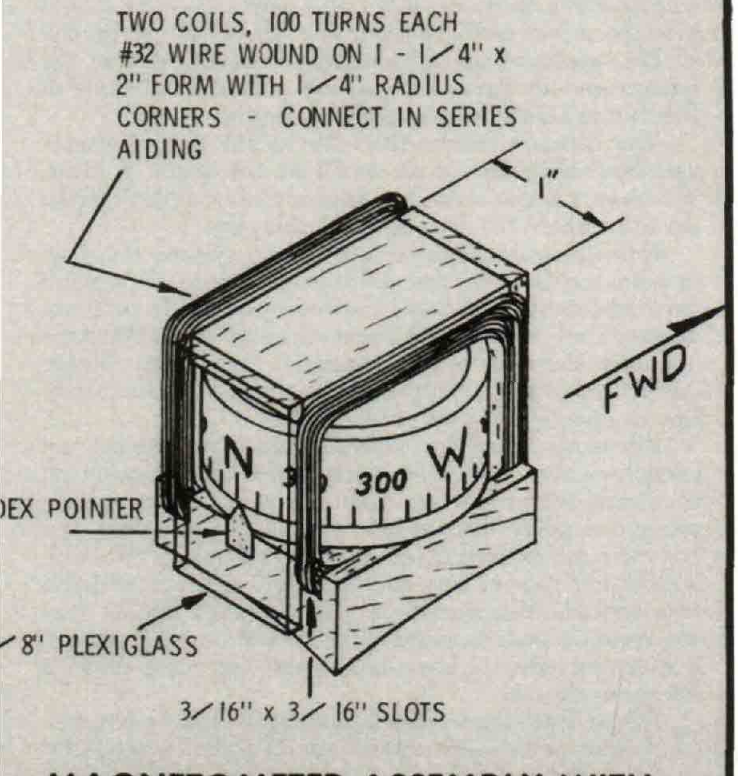

MAGNETOMETER ASSEMBLY WITH COMPENSATION COILS

FIGURE 10

Now the thing that makes all this practical is that we don't need a really precise correction for this northerly-turning-error business. When you are in a turn, the output from the magnetometer need only be good enough to start leveling out the airplane at about the right place to end up on the course you have set it to. As the airplane levels out both the error and the correction go to zero, and you can forget them until the next turn. This lets us get away with several approximations in the process of computing the error signal.

We know both the magnitude and the direction, with respect to the airplane, of that evil portion of the vertical component of the earth's magnetic field that comes in to foul things up during a turn. As we previously noted, the magnitude is equal to the magnitude of the vertical component times the sine of the bank angle and the direction will be parallel to the airplane's pitch axis, or wing span, since the airplane will have rotated about its longitudinal, or roll, axis.

This means that we can always apply the corrective magnetic field parallel to the pitch axis, that the magnitude of the correction will be proportional to the sine of the earth's field, and that the polarity of the corrective field will depend on which way the aircraft is banked. Since the earth's field doesn't vary much over rather long distances, we can set that in as a constant with the provision that we may want to make some adjustments if we fly over Alaska or Brazil.

Now, for reasonable bank angles, we will be operating in a fairly straight portion of the sine curve so we can use the bank angle multiplied by a suitable constant in place of the sine of the bank angle.

Well, we don't really have a measure of bank angle handy but what we do have is our little fluidic rate sensor, measuring rate of turn, and, for a coordinated turn, the bank angle is a function of turn rate. It also happens to be a function of the airplane's speed, which means that we can only calibrate the system correctly for one particular speed, but then, we're getting this whole compensation system almost for free so don't knock it. This bank-angle, turn-rate, speed function is plotted out in Figure 9 because we will need it later to calibrate the compensation system. This function, too, is a bit non-linear but good enough for our purposes.

O.K., so in case you have foundered in this vast morass of approximations, assumptions and half-truths, we will summarize the whole philosophy of this NTE compensator as follows:

1. The errant magnetic field that screws up the magnetometer during a turn comes in from the side of the airplane and is equal to the magnitude of the vertical vector of the earth's magnetic field at that particular location, multiplied by the sine of the airplane's bank angle.

2. If we can generate a magnetic field of this same magnitude, and sock it to the magnetometer in the opposite direction during a turn, we will cancel out that part of the earth's magnetic field that is causing the problem.

3. We can get away with a fairly rough approximation of this corrective field.

4. The voltage signal from the fluidic rate sensor is approximately proportional to the corrective field we need.

5. Our entire compensation system can consist simply of a solenoid coil enclosing the magnetometer. The coil is energized by the voltage output from the fluidic rate sensor applied through a dropping resistor which can be adjusted to take care of all the proportionality factors involved. The axis of this coil will remain at

right angles to the line of flight and parallel to the floor of the airplane (the "Y" axis to you aero engineering types).

Figure 10 shows one way of adding the compensation coils to the magnetometer assembly of Figure 7. The coils are wound on a form, as indicated, and tied with thread until they can be coated with epoxy. The magnetometer assembly is partially enclosed in a box made of one-eighth inch thick plexiglass and wood, and the coils are epoxied to the edges of this box to form a rigid assembly. The coils are connected in series in such a way that they act as one continuous coil, all wound in the same direction. You can get a lot more sophisticated than this but not much simpler.

The coils are driven by IC2B, Figure 8, through R48 and R49 to make the compensating field come out right. We shall get back to the business of calibrating the compensation coils after a brief excursion to get the rest of the autopilot working properly.

The Matter of Getting Everything To Move In An Appropriate Direction

At this time it would be a good idea to start worrying about getting our act together in such a way that the autopilot will correct any errors in the airplane's motion rather than making them worse.

As soon as you get the autopilot in operation, trim up the servo to the center of its travel, set the autopilot on a level surface and switch it to the wing-leveler mode. Now, rotate it slowly and smoothly in one direction, about a vertical axis, and note which direction the servo moves. You can now determine which way to hook up the servo linkage so that your ailerons will move in a direction to correct for this turn rate. Don't forget, a control surface will move in the opposite direction than the motion of the control tab.

The direction of motion of most servos can be readily reversed by reversing their signal leads, but those toy RC servos have this weird signal arrangement that cannot be handled so easily.

If your servo is already irrevocably installed and the linkage cannot be altered, and the stupid thing moves in the wrong direction, you must go back to the beginning and reverse the leads on the thermistors of the rate sensor to set things right. If you must do this, you will also have to repeat the CMR and Bridge Trim adjustments as detailed in Reference 11.

Now check the pilot's trim control and be sure the servo moves in a direction to produce a right turn when the trim knob is rotated clockwise. If not, reverse the connections to either end of the potentiometer's resistance element.

It is now time to check out the magnetometer.

Calibrating the magnetometer compensation system can get a little troublesome as you really need a rate table to do it up right. But this does not have to be done very accurately and if you built your rate sensor and its circuitry to the basic specifications laid out in Reference 11, and your compensation coils just like Figure 10, you can use my values and probably come out pretty close.

The way I do it, after the trim and CMR adjustments are completed (Ref. 11), is to set the rate sensor to the tilt angle to be used for the particular airplane (30 degrees is a good number), set the autopilot on a rate table and put in a yaw rate of three degrees per second (a standard, two-minute turn).

Take a reading of the voltage produced at pin 7 of IC2B, Figure 8, for this rate input. Measure this voltage with respect to signal ground. We get 1.4 volts out for three degrees per second, at a 30 degree tilt angle.

Now, take a look at Figure 9, feed in the cruise speed of your airplane (or the speed at which you expect to make heading changes), and see what bank angle goes with that three-degree-per-second turn rate. One hundred and forty miles per hour will give you an eighteen-degree bank angle, for example.

Next, you want to see how much current it takes through your compensating coils to take care of that eighteen degree (or whatever) bank angle. You will need a fixture something like Figure 11 to check this out. Be sure no magnetic material gets into the construction and that you set it up well away from any magnetic stuff when you use it. A level spot in the driveway is a good place, if you stay well away from bikes, lawn mowers and re-bar.

Mount the complete magnetometer assembly on the little platform of the calibration fixture so that the core lies in a horizontal plane and the axis of the sense coil is perpendicular to the roll axis of the fixture. The roll axis of the fixture should point north and south, and the side of the magnetometer assembly which has the index pointer should face south.

Assuming that you have already made the necessary adjustments to get a good output signal from the magnetometer circuitry, you should now be at one of the null points, and should measure zero output voltage from pin 7, IC3B to signal ground. If not, rotate the fixture slightly, one way or the other, about its vertical axis until this voltage is zero.

Now check to see if you are working with the correct null point. Switch the autopilot to the heading-hold mode, and rotate the fixture about its vertical axis a few degrees to the right (clockwise, looking down from above). The servo should move in the same direction that it does for a rate of turn to the right, applied to the rate sensor. If the servo moves the wrong way, rotate the magnetometer 180 degrees with respect to the index pointer and look for the other null point. This one should give the proper servo motion. The axis of the sense coil should, again, be at right angles to the roll axis of the fixture and to the index pointer, and the roll axis of the fixture should still be pointing north.

You can now attach the scale to the magnetometer assembly with the north mark at the index pointer. You may want to trim this up a bit later so don't make the attachment too permanent at this time.

Now, to calibrate the compensation system, put your foot on the base of the calibration fixture so it can't move about the yaw axis, and rotate the little platform about its roll axis by the amount you got from Figure 9 (eighteen degrees, in our example). The error voltage due to northerly turning error at this worst-case condition can now be read at pin 7, IC3B.

All we need to do now is adjust R48 so that we get just enough current through the compensation coils to bring this error voltage to zero when the rate sensor is being rotated at three degrees per second. Be sure that you have the coil hooked up so as to compensate a left bank with a left rate of turn and a right bank with a right rate of turn. You should be able to adjust R48 so that the error is near zero for either direction. If you have a rate table this is easy, but if not, you must make a few assumptions.

If your rate sensor is built and adjusted just like mine, you will get about 1.4 volts at pin 7, IC2B, Figure 8, for a three-degree-per-second turn rate, as noted before. This means that if you connect a 1.4 volt supply (a slightly-run-down flashlight cell would be close) to R48, R49 and the compensation coils in series, you should be able to adjust R48 to cancel the error due to the bank angle. Again, you must be concerned with polarity, and be sure

(Continued on Page 51)

A MAGNETIC HEADING REFERENCE . . .
(Continued from Page 32)

you get the coils tied into the autopilot in the right direction to cancel the error rather than to make it worse. You will probably want to mount R48 on the Magnetometer assembly so that you can try different adjustments in flight, but be sure you select a pot with no magnetic materials in it.

Well, that's about it for the magnetometer. If you feel you just must build one of the things and use it in an autopilot, please go back and read Reference 11 a couple of times more. There is a lot of basic design and adjustment information in this reference and you may have trouble building a satisfactory system without it.

You can do lots of neat things with this magnetometer besides fly an airplane with it — a shrimp boat autopilot, a digital Boy Scout compass, one input to an electronic navigation system, a flying saucer detector and a lot more.

So, What Else Is New?

During the past year we have managed to accomplish a few other things that may be of interest. We have a prototype sonic air pump in operation which is cheaper and easier to build than the motor-driven pump we've been using to drive the rate sensor; with no wearing parts, it ought to run a lot longer, too.

Most of the really good installations I have seen have used vacuum servos operating directly on the main aileron linkage. This makes a lot more sense than the toy RC servos, and I have gone back and redesigned the vacuum system we used in 1973 flight tests to bring it up to date. I found some great little solenoid valves on the surplus market that seem to be ideal for servo valves. I have also developed the interface circuitry to drive those electric torque-motor servos that Cessna uses in their Navomatic autopilots but haven't had time to wring this one out yet.

You can package the complete autopilot with any of these servo options in the little plastic instrument case called out in the parts list if you squeeze things a bit, and the whole installation, including the magnetometer and one toy servo, weighs in at about twelve ounces.

If you have an electro-fluidic (or any other kind of homemade) autopilot in operation in your airplane, I would be most interested in hearing about it. Please take time to drop me a note and tell me something about the autopilot, the type aircraft, the servo, the type and size of auxiliary control surface, if any. This lets the managment know that I'm earning my pay, and helps a lot when I get questions about installations in specific aircraft.

References

1. Dwiggins, Don: Famous Flyers and the Ships They Flew, Grosset and Dunlap, Inc., 1969.
2. Garner, H. Douglas and Poole, Harold E.: Development and Flight Tests of a Gyro-less Wing Leveler and Directional Autopilot. NASA TN D-7460, 1974.
3. Geyger, William A.: Nonlinear-Magnetic Control Devices, McGraw-Hill, 1964.
4. Geyger, William A.: Magnetic-Amplifier Circuits, McGraw-Hill, 1954.
5. Geyger, William A.: Flux-Gate Magnetometer Uses Toroidal Core, Electronics, June 1962.
6. Geyger, William A.: The Ring-Core Magnetometer, A New Type of Second Harmonic Flux Gate Magnetometer, AIEE Transactions 81 Pt. I, March 1962.
7. Magnetics, Inc., Butler, PA 16001, 412/282-8282.
8. Precision Winding, Inc., 3407 W. Douglas, Wichita, KS 67213. 316/942-2811.
9. Lancaster, Don: CMOS Cookbook, Howard W. Sams and Co., Inc., 1977.
10. Roskam, Jan: Airplane Flight Dynamics and Automatic Flight Control, Roskam Aviation and Engineering Corp., 519 Boulder, Lawrence, KS, 1979.
11. Garner, Doug: Fine Tuning the Electro-Fluidic Autopilot, Sport Aviation, August 1980.
12. Garner, H. D.: Magnetic Heading Reference, U.S. Patent 3,943,763, 1976.

(Photo by Ted Koston)

VH-MIW was one of two Corby Starlets flown in from Australia on the Oshkosh Express. The neat little VW powered, all-wood design will be marketed in the U.S. by HAPI, Inc., Eloy Municipal Airport, RR #1, Box 1000, Eloy, AZ 85231.