FLUXGATE SENSORS FOR A SPINNER MAGNETOMETER

by

ROBERT E. KRIDER

B.Sc. (Physics), University of British Columbia, 1975

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

in

THE FACULTY OF GRADUATE STUDIES

Department of Geophysics and Astronomy

We accept this thesis as conforming

to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA

July, 1985

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the head of my department or by his or her representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Department of Geophysics & Astronomy

The University of British Columbia
1956 Main Mall
Vancouver, Canada
V6T 1Y3

Date 13/SEPT/85

ABSTRACT

Noise sources in a Schonstedt spinner magnetometer, the primary laboratory instrument used in paleomagnetic research, and various methods for dealing with them are investigated theoretically and experimentally. Based on the results, an instrument is designed and constructed which replaces the Schonstedt sensor and magnetometer with the recently developed University of British Columbia ring core sensor and magnetometer, and the Schonstedt analyzer with an adaptive filter.

A 10dB improvement in the signal to noise ratio was achieved in the magnetometer, and a superior method for measuring this signal, using software, is suggested. A technique for faster noise reduction using a time-varying filter is developed analytically.

TABLE OF CONTENTS

Chapter 2 Design Considerations

Chapter 3 Construction and Performance

LIST OF FIGURES

Acknowledgement

For my very presence in graduate school, I am indebted to Don Russell, my thesis advisor, and to Bob Ellis, who, in spite of my somewhat dubious formal credentials, did battle with the bureaucracy on my behalf.

My thanks to Barry Narod and Peter Whaite for their direction and advice, and particularly to John Bennest, for leading me by the hand through the jungle of applied electronics.

Ted Irving and June Wynne at the Pacific Geoscience Center displayed remarkable patience as I struggled with their machines, and Dr. Irving was primarily responsible for financial support for the project. Other support came from NSERC operating grant A-720 held by Don Russell and strategic grant G-1050 held by Don Russell and Tomiya Watanabe. Computing facilities were provided by the Computing Center of the University of British Columbia.

Finally, I must express my appreciation to my fellow students, particularly Michael Schlax, for keeping me excited about The Earth.

INTRODUCTION

1. Spinner Magnetometers

Paleomagnetism, the study of the history of the earth's magnetic field and crustal evolution through rock magnetism, provides fundamental support to the theory of plate tectonics and continues to expand our knowledge of the geological history of the earth. The primary laboratory instrument used to measure the magnetization of rocks is the spinner magnetometer. An oriented rock sample from the field is spun near a magnetic detector, and the resulting signal analyzed to determine the remanent magnetization in the rock. The direction of the remanent field is taken to be an indication of the direction of the earth's field at some particular time; most commonly the time when certain minerals in the rock were cooled below their Curie temperature, hence "freezing" the ambient magnetic field. Spinning the sample (and measuring at the spin frequency only) provides improvement in signal to noise; hence weaker remanences can be measured.

Currently, two types of sensors are in use: cryogenic (SQUID) and flux-gate. The cryogenic are far more sensitive, but are also more than twice as expensive ($50,000 with analogue output) and have higher running costs; consequently, the flux-gate type, as represented by the Schonstedt instrument, is more common.

We may consider flux-gate spinners to consist of three subsystems—a mechanical assembly to hold and spin the sample, electronic hardware to extract the signal, and software for control and magnetic vector orientation calculation. This thesis is primarily concerned with the hardware portion of the system, which again may be divided into three sections—sensor, magnetometer electronics, and analyzer. [1] The magnetometer

[1] The hardware/software boundary is flexible; Schonstedt, for example has an "analogue" instrument and a "digital" instrument, the difference being in whether the analyzer (called by Schonstedt the processor) is implemented by hardware or software.

electronics are the drive and sense circuitry for the flux-gate element, and the analyzer extracts the in-phase and quadrature components at the spin frequency.

2. NASA Ring Core

The University of British Columbia's Geophysical Instrumentation Group has developed a superior ring core flux-gate (Russell and Narod, B.C. Science Council Report, 1982) based on technology originating primarily with M. Acuna of the NASA Goddard Space Flight Centre in cooperation with the Naval Ordinance Laboratory, especially with Jim Jaquet (Gordon et al. 1979). The important feature was the development of a low-noise ring core, because the majority of the noise in flux-gates is due to the sensor core.

Late in 1981, B. Narod substituted a University of British Columbia ring core sensor for the Schonstedt sensor in a spinner magnetometer at the Pacific Geoscience Centre in Victoria and observed in the time domain that the magnetometer output had significantly improved signal to noise ratio. As a result of this observation, a contract was concluded with the Earth Physics Branch, through the Department of Supplies and Services of Canada, to install the University of British Columbia instrument in a Schonstedt analogue spinner magnetometer and provide support for this research.

Chapter 1

Schonstedt

1.0 Introduction

The commercial instrument to be replaced, the Model SSM-1A, is manufactured by the Schonstedt Instrument Company of Preston, Virginia, and is an accepted industry standard. It is referred to as an analogue instrument because the Fourier components are extracted by analogue hardware. The system is controlled by an LSI 11 minicomputer, which turns the spinner motor off and on, reads switch and carriage positions, digitally samples and reads the Fourier components, and calculates the magnetization vector of the sample. The complete system also includes a T.I. Silent 700 terminal. The details of the hardware to be replaced are briefly discussed here (see also Schonstedt, 1972).

1.1 Sensor

A permalloy cylinder wound with a drive coil and a pick-up coil constitute the flux-gate sensor, with the geometry of an elongated ring core. The sense coil axis is parallel to the core axis, unlike the NASA flux-gate, which has the sense coil perpendicular to the core axis. Schonstedt's stated reason for this is that the drive field is always perpendicular to the sensitive direction of the pick-up coil, and hence eliminates contamination of the output with the drive signal; however, a toroidal element should have near zero external field and the second harmonic detect circuitry should, in any case, reject the fundamental drive frequency.

1.2 Magnetometer Electronics

The sensor is driven by an 8 kHz oscillator, carefully regulated to minimize second harmonic distortion. This drive saturates the permalloy core twice each cycle. The signal from the sense coil is first passed through a 16 kHz synchronous filter, whose

output is a square wave at the second harmonic (16 kHz) with amplitude proportional to the 16 kHz content of the modulator's input signal. This signal is amplified and passed through a synchronous detect which recovers the envelope of the square wave. A variable gain operational amplifier controlled from the front panel then feeds the signal to the processor.

1.3 Reference Generation

In order to relate the specimen orientation to signal phase, a slotted disk attached to the shaft strobes a light onto a photodiode once every revolution (about 5 Hz); an additional eight slots on the strobe disk and a second light and photodiode provide an 8 pulse per revolution (about 40 Hz) signal. The 40 Hz signal is divided by flip-flops to provide the processor with two 5 Hz signals, 90° out of phase with each other. The flip-flops are reset by the pulse from the first photodiode once every revolution to ensure valid relative phase at the beginning of each revolution.

1.4 Analyzer

Two identical circuits, each using one of the reference signals described above, extract the 5 Hz in-phase and quadrature components. The concept is identical to the magnetometer detection circuitry—a synchronous filter to produce a square wave at 5 Hz, and a synchronous detector to measure its amplitude. The bandwidth of the filter, and hence the total noise power, is controlled by the "integration time" switch on the front panel. A long integration time implies a slow transient response, narrow bandwidth, and low noise levels.

The output of the processor consists of two voltage levels proportional to the amplitude of the in-phase and quadrature components of the signal. These signal levels are fed to an A/D converter, which the LSI 11 samples at the end of the spin time to provide a direct measure of the component perpendicular to the spin axis of the sample's magnetic field at the sensor.

1.5 Miscellany

One decade of gain switching occurs in the processor, and three decades in the magnetometer. The switching in the magnetometer is done by field effect transistors so that only the D.C. switching levels applied to the field effect transistor's gates need to be run to the front panel, thereby reducing noise pick-up.

Neutralization of the magnetometer is accomplished by passing a current through a coil around the sensor to cancel any ambient field in the direction of the sensor. A front panel nulling meter is provided to set the current.

1.6 Magnetometer Performance

Figure 1 shows the noise spectrum at the output of the magnetometer electronics as measured by a Spectral Dynamics SD 345 spectrum analyzer. This is the signal the processor has to analyze to extract the Fourier Coefficients. The spike at 5.75 Hz is a signal from sample AK1-3 provided by Dr. E. Irving at Pacific Geoscience Centre, and represents the weakest rock the Schonstedt system can reliably measure with a 3 second time constant (corresponding to a 12 second spin time).

The general measurement problem is to determine accurately the magnitude of the two components at 5.75 Hz which produce the spike, that is, to duplicate the spectrum analyzer's calculations for a single frequency, using hardware. We may make some fundamental observations referring to Figure 1.

(1) The spike is spread out somewhat due to variations in shaft speed. Since, in the processor, the measurement is linked directly to the shaft, this source of inaccuracy is eliminated, so that the spike will be as close to a delta function as the bandwidth (governed by spin time) will allow.

(2) The minimum signal that can be measured is determined by the noise floor. Because the signal is deterministic and very nearly a delta function in the frequency

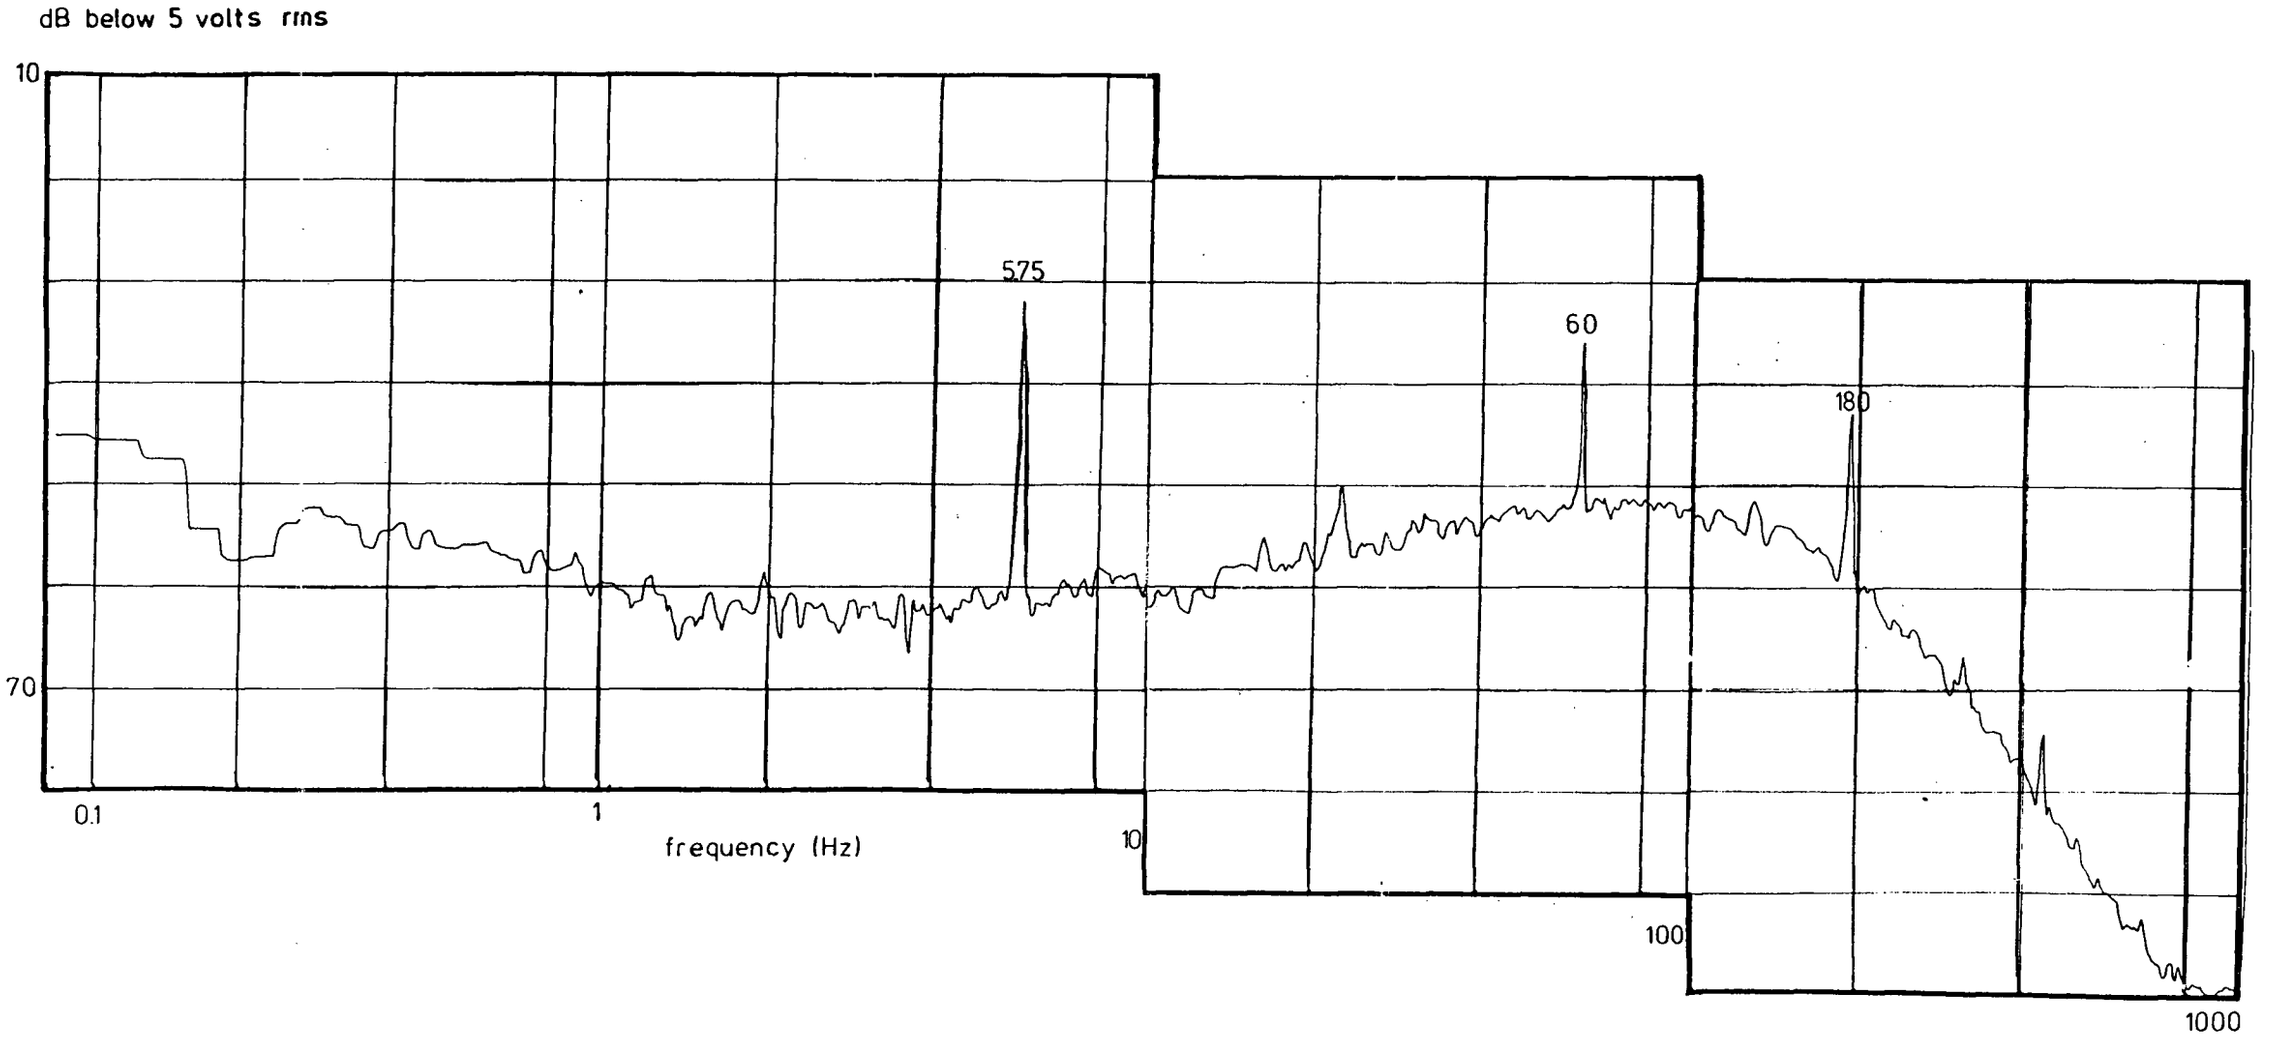

Figure 1: Schonstedt Magnetometer Noise Spectrum. Sample AK1-3 provides signal spike at 5.75 Hz.

domain, and the noise is stochastic, increasing the measurement time (which decreases the width of the frequency "bin" around 5.75 Hz) lowers the noise floor, but leaves the spike unchanged.

(3) In producing the spectrum of Figure 1, where measurement time is not a factor, it is appropriate for display purposes to average several measurements to reduce the variance of the stochastic noise floor. However, in the processor with a limited measurement time, a zero-bias, maximum-variance measurement is required. When using Figure 1, and subsequent spectra, it must be born in mind that from the point of view of the **processor** the mean noise floor will be at a different level relative to the noise spike; and that the variance will be much larger. The spectra remain very useful, however, for relative comparisons.

(4) The spectrum has most noise at low frequencies and falls off rapidly above 100 Hz. The most unexpected feature is the dip centered on the spin frequency. All the sensors assembled and tested by University of British Columbia's Geophysical Instrumentation Group to date have noise curves monotonically decreasing with frequency. The implication of this is that, while our sensors have substantially lower total noise power, at the frequency of interest, 5.7 Hz, the two noise levels are the same within a few dB. (More detailed comparisons appear in Chapter 3).

On the whole, the non-monotonic behavior of the noise in Figure 1 is quite difficult to understand. We have considered three possible explanations, as follows:

(1) We might be observing a monotonic noise curve through the transfer function of the magnetometer electronics, which may not be flat. This was tested by generating a signal near the sensor with a signal-generator driven coil over a range of frequencies. The resulting transfer function (Figure 2) is very nearly flat to 100 Hz, eliminating this possibility.

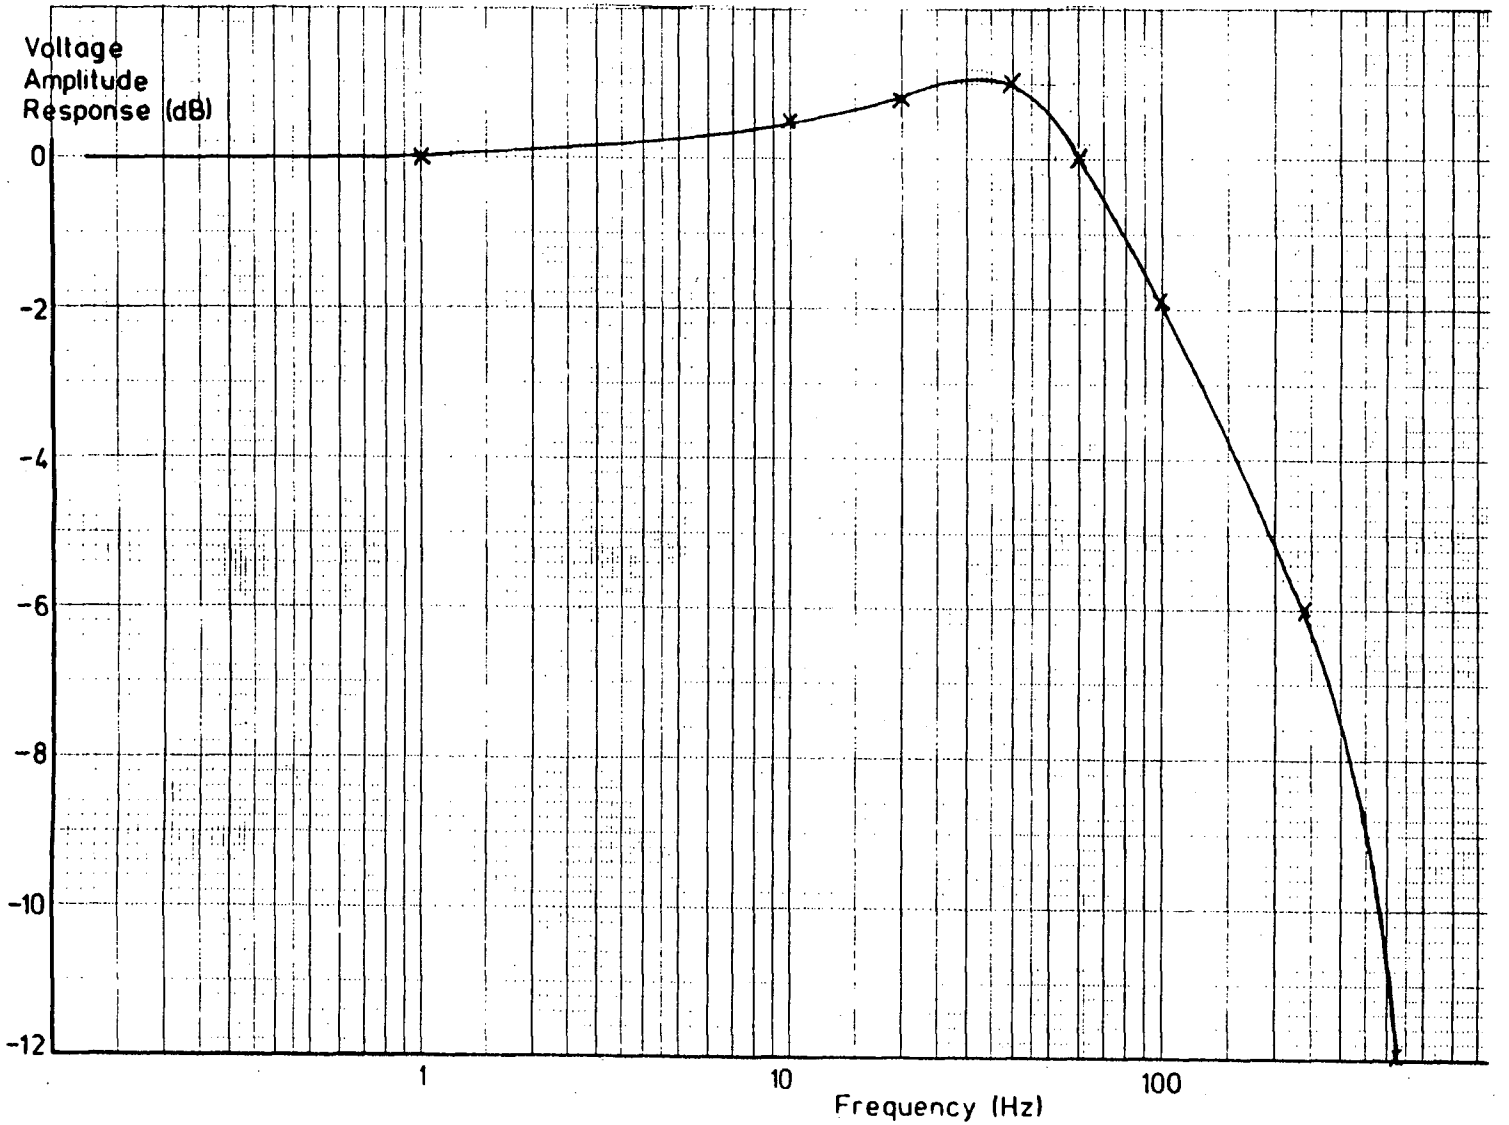

Figure 2: Schonstedt Magnetometer Transfer Function

(2) The magnetometer electronics may add significant noise to the sensor noise everywhere except at the spin frequency. This is difficult to test; however it is considered unlikely because normal noise levels due to electronic circuitry are generally much lower than the noise observed. It is generally assumed that the major contribution to the noise is the sensor.

(3) Figure 1 may represent the true sensor noise curve. If this is the case, and if it is engineered, then Schonstedt knows a lot more about noise mechanisms than is published; if it is fortuitous, it is plain to see why the spin rate is chosen near 5 Hz. It is difficult to imagine a physical mechanism producing such a complicated noise spectrum—further investigations would certainly be worthwhile.

1.7 Analyzer Performance

Because the processor heterodynes the 5.7 Hz spin frequency to D.C., it is of interest to observe the output spectrum at very low frequencies. The overall transfer function resembles a low-pass filter, with (for a time constant of 3 sec.) a corner at about 0.1 Hz and a 12 dB/octave roll off, consistent with the two sequential first order sections in the processor.

Three possible noise sources are identified, and for convenience are labelled wow, rumble and flutter in analogy with audio electronics.

Wow: variation in shaft speed which produces a broadened peak in the magnetometer output. It is assumed that the reference signal generation technique is adequate to deal with this source.

Rumble: movement of the sample holder along its axis while spinning. This can be reduced by carefully adjusting the shaft bearings. This noise will decrease with distance from the sensor, and should become negligible at large (> 10 *cm.*) sample-sensor separations.

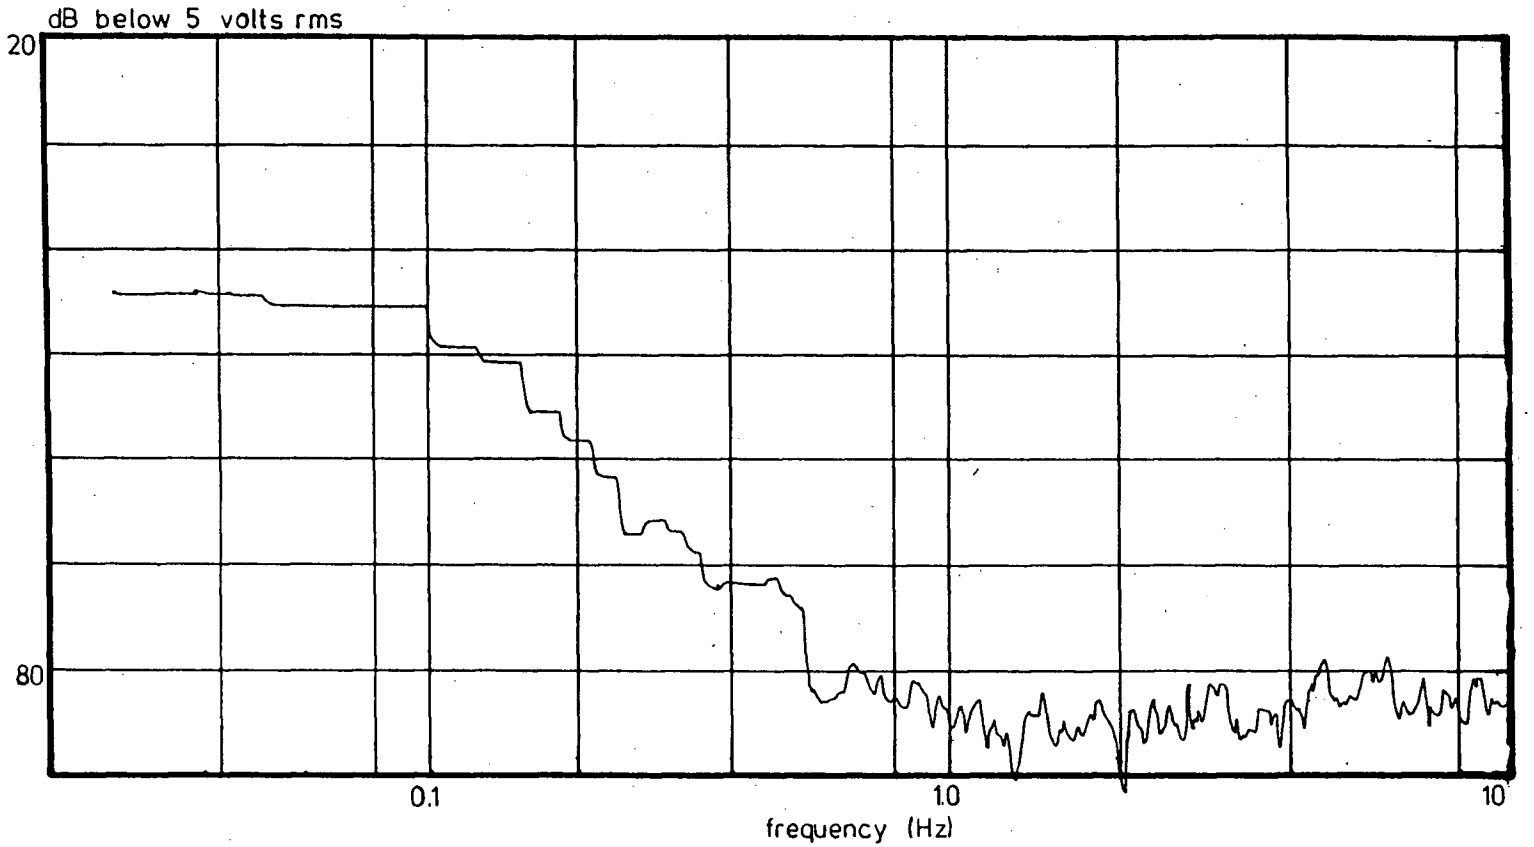

Figure 3: Flutter Test: Channel X , nail at maximum distance (20 cm.) as source.

Flutter: inconsistent reference signals (arising from, for example, variations in the photodiodes turn-on behavior from one slot to the next or noise on the reference line).

1.8 Flutter Test

In order to reduce rumble for this test, a strong sample (consisting of a nail taped to the sample holder) was spun at the maximum distance (approximately 20 cm) from the sensor that provided a good S/N ratio. Figure 3 is the spectrum of the Channel X output, which was indicating a D.C. level of 6.8 Volts. To make a similar measurement without using the shaft encoder, the configuration shown in Figure 4 was used. The analyzer received a very stable 45 Hz reference from a signal generator, and divided it to 5.6 Hz. The 5.6 Hz signal was tapped and used to trigger a second signal generator to produce a perfectly synchronous sinusoid, which drove a source coil placed near the sensor. The system was adjusted until the X Channel output measured 6.8 V, and the spectrum (Figure 5) taken.

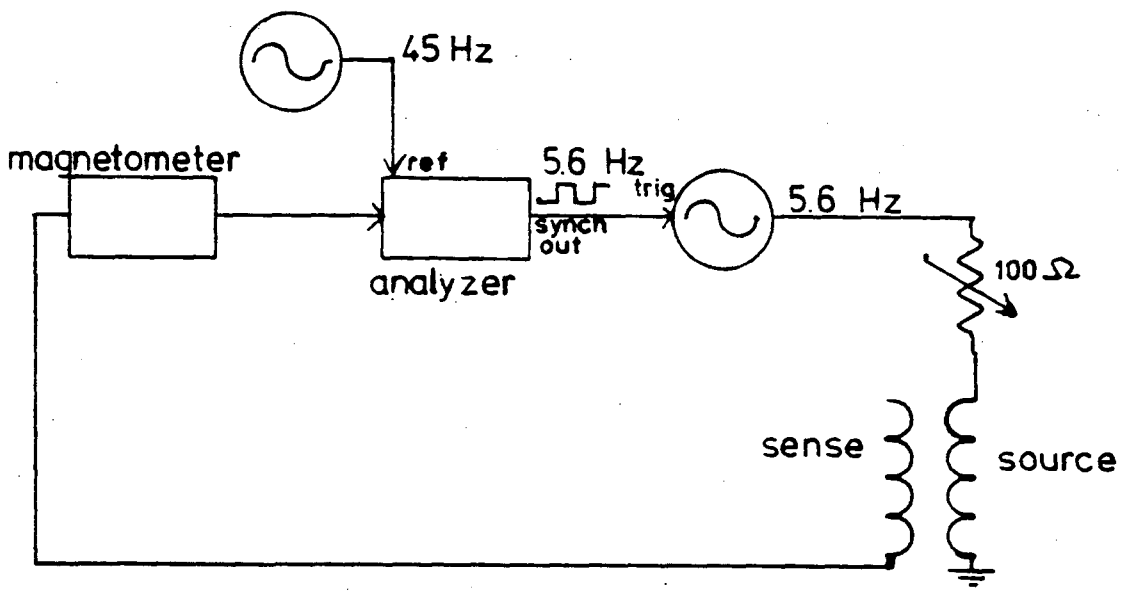

Figure 4: Synthetic Reference Signal Generation

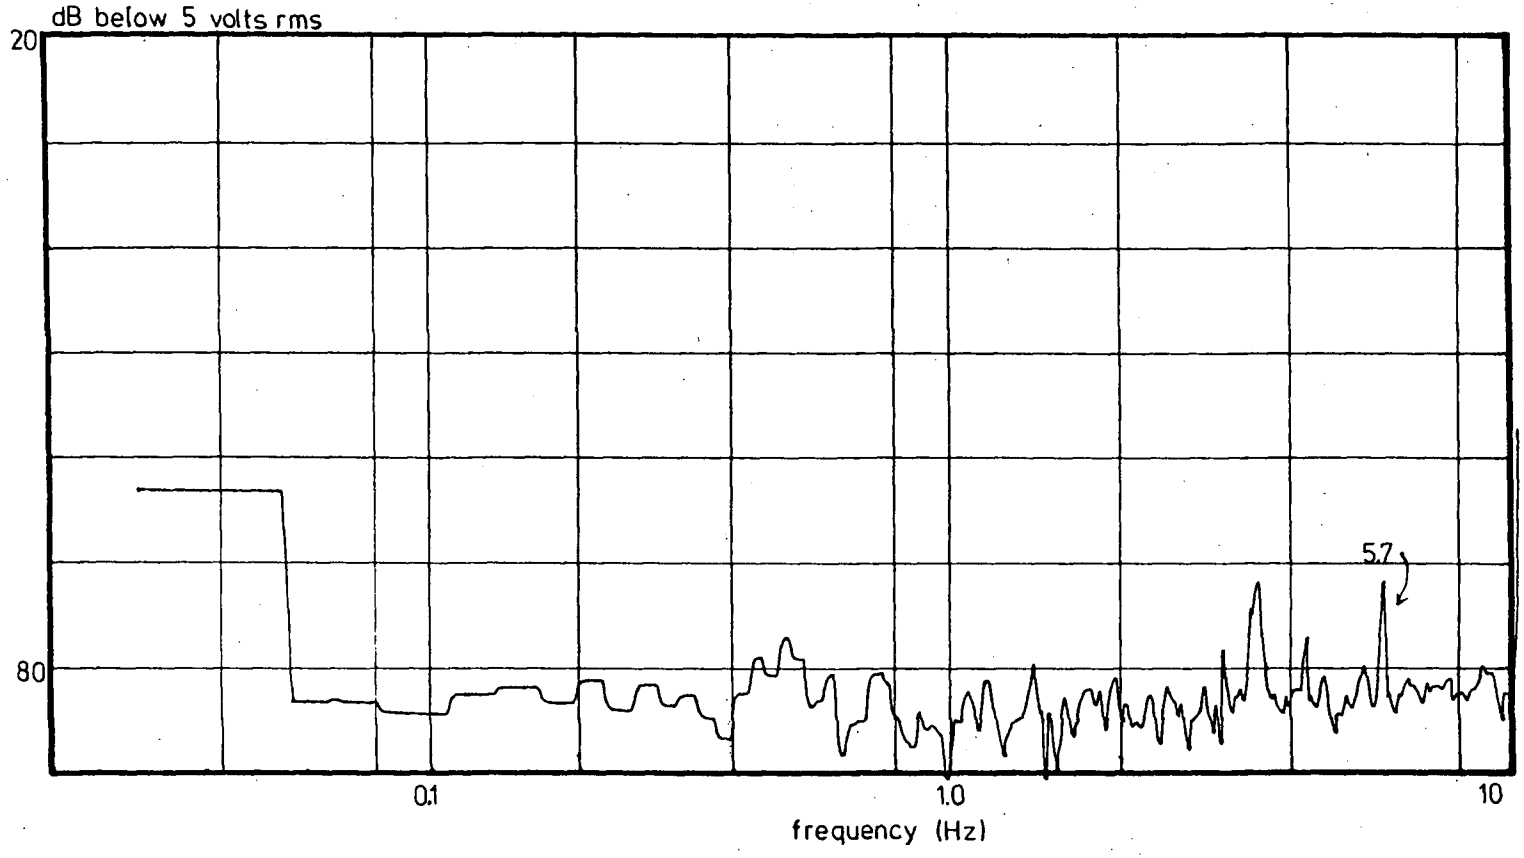

Figure 5: Flutter Test: Channel X, synthetic source.

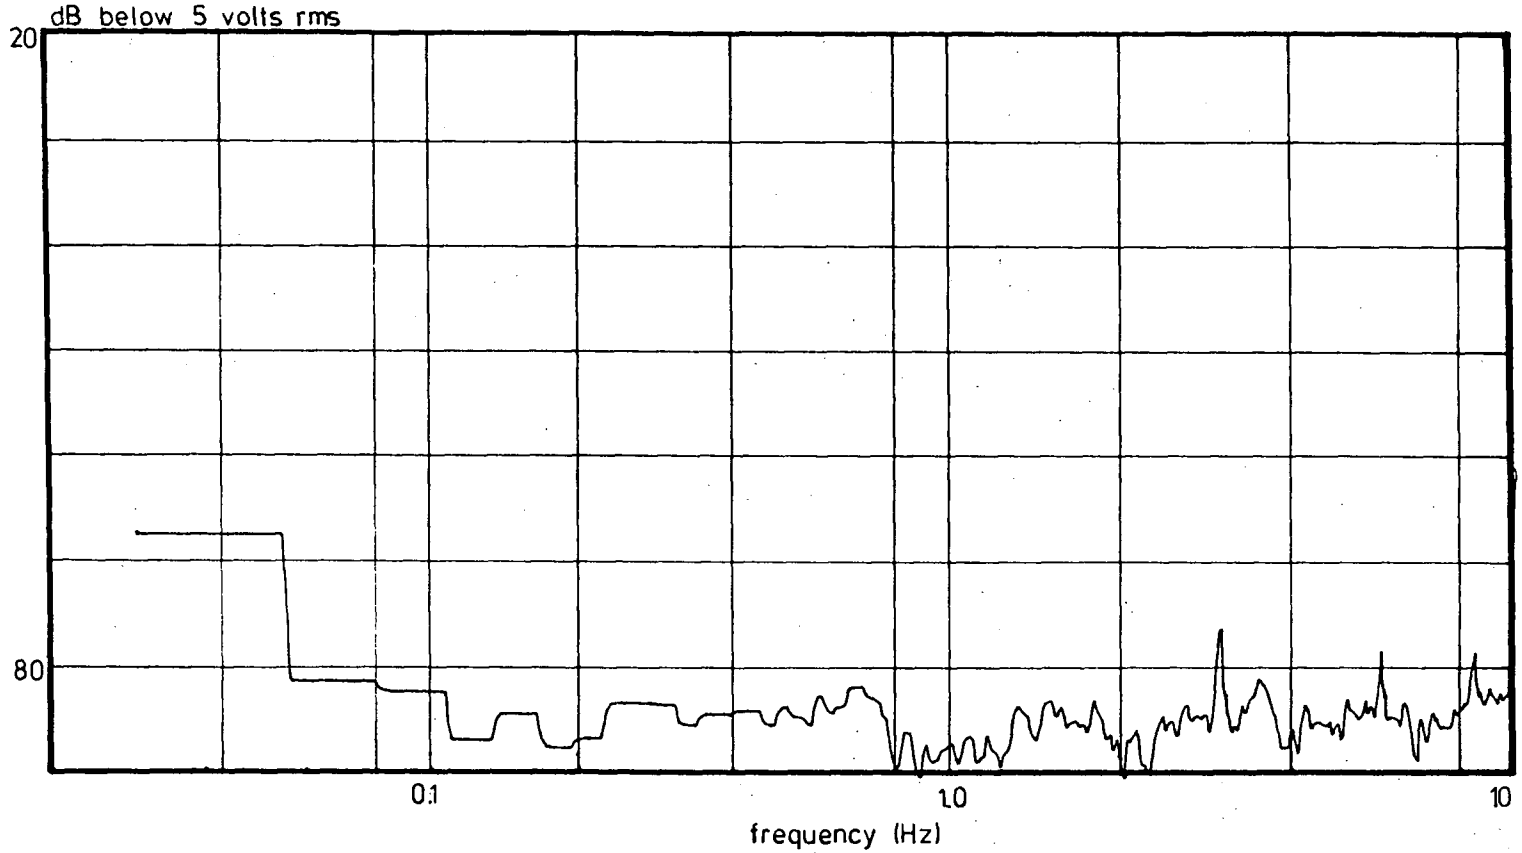

Figure 6: Amplitude dependence of Flutter: Shaft encoder case.

Comparison with Figure 3 clearly shows that the shaft encoder is a measureable noise source. Because the reference signal essentially multiplies the sensor signal, it might be expected that the noise in Figure 3 is signal-amplitude dependent. To test this, the experiments were repeated with no source, i.e. zero signal level; for the shaft encoder case (Figure 6), this meant no sample in the holder, and for the synthetic timer case (Figure 7), the source signal generator amplitude was reduced to the minimum possible. In both cases, the noise is seen to be comparable to Figure 5, indicating that flutter noise is in fact signal amplitude dependent; it only becomes noticeable when the signal is large (affecting the S/N ratio only minimally) and, therefore, there is little to gain by attempting to reduce it.

1.9 Rumble Test

In order to look for noise produced by horizontal motion of the shaft, spectra were taken of small and large sample-sensor separations. Because this noise will clearly be dependent on signal-strength, it is important to ensure that the signal **at the sensor** is the same in both tests. This was done by measuring a rock near the sensor (Figure 8), and then duplicating the processor's output with a nail far (approximately 20 cm) from the sensor (Figure 9). The spectra thus obtained are identical except that the near spacing has more energy in the second harmonic, and the far spacing more energy at the third. This discrepancy is difficult to relate to the sensor-sample distance; it is more likely due to either the different samples used, or a statistical glitch. In any case, the higher harmonics are 40 dB below the fundamental and therefore of little concern.[2]

An interesting feature of the spectra in Figures 8 and 9 are the smaller peaks approximately 1 Hz to either side of the fundamental. Because the belt period is 1 Hz, it was suspected that they were FM sidebands produced by a non-uniform belt. This was confirmed by running a finger along the belt to straighten it out, which caused the

[2] More is said about harmonics in design considerations.

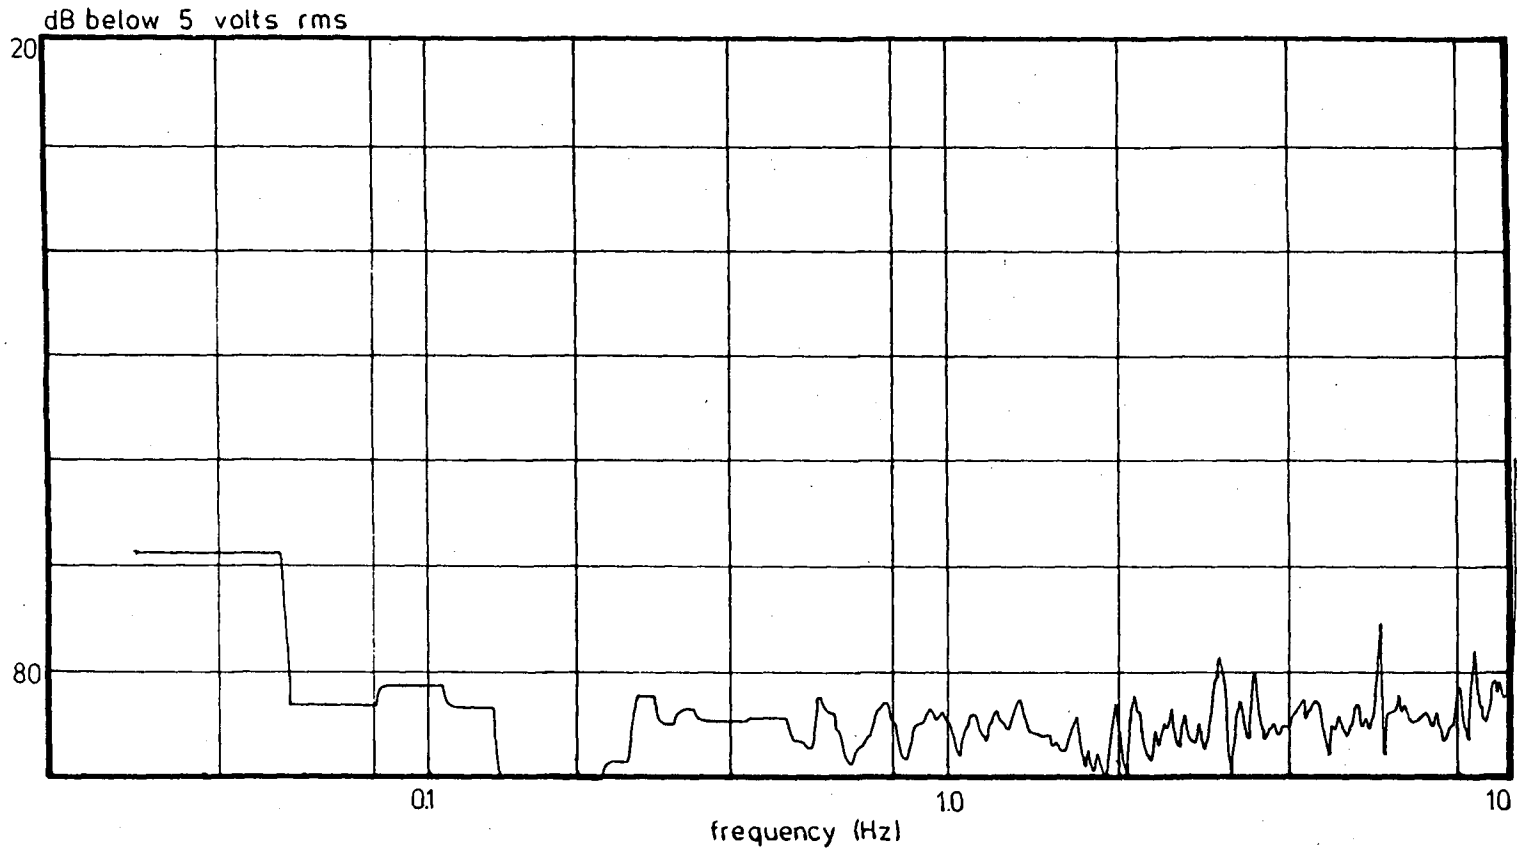

Figure 7: Amplitude Dependence of Flutter: Synthetic timer case.

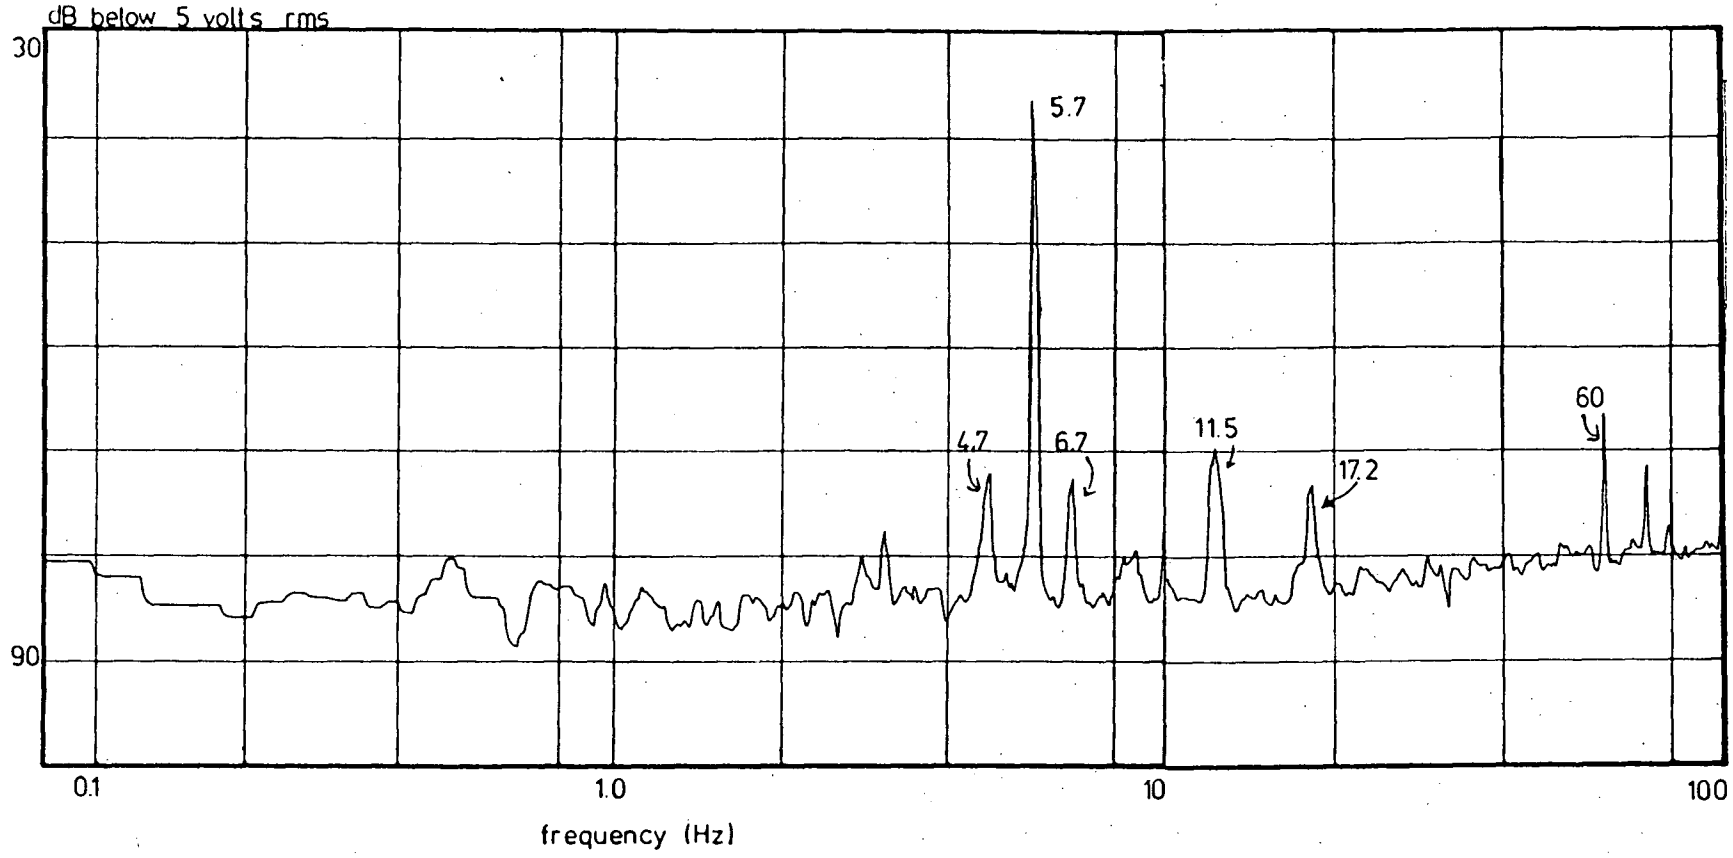

Figure 8: Rumble Test: Rock source near sensor.

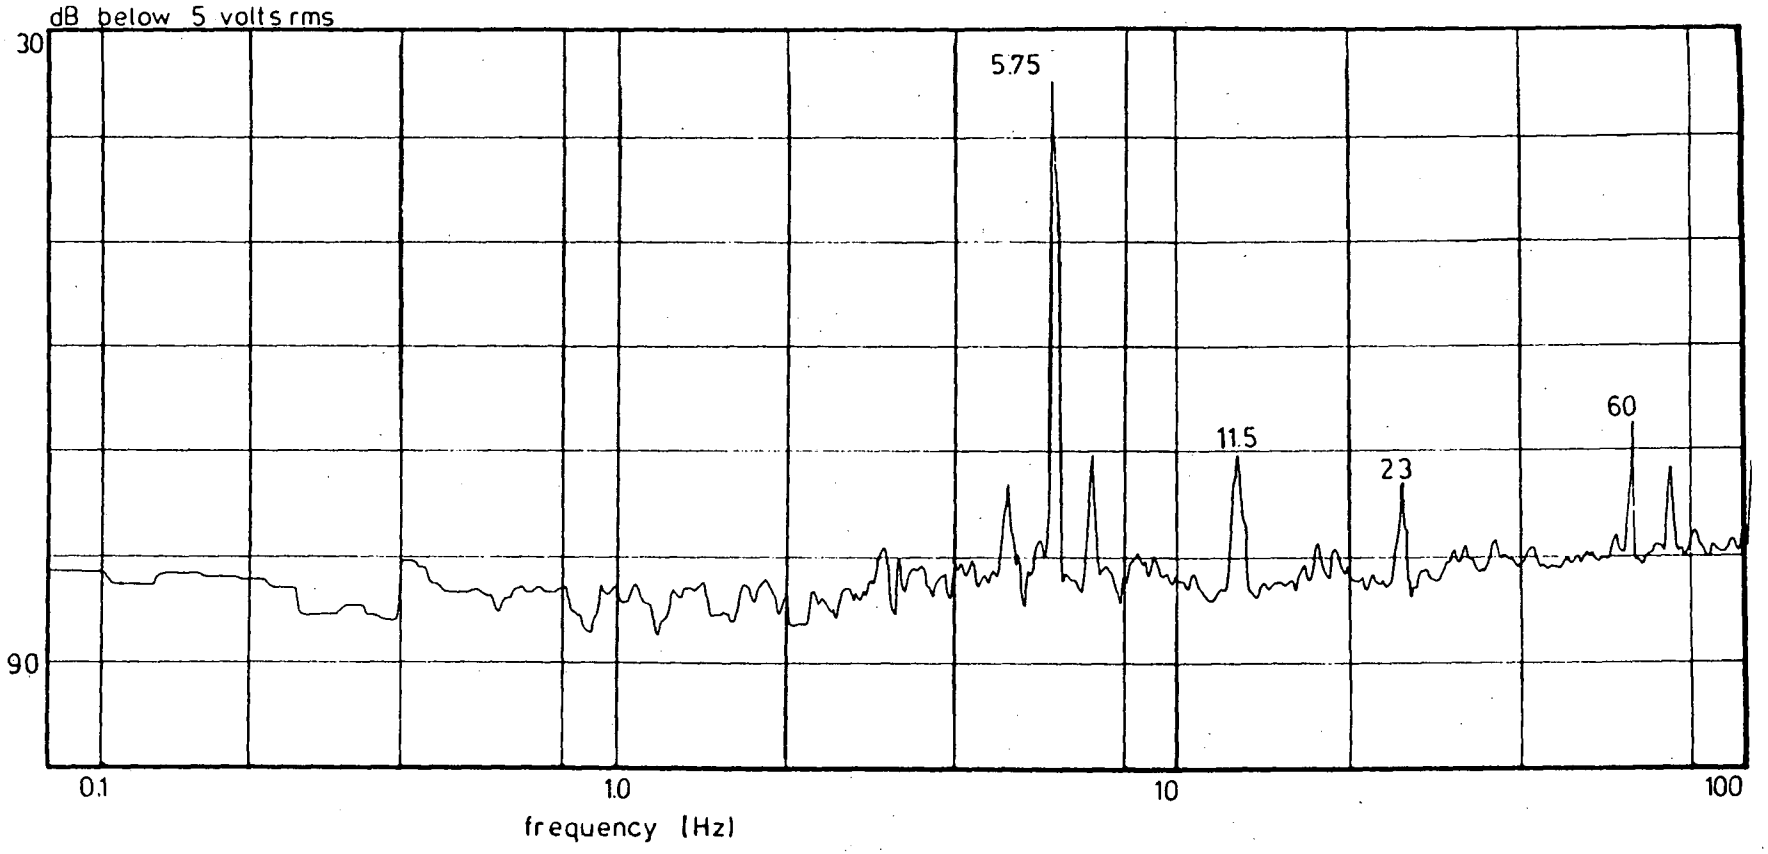

Figure 9: Rumble Test: Nail source far from sensor.

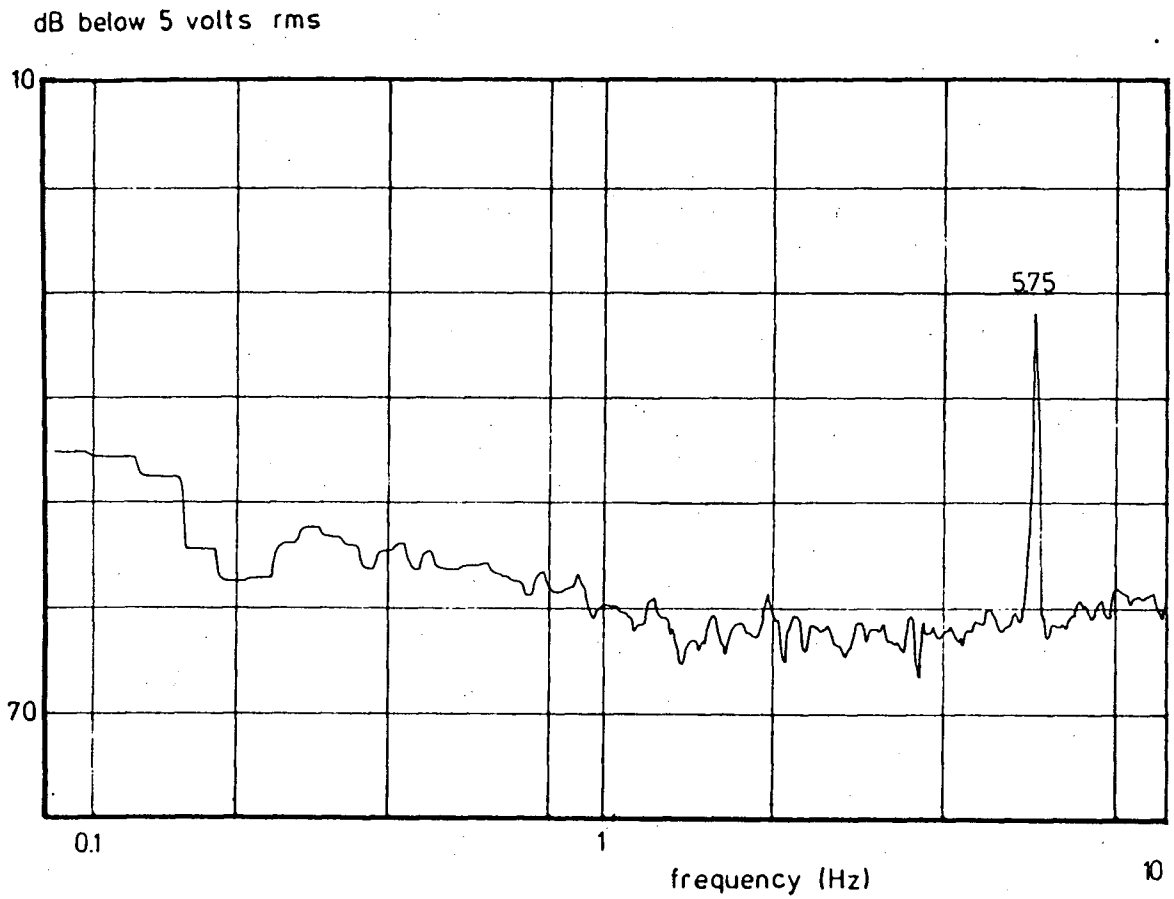

Figure 10: Disappearance of Sidebands.

sidebands to disappear (Figure 10). After a few minutes, however, the belt tends to twist up a bit and the sidebands return. The processor never sees this "wow", of course, because it is referenced directly to the shaft.

Chapter 2

Design Considerations

2.0 Introduction

Careful attention to circuit design in order to minimize noise and distortion became much more important after the discovery of the unexpected dip in the Schonstedt sensor noise curve. In this chapter the major constraints and important design options are discussed.

2.1 Sensor

The objective of installing the existing University of British Columbia magnetometer design dictated that the basic Goddard Space Flight Center flux-gate sensor geometry, with the sensitive direction perpendicular to the ring core axis, be retained. The ceramic mount for the element, however, has small protrusions on it which keeps the spinner sample holder a minimum of 3 or 4 mm. from the sense element. If the spinning holder can almost touch the sense coil, the resulting S/N, (where N=sensor noise) can be improved; consequently, sensor mount redesign was undertaken. Easy adjustment of the position of the mount along the spin axis, and its angular orientation about the spin axis for phase calibration is also important, and these features were incorporated into the design. Also, the mount must obviously be made from non-magnetic material, and acrylic plastic was chosen because it is easy to work with.

The Schonstedt instrument has the magnetometer electronics in a separate package beside the shield to minimize lead length between the sensor and the electronics. The alternative approach, used in the University of British Columbia design, is to allow long sensor leads to the main instrumentation package, but use shielded twisted pair cable to minimize inductive pick-up.

2.2 Magnetometer

The gain of the system needs to be very high, because of the weak signals to be detected. Gain stages after the magnetometer will amplify any electronic noise generated by the magnetometer, so that the general practice of concentrating as much gain as possible in the first stage of amplification needs to be followed. The University of British Columbia's Geophysical Instrumentation Group magnetometer, as described in the B.C. Science Council Report (Russell and Narod, 1982), is optimized for a feedback resistor of 33K. This was increased to 1M, for a thirty fold increase in magnetometer gain, resulting in a sensitivity of $100\text{mV}/\gamma$.

In order that the magnetometer retain approximately the same bandwidth (D.C. to 60 Hz), the integrating capacitors on the sense amplifier were changed from $0.47\mu\text{F}$ to 33nF. The final change was to remove the on-card power supply, PS 1 (the original design was for a battery-operated portable version) and make the appropriate connections for the ± 15 V rails.

2.3 D.C. Servoing

The magnetometer output to be detected is a sinusoid at the spin frequency, so it is important that the D.C. level be zero to take advantage of the full output range (± 14 V). Any but the smallest ambient field inside the shield could saturate the amplifier.

Schonstedt provides an operator-controlled nulling current through a separate coil to zero the D.C. field. Another option is to provide automatic nulling by means of a D.C. servoing feedback loop around the magnetometer coil, so that the sense coil is used to generate the nulling field. The amount of current required to null the field must still be monitored by the operator in case the shield becomes magnetized, or a magnetized object (such as a screwdriver) is left inside the shield, because there will be a limit to the field that can be nulled. This approach was chosen because it is simple and it frees the operator from continually adjusting the nulling current.

2.4 60 Hz Noise

Tests at Pacific Geoscience Centre indicated that 60 Hz noise is the peak interference frequency (see previous spectra). Because the signal requires a further three to four orders of magnitude amplification after the magnetometer output, the 60 Hz can saturate the amplifier, causing a loss of information. It would therefore be advantageous to reject the 60 Hz before further amplification, and a standard active second order notch filter was built and tested, both alone and with a second identical filter as a second stage to deepen the notch. A single second order stage was determined to be sufficient.

2.5 Analyzer

After amplification and filtering, it is necessary to detect and measure the amplitude and phase of the signal at the spin frequency, and provide D.C. output levels to the minicomputer's A/D Converter proportional to the in-phase and quadrature components. If we consider the operation as simply heterodyning the spin frequency to D.C. and then low-pass filtering, it is clear that to improve signal to noise, the corner of the low-pass filter must be reduced. However, the more severe the cutoff, the slower the transient response, and the longer it takes for the output of the filter to achieve its steady state D.C. value—a decision needs to be made on the trade-off.

Because the system at Pacific Geoscience Centre normally is run with a 12 second spin time, and the software expects at least a 12 second spin time, it was decided to take a 12 second maximum convergence time as a design constraint, which, in turn, imposes a minimum noise constraint on the system.

Once the time of sampling relative to start-up (12 seconds) is fixed, the "optimal" system time constant can be chosen. This is a question of how close to the final value the output should be after 12 seconds, which in turn depends on the amount of noise on the output value (see Appendix B).

It is possible under certain conditions to improve on the bandwidth/convergence time trade-off by changing the time constant (τ_s) as the system converges—that is, to start off with a shorter τ_s at the beginning to get the system near where it should be and then increase τ_s to reduce the noise level toward the end of the spin period. This can be broken down into two separate sequences: an "outer convergence", where the difference between the output value and its final asymptotic value is large compared to the standard deviation of the noise on the signal; and an "inner convergence" where the output is close to its final value, relative to the noise level (Bruggemann, 1980). For S/N>1, a simple outer convergence can be implemented to speed up the initial rise of the output, and thus allow a slightly greater τ_s (and hence greater S/N) at sample time, than with a single τ_s for the entire 12 seconds.

Unfortunately, using inner convergence to push the measuring ability of the instrument down to smaller signals gives only minor improvements for this system (Appendix B).

An alternative to Schonstedt's synchronous detect processor for the analyzer is an adaptive filter (Appendix C). Because the synchronous detect uses a square wave to look for a sinusoid, it brings in sensor noise at all the odd harmonics. If a reference signal closer to a sinusoid is used, some noise reduction can be achieved. Because the adaptive filter lends itself to more general reference signals, and because the lab has extensive experience with them, it was decided to use an adaptive filter as an analyzer.

Frydecke's (1980) analogue adaptive filter design was taken as a starting point. The analogue multipliers were replaced with more accurate multiplying digital-to-analogue converters and the digital reference signal applied to the DAC's was generated from a pair of UV-eraseable EPROMS. Because the strobe wheel on the spinner shaft has eight slots on its perimeter, the reference waves of choice are an eight-sample sine, and an

eight-sample cosine. [3] The sampling points were chosen to minimize harmonics, with the resulting waveform shown in Figure 11, whose first non-zero harmonic is the seventh—i.e., the third and fifth have been eliminated—at 1/7 the amplitude of the fundamental. If we assume the sensor noise is inversely proportional to frequency, then it also will be 1/7 of the value at the fundamental. Hence the sensor noise contribution due to the lowest harmonic of the waveform is 1/49 of the contribution at the spin frequency and can be ignored. In other words, Figure 11 is a practical sinusoid.

Another advantage of the adaptive filter is that (unlike the synchronous detector) spin frequency can be changed without changing circuit components. A faster spin implies operating at a lower noise level if the sensor noise is 1/f, and consequently it is advantageous to spin faster. However, the bandwidth of the instrument and the mechanical stability of the spinning assembly impose upper limits. In practice the latter is the more severe constraint. In testing, it was found that the spin rate could be tripled without causing obvious mechanical problems.

The final design choice was to return the adaptive filter's feedback directly to the sense coil, so that it became the summing node. The result is that the entire system becomes a major feedback loop, with the forward branch comprising the magnetometer electronics, the notch filter, the gain stages, and the adaptive filter up to the integrator outputs (see Appendix A). In this configuration, the adaptive filter operates as a null detector, so that in the steady state, the signal in the forward branch (and hence in all of the above mentioned hardware) is as close to zero as the noise will allow. The advantage of this is that the system becomes insensitive to phase changes imposed by components

[3] 256 sample sinusoids were tested and rejected because of the necessity of using phase-locked-loop frequency multiplication to generate the addressing pulses from the shaft encoder. It was found that the "hunting" of the PLL introduced an unacceptable amount of flutter type noise.

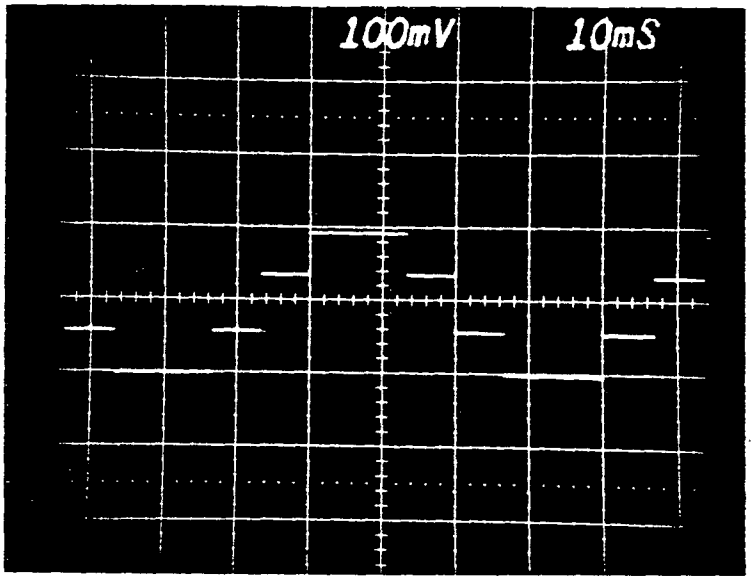

Figure 11: Eight Sample Sinusoid

in the forward branch, so that any changes in component values (either long term or temperature dependent) do not affect the measured phase of the signal. It also adds points where the system time constant can be adjusted, because τ_s now becomes gain dependent (Appendix A). The added complication of having to switch forward gain and reverse attenuation simultaneously to maintain the **same** τ_s is simply solved by using a CMOS transmission gate multiplexing chip as the switching element.

2.6 Interfacing

The motor control relay and the circuitry to indicate carriage and switch positions to the LSI 11 are contained in the Schonstedt chassis to be replaced, and therefore needed to be duplicated in the University of British Columbia instrument.

Because the LSI 11 requires TTL levels, this is the only logic hardware not implemented in CMOS.

Chapter 3

Construction and Operation

3.0 Introduction

The final design, based on the previously discussed constraints, is described in this chapter. Circuit diagrams appear in Figures 14, 15 and 16, on pages 35, 36 and 37.

3.1 Hardware Organization

The hardware is distributed on five cards and mounted in a chassis with a hinged front panel containing the gain select switch, monitoring voltmeters, power switch, fuse, and indicator light. Following is a brief description of each card, with the numbers in parentheses referring to the section number under "System Operation".

CARD 1 Magnetometer Electronics (3.2.1)

Voltage Requirements: ±15 volts and common from system power supply.

CARD 2 Intermediate Board

- magnetometer feedback resistor (3.2.1)
- D.C. servoing (3.2.2)
- 60 Hz notch filter (3.2.3)
- gain control (3.2.4)
- forward gain (3.2.5)
- feedback attennuator (3.2.6)

Voltage Requirements: ±15 volts and common from system power supply and +5 volts and 0 volts from LSI 11.

CARD 3 Analyzer Board

- shaft reference signal conditioning (3.2.7)
- system τ switch timing (3.2.8)
- sinusoidal reference signal generation (3.2.9)
- multipliers (3.2.10)
- integrators (3.2.11)

Voltage Requirements: ± 15 volts, common, +5 volts and 0 volts from system power supply.

CARD 4 Power Supply #1 and Interface Board

- power supply #1 (3.2.12)
- motor control relay (3.2.13)
- carriage position indicator (3.2.14)

Voltage Requirements: 120 volts AC; +5 volts and 0 volts from LSI 11.

CARD 5 Power Supply #2 (3.2.12)

Voltage Requirements: 120 volts AC.

Circuit common is tied to LSI 11 ground only at CARD 4. LSI 11 ground and +5 volts are also used to indicate switch positions and provide the logic levels to one quad NAND gate on the intermediate board, and one TTL quad NAND gate on the interface board. These levels are isolated from the rest of the circuit except via the power supply board, as mentioned above. If circuit common is not tied to LSI 11 ground, one NAND gate on the intermediate board (used to gate the τ_s switching signal), two NAND gates on the interface board, and the motor relay will fail to operate.

3.2 System Operation

3.2.1 Magnetometer Electronics

The sensor and magnetometer board are described in detail in the B.C. Science Council Report (Russell and Narod, 1982). The redesigned sensor mount appears in Figure 12. Other modifications include a one megohm feedback resistor, located on the intermediate board, so that if it is necessary for the magnetometer card to be operated alone, an external feedback resistor is required. On the magnetometer card, capacitors C24, C25, C27, and C28 are changed to 33nF to maintain bandwidth.

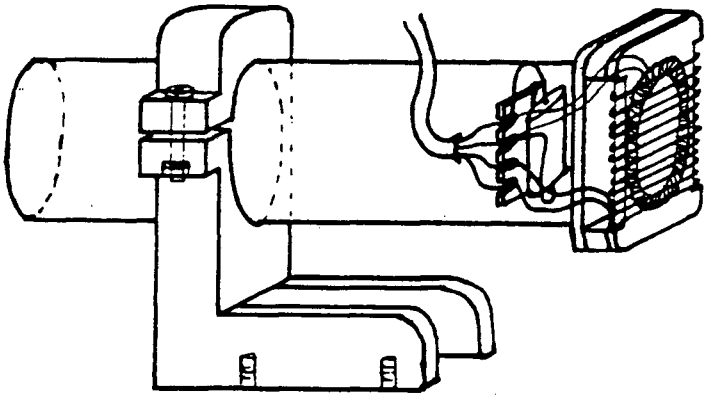

Figure 12: Sensor Mount

The drive and sense terminals of the board attach to the sensor via two banana jacks at the rear of the chassis.

3.2.2 D.C. Servoing

Automatic neutralization of any D.C. field up to 150 gamma is provided by an integrating feedback loop around the magnetometer electronics. The output of the integrator is inverted and applied to the sense coil through a one megohm resistor, in parallel with the 1M feedback resistor. The D.C. field meter on the front panel monitors the integrator output voltage required to null the field, and is calibrated in gamma. Should a

field greater than 150 gamma require neutralization (due to, for example, a bad shield) the input resistor to the inverting operational amplifier may be decreased in size. The 1M feedback and voltage-to-current resistors should not be changed, as that will affect system response.

3.2.3 60 Hz Notch Filter

Before further amplifications, any 60 Hz in the signal is removed by a two pole active notch filter, with two 100nF capacitors providing the reactive elements. The performance of this filter may be tested (with the intermediate board out and powered separately) by applying a sinusoidal signal to pin (P) and observing test point 18. The two resistors labelled SOT are iteratively trimmed to shape and locate the notch.

3.2.4 Gain Control

The front panel gain select switch creates a digital code which is read by the LSI 11 as follows:

GAIN	*U*	*V*	*W*	*X*
1000	0	1	1	1
100	1	0	1	1
10	1	1	0	1
1	1	1	1	0

Two NAND gates in a 4011 [4] translate this into a two bit code which is used to address a 4052 CMOS transmission gate multiplexing chip, which in turn switches the forward gain and reverse attenuation in unison. The address code to the 4052 is as follows:

GAIN	*A*	*B*
1000	0	0
100	0	1
10	1	0
1	1	1

[4] In this section, four digit numbers (4011, 4052, OP-07) represent electronic industry generic names of integrated circuits.

The 4052 maximum supply voltage requirement of ±9V is provided by a Zener diode network from the ±15 volt rails. Care is also taken to ensure no signals exceed ±9 volts at any inputs of the 4052 at any time.

3.2.5 Forward Gain

Following the notch filter, the signal is RC coupled to an OP-07 amplifier configured as a follower. The gain setting is set by switching one of four resistors to ground using the 4052. A reed relay can switch the feedback resistance as well, to change the gain by a factor of 9. The relay is operated via a buffer amp from a gated timing signal generated on the analyzer board. The net effect is to create a high forward gain during the first one second of spin time, unless the system gain switch is already at maximum (x1000), in which case the NAND gate prevents the high gain select signal from reaching the relay. While the closed loop gain is not affected by the forward gain, the system transient response becomes nine times faster during the first one second of spin (see outer convergence discussion in Appendix B).

Because the signal to the analyzer board must have no offset, the low-offset OP-07 amplifier was chosen; a trimpot is provided to further reduce offset if necessary.

3.2.6 Feedback Attenuator

The signal attenuation in the feedback loop directly determines system gain and is selected by the 4052 as described previously. Four decades of gain switching imply a range of 0.5nA to 5μA of current to be applied, as feedback, to the sense coil. A separate operational amplifier supplies the voltage for each gain range to the voltage-to-current resistor network, so that each one is operating in the range 0.5 volts to 5 volts. This minimizes error due to operational amplifier offsets and 4052 leakage current, which may be as great as 50nA, and also keeps the terminals of the 4052 at virtual ground at all times.

The input to the 4052 includes two trimpots: one to balance the sine and cosine error signals, and one for system gain calibration. The attenuator output is applied directly to the sense coil via the magnetometer feedback terminal.

3.2.7 Shaft Reference Signal Conditioning

Reference signals are generated by a photodiode pulling a logic "1" line to ground when a strobed light falls on it. In order to sharpen the transition and reduce sensitivity to noise, both signals are put through a hysteresis section: the 18 Hz (0^0 reference) signal through an exclusive or gate configured as an inverter with built-in hysteresis, and the 144 Hz signal through a CA 311 comparator with hysteresis by positive feedback. Two trimpots control the switching levels for the CA 311. The 50K trimpot sets the amount of hysteresis—that is, the difference between the rising and falling switching levels—and the 100K trimpot sets the reference level—the average of the rising and falling levels. These trimpots can be used to provide some phase correction, and eliminate the need for shaft encoder adjustment. One and a half to two volts of hysteresis is adequate.

3.2.8 System Time Constant Switch Timing (τ_s)

Generation of the switching signal to the reed relay on the intermediate board is accomplished by switching a logic level from high to low after the first 24 revolutions of the shaft. The circuit, located on the analyzer board, uses a monostable multivibrator (4047) to determine if the shaft is spinning (refer to the partial circuit diagram in Figure 13). As long as the 4047 is receiving the 0^0 reference pulse at 18 Hz (or somewhat slower) at its retrigger input (B), the Q output (C) is held low. A 4024 ripple counter may then count shaft revolutions by being clocked by Q3 of the addressing 4024. After 24 revolutions, Q6 (E) goes high and clocks a "0" to (F).

When the shaft stops spinning, and the 4047 stops receiving pulses on (B), line (C) goes high approximately 1/4 second after the last pulse, resetting the 4024 and 4013 and pulling (F) high.

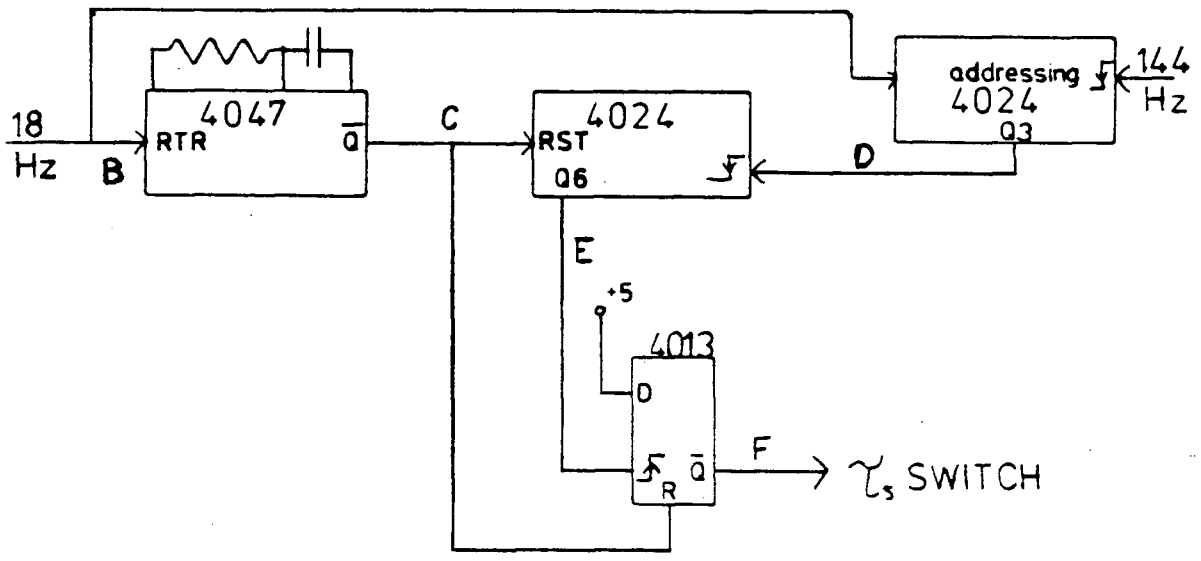

Figure 13: System τ Switch

3.2.9 Sinusoidal Reference Signal Generation

An eight-sample-per-cycle, rectified digital sine and cosine are stored in two 2716 EPROMS, which are addressed by a 4024 ripple counter. The 4024 is clocked from the 8 pulse-per-cycle (144 Hz) reference signal from the shaft encoder, so that the digital sine stays in synch with the spinning shaft; phase alignment is maintained by resetting the 4024 with the 0^0 reference (18 Hz) pulse from the shaft encoder. The 2716's allow 8 bits of accuracy for each sample, and an effective ninth bit, a sign bit, is generated from the address lines. For the in-phase or sine reference, the sign bit is supplied directly by the high order address line, and for the cosine, or quadrature reference, by an exclusive-or of the two highest order address lines.

The high order address line also provides a buffered 18 Hz reference signal for synchronizing external instruments during testing.

3.2.10 Multipliers

Each of the four multipliers consist of a DAC 1020 multiplying digital-to-analogue converter and two operational amplifiers. The first operational amplifier is contained in the low-offset LF 347 quad package, and the second is a single OP-07. To provide symmetrical bipolar operation, two 10K resistors, matched to at least 0.1% are used, with a second resistor selected on test to provide final trimming.

The digital sinusoid reference multiplies the amplified magnetometer signal, or the Fourier coefficient output signal, applied to the reference input (pin 15) of the DAC, with the result appearing at the output of the OP-15 (pin 6).

3.2.11 Integrators

An integrator with an RC product of 3 seconds using an OP-07 amplifier provides the final system output for each component (in-phase and quadrature). Multiplication of the magnetometer signal by a sinusoidal reference, followed by integration, provides the in-phase Fourier component, F_I. The output F_I is then multiplied by the reference sine to produce a correction signal C_I, which (along with the quadrature component C_Q) is subtracted from the sensor via the feedback attenuator:

$$C = C_I + C_Q = F_I \sin \omega t + F_Q \cos \omega t$$

When the current at the sensor due to C is as close as possible (in the least squares sense) to the negative of the current produced by the spinning rock, the system is in steady state, and F_I and F_Q, the integrator outputs, are measured by the LSI 11.

Once the trimming resistors on the individual multipliers are properly selected, the four multipliers and two integrators may be trimmed as a system using the trimpots on the OP-15's, by minimizing the ripple on the error signal outputs with no signal input.

3.2.12 Power Supplies

A Hamond power supply operating from 120 volts AC provides the ±15V, and +5 volts, with separate common terminals. The common for the +5 volt supply is maintained as the logic "0" level, separate from the ±15 circuit common, throughout the system. Circuit ground is tied to LSI 11 ground at this power supply.

The second power supply, on CARD 5, has the sole function of powering the shaft encoder lights, which draw almost as much current as the rest of the system.

A power switch and light are located on the front panel, and everything is protected by a 2 amp fuse.

3.2.13 Motor Control Relay

The LSI 11 "Motor On" signal operates a 12 volt relay; the relay provides 120 volts AC to the spinner motor.

3.2.14 Carriage Position Indicator

The four position switch on the carriage sets, via two NAND gates (7401), a two bit digital code, which is read by the LSI 11 on each run.

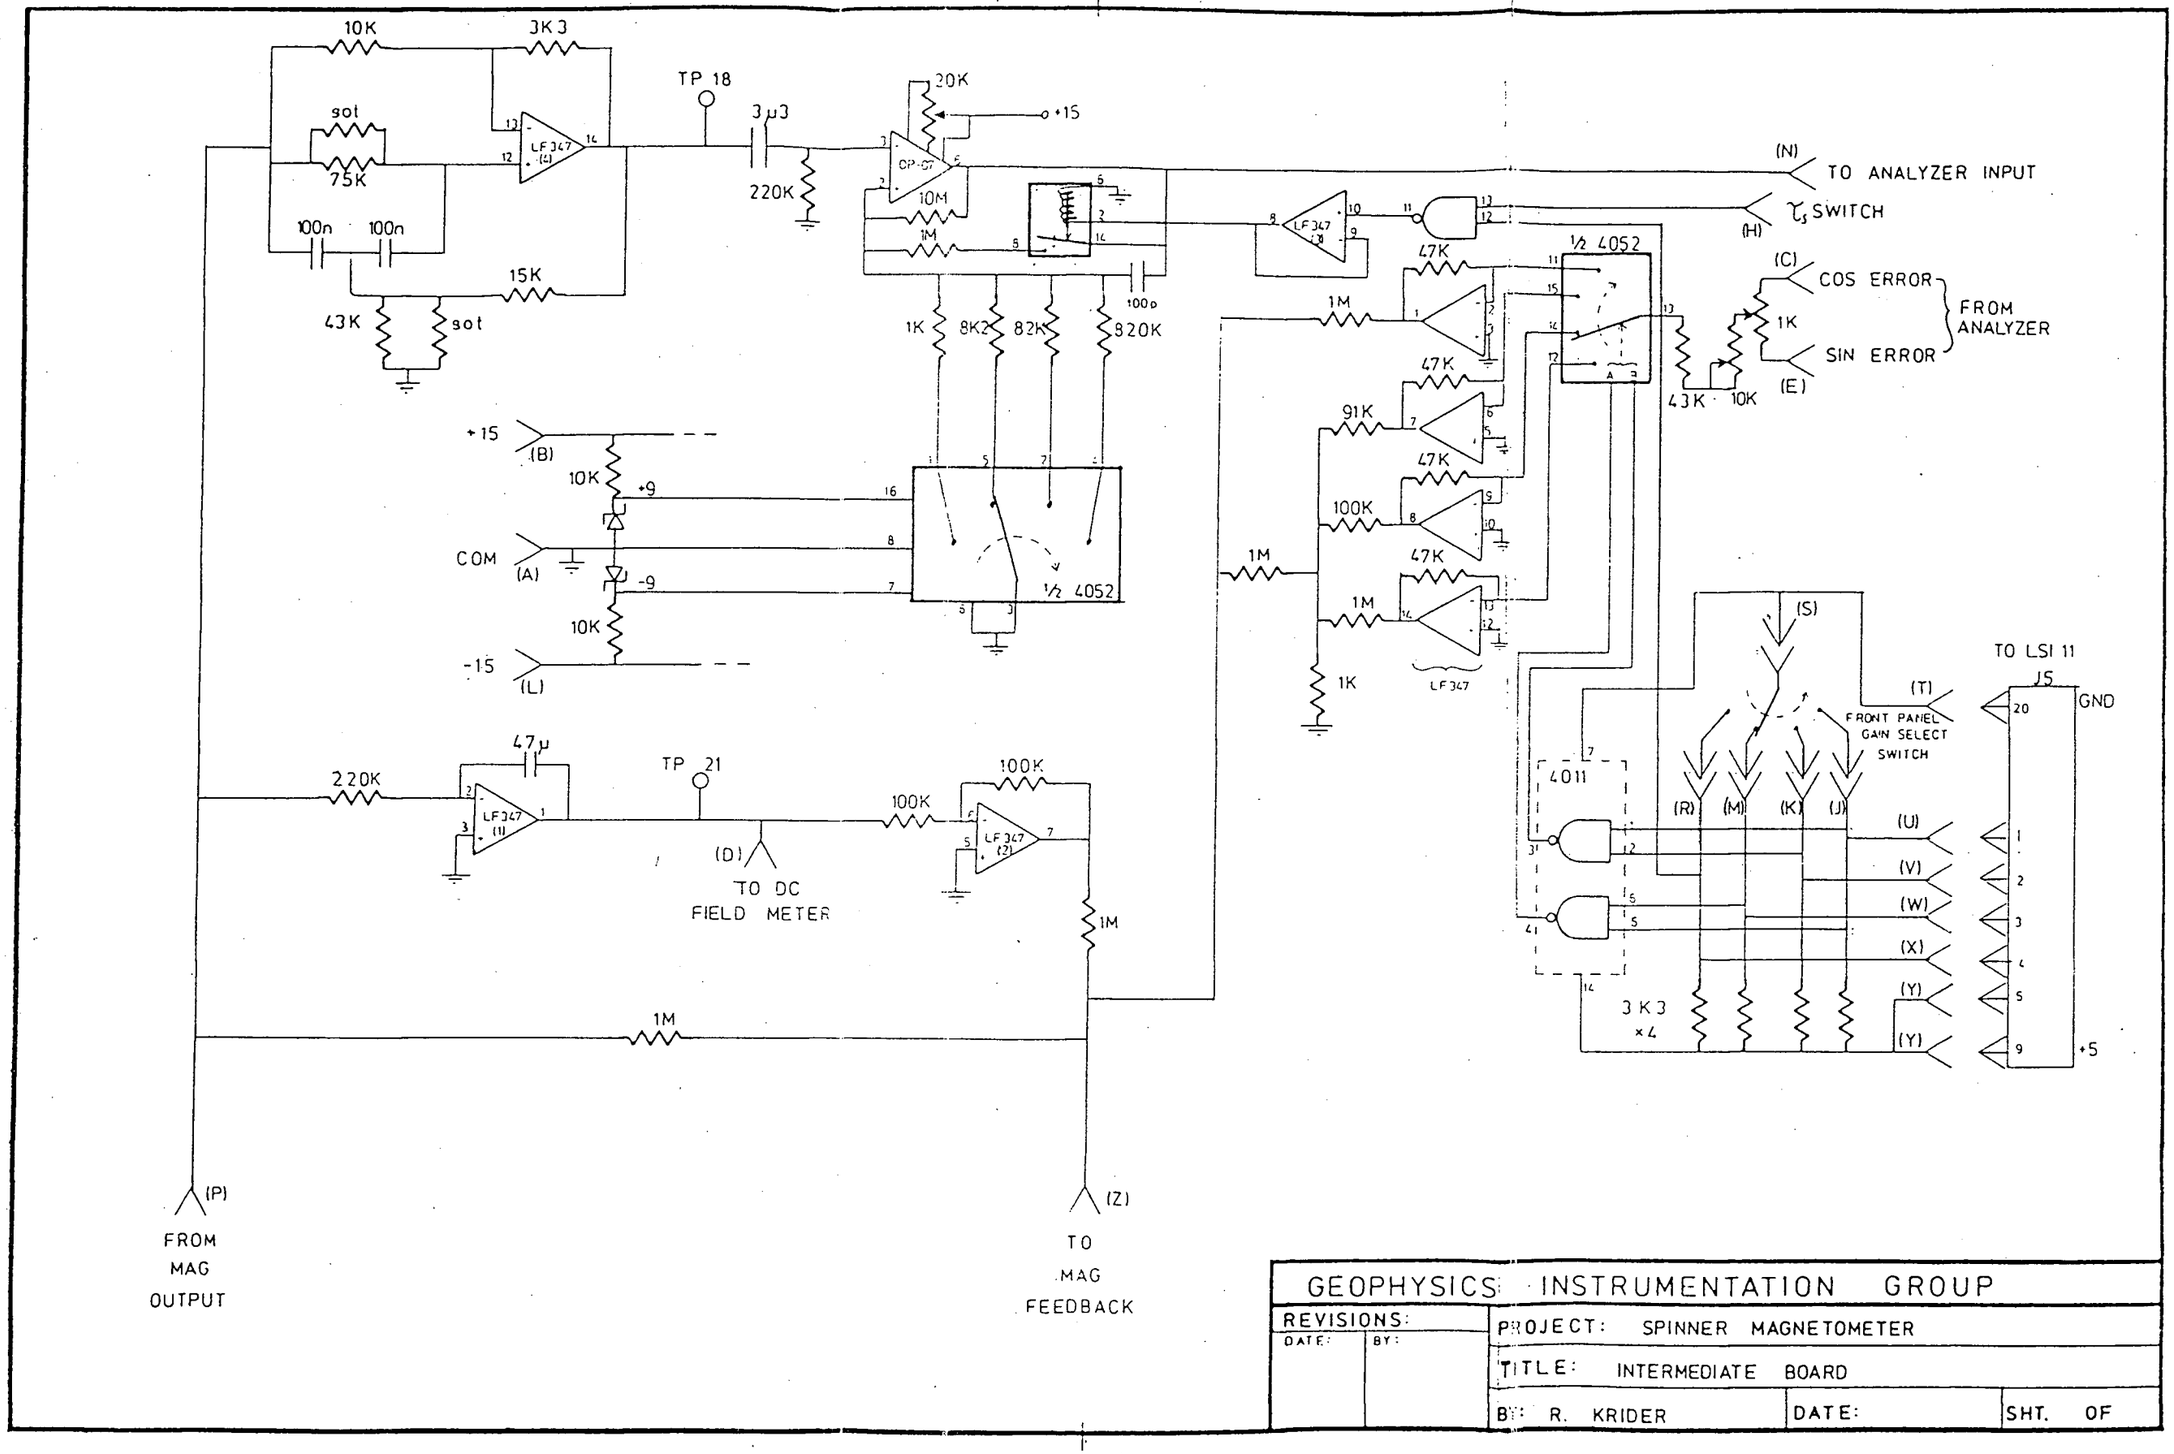

GEOPHYSICS INSTRUMENTATION GROUP			
REVISIONS: DATE:	BY:	PROJECT: SPINNER MAGNETOMETER	
		TITLE: INTERMEDIATE BOARD	
		BY: R. KRIDER	DATE: SHT. OF

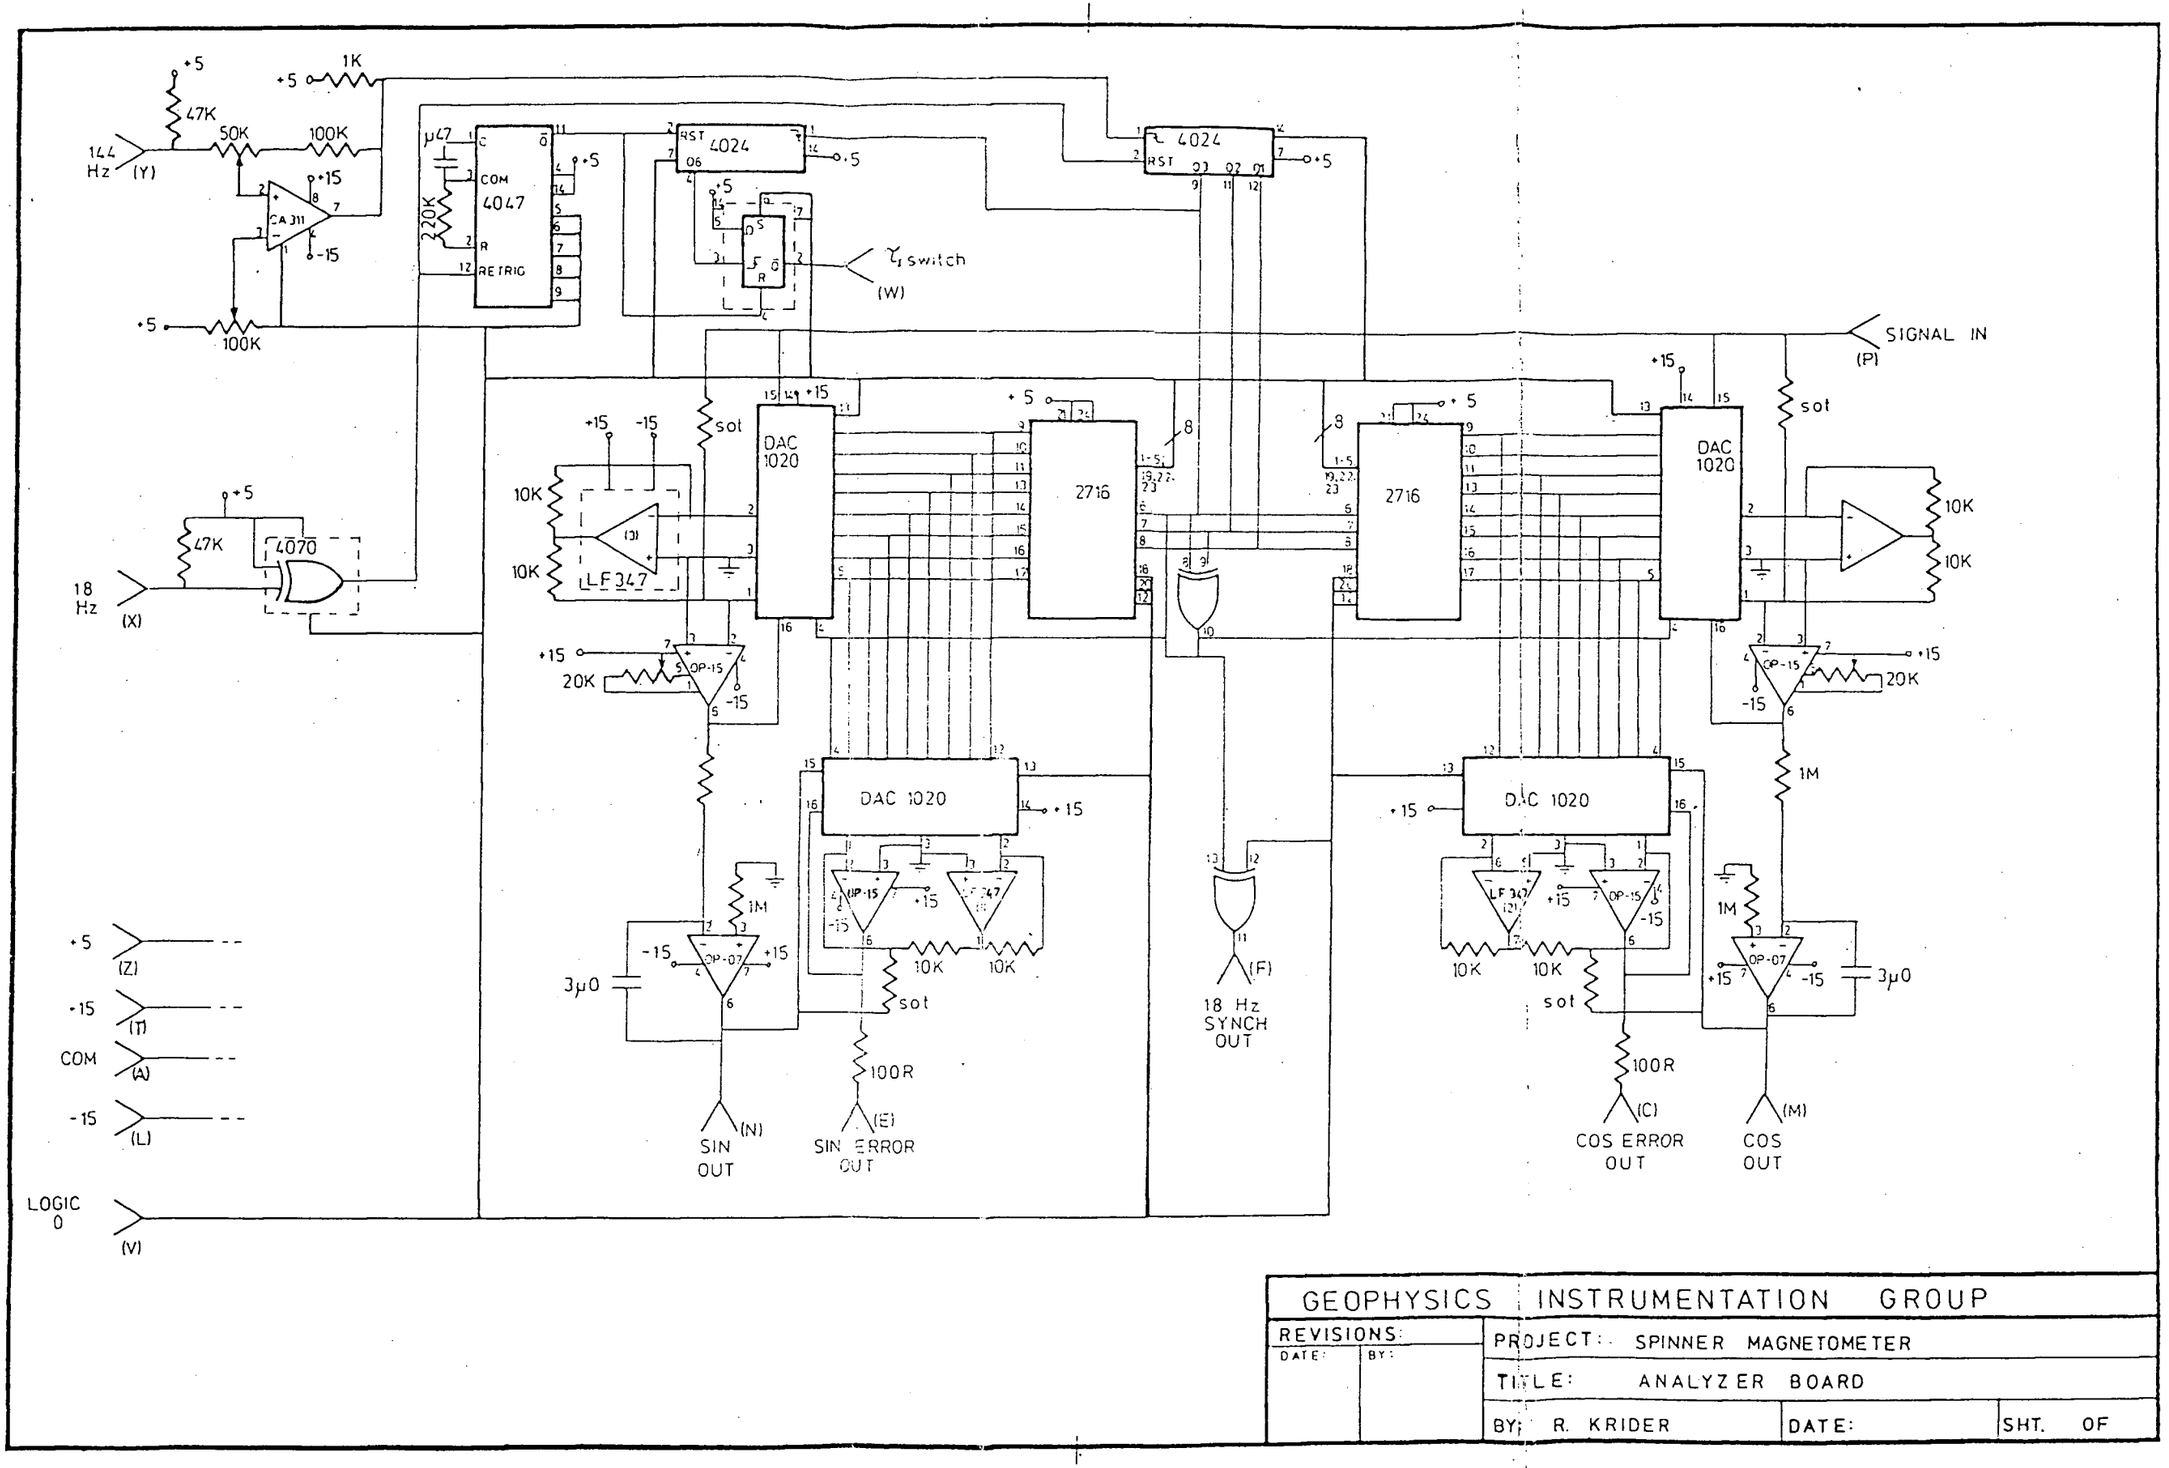

GEOPHYSICS INSTRUMENTATION GROUP		
REVISIONS: DATE:	BY:	PROJECT: SPINNER MAGNETOMETER
		TITLE: ANALYZER BOARD
		BY: R. KRIDER — DATE: — SHT. OF

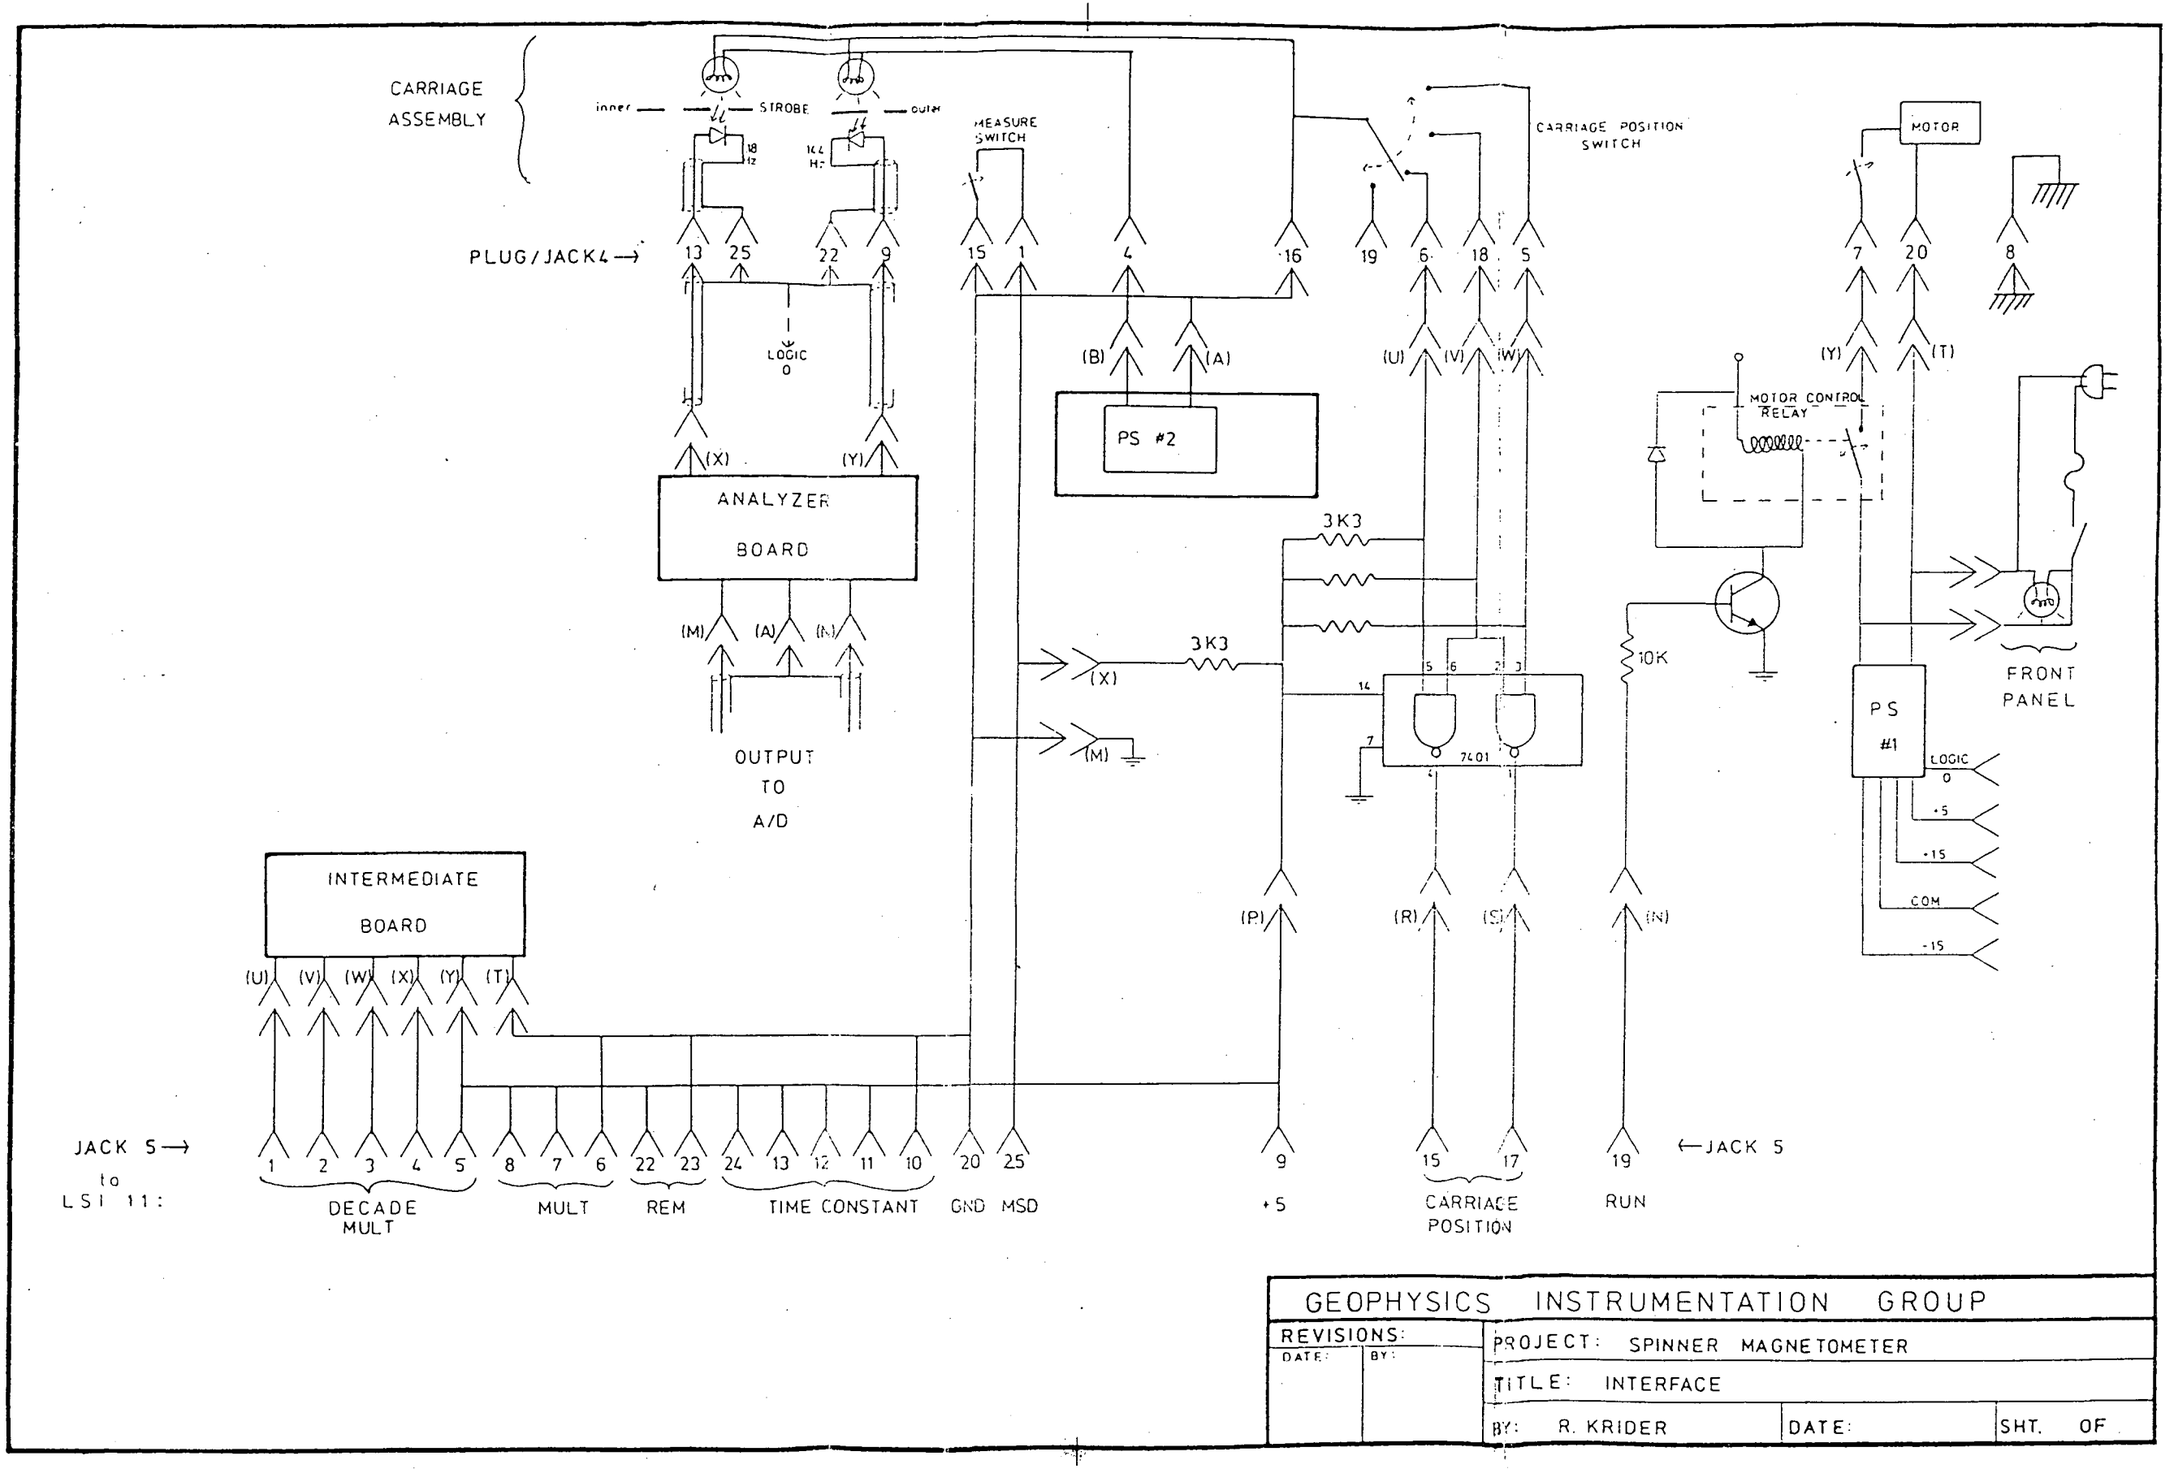

GEOPHYSICS INSTRUMENTATION GROUP			
REVISIONS: DATE:	BY:	PROJECT: SPINNER MAGNETOMETER	
		TITLE: INTERFACE	
		BY: R. KRIDER	DATE: SHT. OF

Chapter 4

Performance and Recommendations

4.1 Magnetometer

A 10dB reduction in the noise floor at the spin frequencies was achieved at the magnetometer output (Figure 17). Approximately half of this was due to increasing the spin frequency, and half from sensor noise reduction at 5.7 Hz.

4.2 Analyzer

The analyzer was unable to accurately measure the smallest signals, primarily due to the extremely small currents ($\simeq$1nA) that needed to be fed back to the sense coil.

Capacitive coupling from higher signal level portions of the circuit tended to overwhelm these feedback currents. The problem was reduced by using shielded twisted pair for the feedback lines and paying attention to physical layout, but could not be eliminated. Capacitative coupling could have been attacked by complete rearrangement of the hardware, or by returning the feedback signal to a point where the signal levels were higher.

4.3 Recommendations

The best configuration for the instrument would be to digitize the output of the magnetometer with a 16 bit ADC and implement the analyzer in software, thereby eliminating all the problems inherent in analogue circuitry at low signal levels. Furthermore, discrete signal averaging provides the best trade-off between convergence time and bandwidth, and provides the opportunity for automatic control of spin time as a function of signal strength.

Work is proceeding at University of British Columbia in understanding sensor noise mechanisms, and on the first iteration, a sensor has been produced which equals the

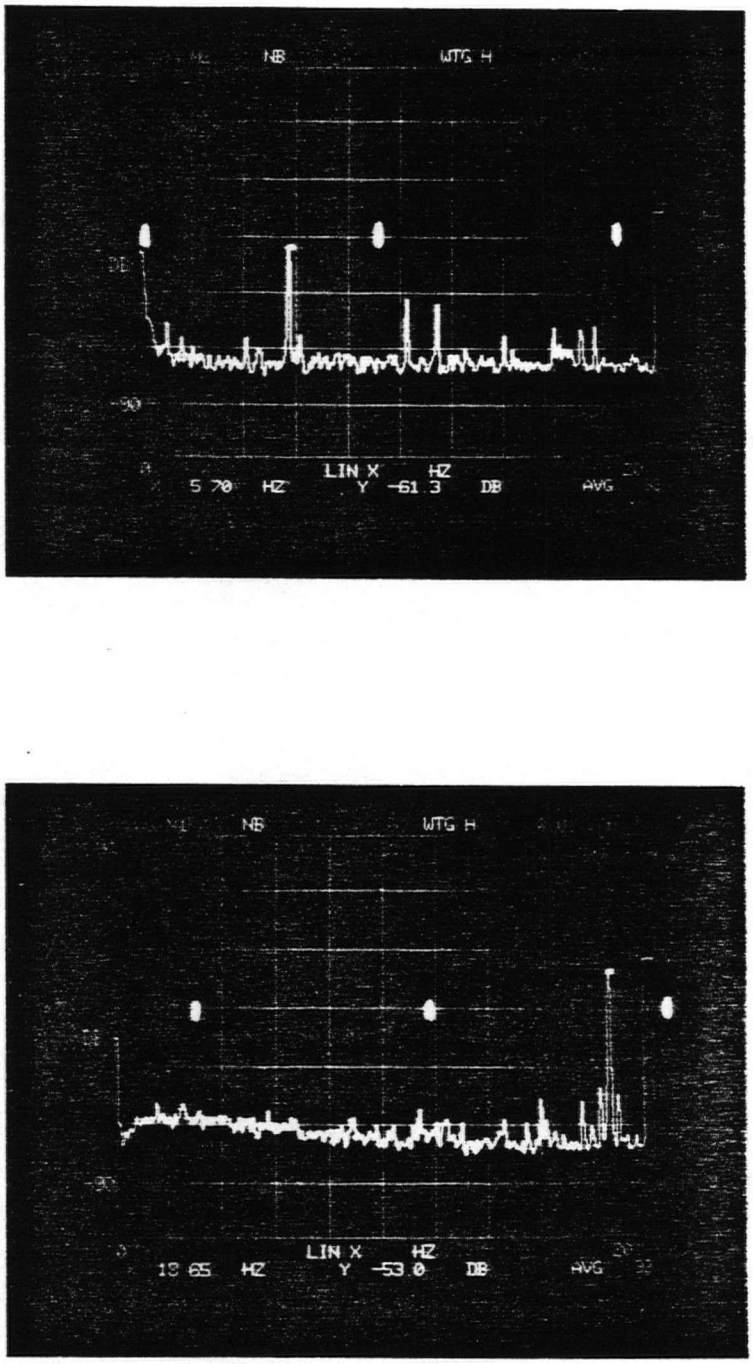

Figure 17: Magnetometer Outputs. (a) Schonstedt; (b) University of British Columbia's Geophysical Instrumentation Group. The spike is the signal from a weak rock for calibrating the two spectra, and indicates a 10dB improvement in S/N. Note that increasing the spin frequency to 18 Hz moves the signal to a lower noise region of the University of British Columbia spectrum, while the Schonstedt has its signal spike in the lowest portion of the noise spectrum already.

Infinetics core used in this project.

It is possible that a sensor with a lower noise floor at 5.7 Hz may be available in the near future. If this occurs, the spinner magnetometer should be redesigned using a software analyzer, and the software should be developed in close cooperation with the end users.

Bibliography

Bruggemann, H., **Temporal Filtering Using Pixel Incrementing**, *SMPTE Journal*, August, 1980.

Frydecky, I. I., **The Induced Polarization Receiver**, M.Sc. Thesis, University of British Columbia, 1980.

Gordon, D. I., Lundsten, R. H., Chiarodo, R. A. and Helms, H. H., **A Fluxgate Sensor of High Stability for Low Field Magnetometry**, *IEEE Trans. Magn.*, **Vol.MAG-4**, 397–401, 1968.

Russell, R. D. and Narod, B. B., **Development of an Improved Fluxgate Magnetometer**, Final Report to Science Council of B.C., 1983.

Instruction Manual SSM-1 and SSM-1A Spinner Magnetometer, Schonstedt Instrument Co., Reston, Virginia, 1972.

Widrow, B. et al., **Adaptive Noise Cancelling: Principles and Applications**, *Proc. IEEE*, **Vol.63**, No.12, 1692-1716, December, 1975.

APPENDIX A

Consider the simplified model of the system in Figure A-1.

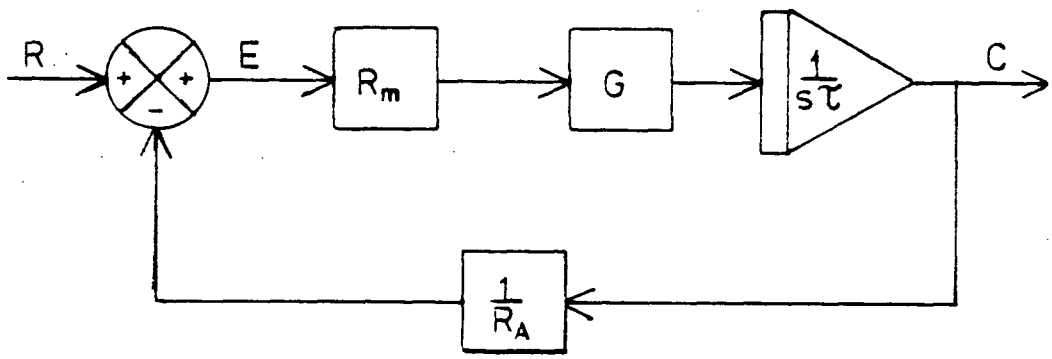

Figure A-1: Control Theory Model

The summing node represents the sense coil with magnetic field generated current input R; E is the current to the magnetometer electronics, which here is represented by the magnetometer feedback resistor R_m (the magnetometer output is a voltage with magnitude $R_m \cdot E$). Forward gain and everything else, except the integrator, is represented by G. $\frac{1}{R_A}$ is the attenuator, which performs a voltage to current conversion to be fed back to the sensor. The multipliers of the real system are used for heterodyning before the integrator and attenuator, and, as long as the multiplication is by a constant amplitude and frequency sinusoid, may be ignored for the time being (see Appendix C for adaptive filter behaviour in this configuration).

$$C = E \cdot \frac{R_m G}{s\tau_I} \qquad (A1)$$

$$E = R - \frac{C}{R_A} \tag{A2}$$

Solving,

$$C = \frac{R_A}{1 + \frac{R_A}{R_m G} \cdot s\tau_I} \cdot R \tag{A3}$$

(A3) represents a low pass filter with gain dependent only on R_A, the feedback attenuator. The system time constant τ_s is

$$\tau_s = \frac{R_A}{R_m G} \cdot \tau_I \tag{A4}$$

so that if R_A is switched to change system gain, G must be switched simultaneously to prevent τ_s from changing. Similarly, to change the system time constant without affecting gain, one may switch G only.

APPENDIX B

I Optimum Single Time Constant

In a system where an analogue integration is used to improve signal to noise, it is clear that the longer the integration time, the better the S/N will be. The output of the integrator will have an expected value which converges exponentially to the true value, at a rate determined by the time constant of the integrator. The actual output will also have random fluctuations about the expected value, determined by the input noise of the system and the bandwidth of the integrator—the narrower the bandwidth, the less input noise is allowed through (equivalently, the longer the integration time, the less the noise).

The first question that might be asked is (Figure B1), "in a given S/N environment and for a given time constant, when should the output be sampled?" One doesn't want to sample too early (before the expected value has converged to the true value) nor is there much point waiting a long time (when the random fluctuations are much greater than the difference between the expected and true values).

A reasonable arbitrary choice would be to sample when the difference between the expected and true values is the same as the standard deviation of the fluctuations. If the input noise spectrum is known the time constant will give us, via the bandwidth, the standard deviation of the fluctuations, and the sample time can be calculated.

In this thesis, a slightly different problem is posed: given a fixed sample time (twelve seconds) what is the best time constant? It should be long enough to reduce the variance, but short enough to be "reasonably well converged" within the twelve seconds. The problem is approached by setting the sampling time t_f as the time when the expected value of the output $E[F(t)]$ is one standard deviation σ away from the true value F_0, expressing σ and $E[F(T)]$ as a function of the time constant τ, and solving for τ. Recalling

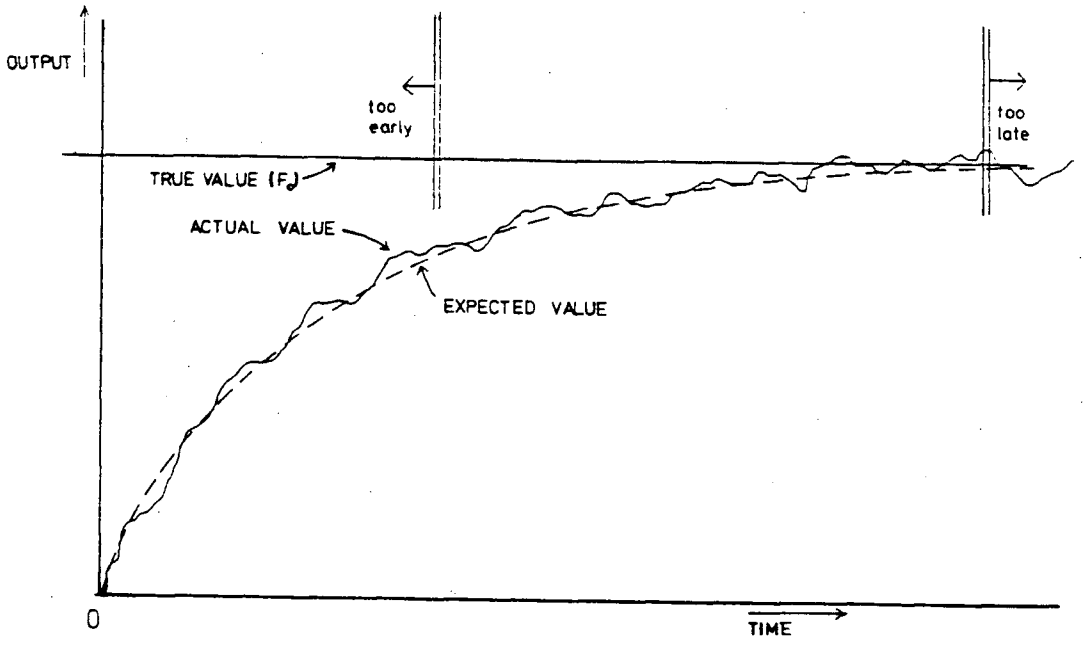

Figure B1: When to Sample? The expected value should be close to the true value, but there is nothing to be gained once their difference is much less than the noise levels.

that a spinner magnetometer heterodynes the input signal, so that the spin frequency appears at D.C., and then low pass filters the result, the time constant we are actually dealing with is the effective low pass filter system time constant, τ_s, rather than a true integration time constant.

Assume (Figure B2):

1. Rectangular input noise power density (per unit Hz) spectrum of height ν_{in}^2 and bandwidth Δf. The actual input noise is bandlimited on the low end by the D.C. servoing and on the high end by the $\frac{1}{f}$ nature of the sensor noise and the high cut of the magnetometer, so Δf should approximate this.
2. $S/N = 1$ at input. This is not only numerically convenient; it defines the transition between outer and inner convergence, discussed in Part II. If signal power is denoted

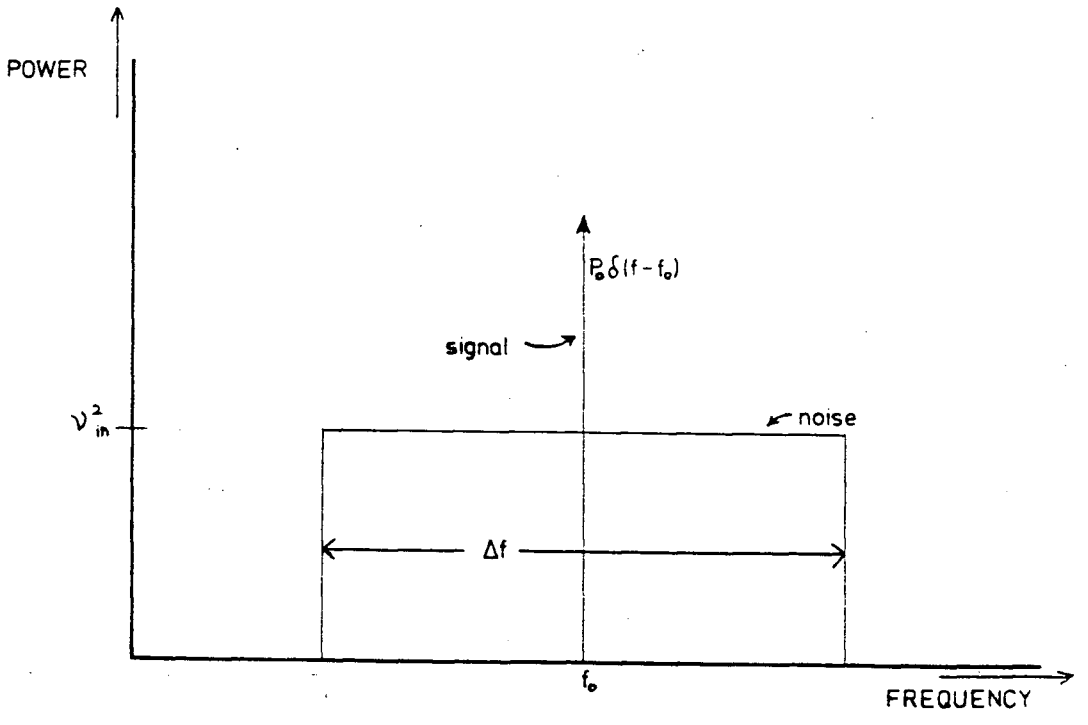

Figure B2: Input Model

as P_0 (with spectrum $P_0\delta(f - f_0)$, where f_0 is the spin frequency) we have

$$P_0 = \Delta f \nu_{in}^2 \qquad (B1)$$

3. The spin time is long enough to do some good. If the noise bandwidth is approximately the same as (or, worse, narrower than) the filter bandwidth, noise averaging won't accomplish anything. In other words, for a given noise spectrum, the filter time constant must be long enough to provide a narrow enough bandpass to cut out some of the noise. This seemingly obvious, implicit assumption becomes important in Part II of this appendix.

Under assumption 3, the noise power spectrum passed by the filter is

$$\frac{1}{1+\omega^2\tau^2} \cdot \frac{\nu_{in}^2}{2\pi} \tag{B2}$$

where the factor $\frac{1}{2\pi}$ converts density per unit Hz to density per unit radian, and the total noise power is

$$\int_{-\infty}^{+\infty} \frac{\nu_{in}^2}{1+\omega^2\tau^2} \cdot d\omega$$

$$= \frac{\nu_{in}^2}{2\pi} \cdot \frac{\pi}{\tau} = \frac{\nu_{in}^2}{2\tau} \tag{B3}$$

For a purely time domain development of this result, see Papoullis, p.436. Let $F(t)$ = the value of the output Fourier coefficient; and $E[F(t)]$ = the expected value, i.e.

$$E[F(t)] = F_i + (F_0 - F_i)\,(1 - exp(-t/\tau)) \tag{B4}$$

where F_i is the initial value, F_0 is the true value, and $F_0^2 = P_0$.

Let the standard deviation (of $F(t)$ from $E[F(t)]$) $= \sigma$; the noise power is

$$\sigma^2 = \frac{\nu_{in}^2}{2\tau} \tag{B5}$$

At the sample time, t_f, we would like

$$\sigma = F_0 - E[F(t)]$$

$$= (F_0 - F_i)(exp(-t/\tau)) \tag{B6}$$

Squaring and using (B5),

$$2(F_0 - F_i)^2 exp(-2t_f/\tau) = \frac{\nu_{in}^2}{\tau} \qquad (B7)$$

which may be solved for a given initial condition and S/N environment. If we assume $F_i = 0$ (the usual physical situation) and recall that $F_0^2 = P_0$, Equation (B1) gives

$$2\tau\Delta f = exp(2t_f/\tau) \qquad (B9)$$

To approximate the designed system, use $t_f = 12s$ and $\Delta f = 20s^{-1}$ and solve numerically for τ. This gives $\tau = 5.2s$, and

$$1 - exp(-12/6.5) = 0.9 \qquad (B10)$$

i.e., the output will be 90% converged at the sample time.

To check assumption 3, approximate the bandwidth as $\frac{1}{2\tau}$, or about 0.1Hz, which is much less than the noise bandwidth of 20Hz. In terms of the input signal power spectrum this is essentially a bandpass filter of bandwidth 0.1Hz centered on f_0, the spin frequency, which for a Δf of 20Hz substantially reduces the noise.

The 10% difference between the true value and the expected value can be taken care of by system calibration, provided we always start at $F(t) = 0$. This will be approximately true as long as a few seconds (e.g., 12) are allowed between measurements.

II Variable Time Constant

The rather unfortunate phrase "variable time constant" is meant to imply that we allow the τ_s of the system (for example the RC product of the filter) to change as the system output converges to its final value; that is, τ_s is no longer constant but a function of time. The rationale is that in the early part of the convergence sequence, τ_s can be short in order to move the output quickly into the region of the true value and then lengthened to reduce the variance, with the resulting variance at sample time lower than could be achieved with a single τ_s. The problem is to find a $\tau = \tau(t)$ which is optimum in the sense of providing the smallest variance at the sample time t.

A similar problem is addressed by Bruggemann (1981) in signal enhancement of digital video signals. Although his approach is not applicable here (due to the analogue signal and constrained sampling time) the concept of inner and outer convergence is roughly parallel. When the S/N ratio is much greater than one, we may very rapidly move towards the true value; once the output is within the noise levels of the true value some sort of averaging is necessary to further reduce noise, and hence a slower convergence is necessary. These are referred to as outer and inner convergence respectively.

The implementation of the outer convergence is straightforward—the output simply rises towards the true value as quickly as possible. In a digital system this would be a sample and hold; in an analogue system, as close to unfiltered (or $\tau = 0$, or infinite bandwidth) as possible. The end of outer convergence occurs when the output is, in some statistical sense, within the noise level of the true value, and the inner convergence algorithm must take over. Of course, if the initial S/N is less than 1, there is no outer convergence.

At the start of inner convergence, let us make the reasonable assumption that our noisy signal is now centered on the true value, F_0. We then have a stationary stochastic signal with mean F_0 input to a time varying linear system. The problem is to find what

time variation (in terms of $\tau(t)$) produces the greatest noise reduction within a given fixed interval, $[0, t_f]$.

The approach taken in the following is to partition the time interval, with $\tau = \tau_i$ a constant within each subinterval Δt_i, so that the linear filter is now time invariant within each subinterval. The price paid for this is that the stochastic input is no longer stationary; this, however, will be easy to deal with.

Formulate the problem as follows:

1. Partition a fixed time interval $[0, t_f]$ into a series of subintervals Δt_i,

$$\Delta t_i = t_i - t_{i-1}, \qquad i = 1, f \text{ and } t_0 = 0$$

2. Let L be a first order time varying low pass filter whose time-dependent characteristics on the interval $[0, t_f]$ are described by $\tau(t)$,

$$\tau(t) = \tau_i, \qquad t_{i-1} \leq t < t_i$$

where each τ_i is a constant. We consider L to consist of a series of filters L_i, each of which is turned on at t_{i-1}, is time-invariant for δt_i, and is then turned off at t_i.

3. Let the input to L be a stationary, Gaussian bandlimited noise signal with variance

$$\sigma_0^2 = \Delta f \nu_{in}^2,$$

and mean F_0, with notation as described in Part I.

4. Let the output be $n(\zeta, t)$, where ζ is the parameter describing which member of "outcome space" we have at each t, and serves to remind us that $n(\zeta, t_i)$ is a random

Gaussian variable at each t_i. Note that $n(\zeta, 0)$ is precisely the input signal; that is, we turn on L at $t = 0$:

$$E[n(\zeta, 0)] = F_0$$

$$E\left[(n(\zeta, 0) - F_0)^2\right] = \sigma_0^2$$

To analyze what happens in each subinterval Δt_i, we start with initial conditions

$$n_{j,i-1} \equiv n(\zeta_j, t_{i-1})$$

where the statistics of $n_{j,i-1}$ have been determined by the preceding intervals (if any), and turn on L_i. We then find the statistics of $n_{j,i}$ at the end of the subinterval, t_i, and iterate on through the interval.

Let us momentarily confine our attention to a particular j at time t_{i-1}, and follow its progress. We may then generalize by integrating over all possible $j's$, whose probability distribution is a known Gaussian. Let us take, for example, the point one sigma away from the mean:

$$n_{s,i-1} = F_0 + \sigma_{i-1} \tag{B9}$$

When we turn on L_i, $n_{s,t}$ will tend towards F_0 at a rate governed by τ_i. However, even though we have started at a predetermined value of n, as soon as the signal leaves t_{i-1}, it becomes indeterminate, because of the noise; it will tend to spread statistically from its "expected" exponential trajectory, at a rate determined by the amount of noise L_i allows through.

We may equivalently consider L_i' to be a time invariant filter (always turned on) and split $n_{s,t}$ into a deterministic signal and a stochastic non-stationary signal as follows:

$$\begin{aligned} n_{s,t} &= A(t) + B(t) \\ A(t) &= \sigma_{i-1}\,(U(t_{i-1}) - 1) + F_0 \\ B(t) &= 0, \qquad t < t_{i-1} \\ &= n_0(\zeta, t), \qquad t \geq t_{i-1} \end{aligned} \tag{B10}$$

where $n_0(\zeta, t)$ is the original input white noise signal with autocorrelation $\nu_{in}^2 \delta(\tau)$.

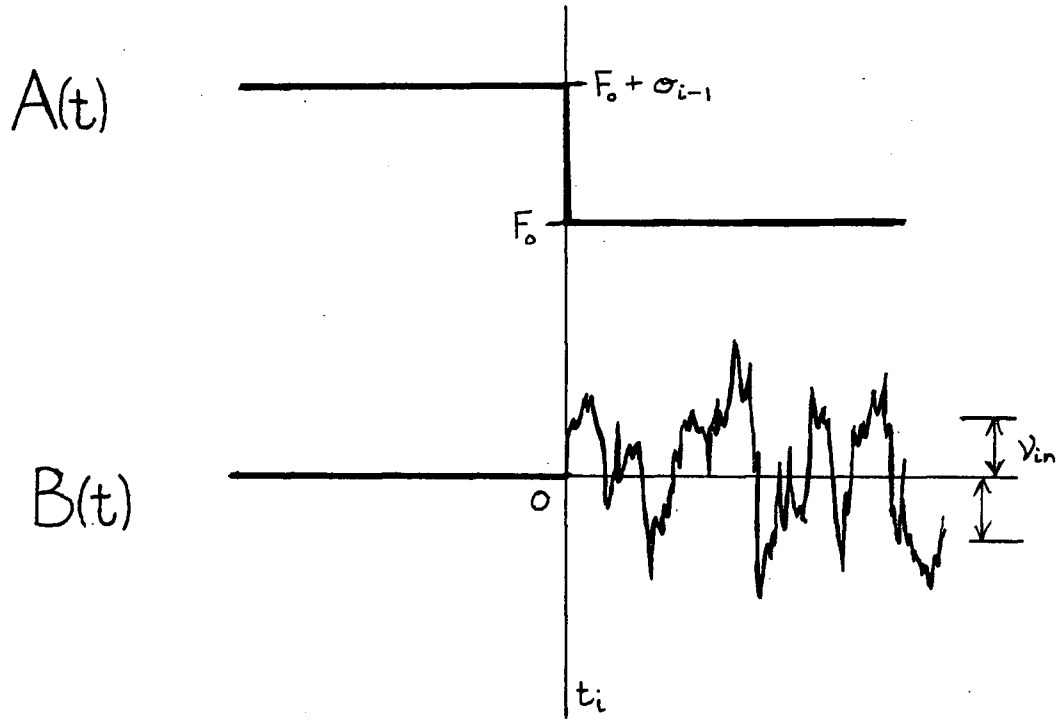

Figure B3: Noise Resolution of $n_{s,t}$ into $A(t) + B(t)$.

We now find the response of L_i' to $n_{s,t}$.

$$L_i'[n_{s,t}] = L_i'[A(t)] + L_i'[B(t)] \tag{B11}$$

The response to $A(t)$ gives the "expected trajectory":

$$L_i'[A(t)] = \sigma_{i-1} exp(-(t - t_{i-1})/\tau_i) + F_0$$
$$= \sigma_{i-1} exp(-\Delta t_i/\tau_i) + F_0, \quad \text{at} \quad t = t_i \qquad (B12)$$

The response to $B(t)$ gives the statistical spreading, and is given by Papoulis (1965, Ch.12) in terms of the mean

$$E\,[L_i'[B(t)]] = L_i'\,[E[B(t)]] = 0 \qquad (B13)$$

and the autocorrelation

$$R_{vv}(t,t') = \frac{\nu_{in}^2}{2\tau_i}\left(1 - exp(\frac{-2t'}{\tau_i})\right) exp(t' - t) \qquad (B14)$$

The variance is $R_{vv}(t,t)$, and, at the end of the subinterval is

$$v^2(t_i) = v_i^2 = \frac{\nu_{in}^2}{2\tau_i}\,(1 - exp(-2\Delta t_i/\tau_i)) \qquad (B15)$$

These results hold for any starting point at time t_{i-1}; that is, for any ς_j, we may break $n(\varsigma_j, t_{i-1})$ into $A_j(t)$ and $B_j(t)$. Now, all the B_j's are the same; and all the A_j's will map the normally distributed random variable

$$n(\varsigma, t_{i-1})$$

onto a new compressed normally distributed random variable

$$n'(\zeta, t_i)$$

which will have mean F_0 and variance

$$\sigma_i'^2 = \sigma_{i-1}^2 \; exp(-2\Delta t/\tau_i) \qquad (B16)$$

from (B12). The situation for the first interval is depicted graphically in Figure (B4), with three specific starting points.

Hence, the probability density at t_i will be the convolution of the "compressed" Gaussian due to $A_j(t)$ with the "spreading" Gaussian due to $B_j(t)$.

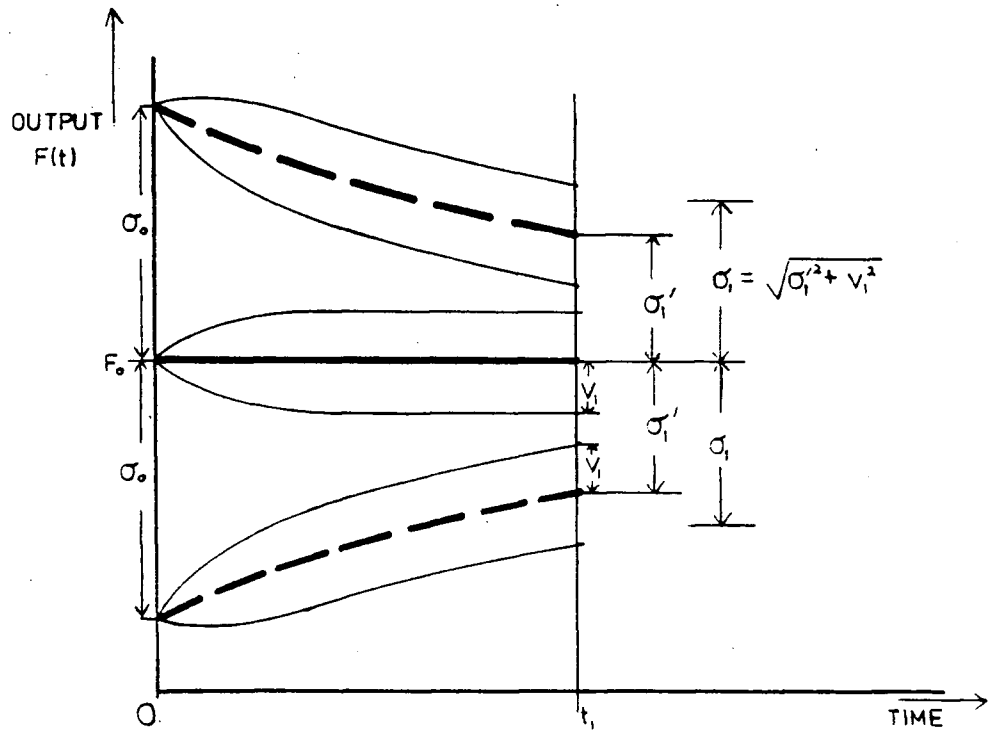

Figure B4: Noise Reduction Model

The convolution of two Gaussians of variance $\sigma_i'^2$ and v_i^2 is a Gaussian of variance

$$\sigma_i^2 = \sigma_i'^2 + v_i^2 \qquad (B17)$$

We want to make σ_i small, which implies that we want to make both σ_i' and v_i small. But to make σ_i' small requires a rapid convergence of the signal, and hence a small τ_i. To make v_i small requires a narrow bandwidth, and hence a large τ_i. We have now returned to the same tradeoff mentioned in Part I, only now described somewhat more tidily as the sum of two variances which depend on τ_i in opposing ways.

Given a partitioning t_i of the interval $[0, t_f]$, we may start with input noise

$$\sigma_0^2 = \Delta f \nu_{in}^2 \qquad (B18)$$

and iterate through the interval, minimizing σ_i^2 to find τ_i in each subinterval, to find the final $\sigma_f = \sigma(t_f)$.

To simplify the notation, let us take $F_0 = 0$ and expand (B17) using (B12) and (B15):

$$\sigma_i^2 = \sigma_{i-1}^2 exp(-2\Delta t_i/\tau_i) + \frac{\nu_{in}^2}{2\tau_i}\left(1 - exp(-2\Delta t_i/\tau_i)\right) \qquad (B19)$$

To further simplify the notation note that if we use (B18), ν_{in}^2 factors out of the right hand side of (B19). Since we will be concerned with relative variances, let us normalize σ_i^2 to ν_{in}^2. Rather than create new notation, set $\nu_{in}^2 = 1$, and remember the final σ_f is normalized.

Setting

$$\frac{d\sigma_i^2}{d\tau_i} = 0$$

gives

$$\frac{1}{2\tau_i^2}\left[-1+exp(-2\Delta t_i/\tau_i)\left(1-\frac{2\Delta t_i}{\tau_i}+4\Delta t_i\sigma_{i-1}^2\right)\right]=0 \qquad (B20)$$

We are not interested in the case $\tau_i \to \pm\infty$, so we want to solve

$$exp(-2\Delta t_i/\tau_i)\left(1-\frac{2\Delta t_i}{\tau_i}+4\Delta t_i\sigma_{i-1}^2\right)=1 \qquad (B21)$$

for τ_i in each subinterval.

Given a partitioning $\vec{t}$ of the interval, the best time constants $\vec{\tau}$ may be calculated according to (B21) and the final variance σ_f^2 found. It is then possible to do a numerical search to find the best **partitioning** to minimize σ_f as a function of $\vec{t}$.

Numerical Calculations (Newton Method)

The partial derivatives $\frac{\partial\sigma_f}{\partial t_i}$ are calculated by varying each of the input t_i's by a small amount ∂t_i and recalculating σ_f to get $\delta\sigma_f$:

$$\frac{\partial\sigma_f}{\partial t_i}\simeq\frac{\delta\sigma_f}{\delta t_i}$$

Now each $\frac{\partial\sigma_f}{\partial t_i}$ is a function of all the t_i's:

$$\frac{\partial\sigma_f}{\partial t_i}=f_i(t_1,\ldots,t_f)=f_i(\vec{t}) \qquad (B22)$$

Let $\vec{t_g}$ be the guessed input partitioning and expand:

$$f_i(\vec{t})=f_i(\vec{t_g})+f_i'(\vec{t_g})\cdot\Delta\vec{t}+\ldots \qquad (B23)$$

Since we want to minimize σ_f, we have $f_i(\vec{t}) = 0$, and to two terms, (B23) becomes

$$f_i(\vec{t}_g) \simeq -f_i'(\vec{t}_g) \cdot \Delta\vec{t}, \quad i = 1, \ldots, f \tag{B24}$$

which is f equations in f variables,

$$\vec{y} = A\vec{x} \tag{B25}$$

where A is the matrix of second partial derivatives,

$$\frac{\partial^2 \sigma_f}{\partial t_i \partial t_j}, \quad i, j = 1, \ldots, f$$

The matrix A is evaluated numerically and (B25) is solved by singular value decomposition in order to keep track of the eigenvalues to ensure a stable inversion. Should the convergence fail, it is useful to know whether or not it is because of an unstable matrix. The resulting $\Delta\vec{t}$ (the $\vec{x}$ of B25) is used to refine the input partitioning and the process repeated until the optimum partitioning is found.

The input variables are the total spin time, t_f, the number of partitions and the input noise bandwidth. The time constants $\vec{\tau}$, partitioning $\vec{t}$, and standard deviations $\vec{\sigma}$ are output at each iteration. A conservative criterium for completion of convergence is

$$\frac{\partial \sigma_f}{\partial t_i} < 0.00001$$

The effective bandwidth associated with each τ_i is $\frac{1}{2\tau_i}$. In order for the bandpass filter to do any good in an interval, its bandwidth must be less than the input bandwidth, Δf. This provides a check on the reasonableness of an optimum solution: $\tau_i < \frac{1}{2\Delta f}$. Compare this with assumption 3 in Part I.

Following is a summary of conclusions arrived at after many runs.

1. Optimum partitionings start with short and end with long subintervals, with correspondingly increasing τ_i.
2. The matrix inversion is generally stable with condition numbers less than 1000.
3. In order to determine how much good partitioning does, define a percentage improvement factor of a partitioned over an unpartitioned interval:

$$IF = \left(1 - \frac{\sigma_{fp}}{\sigma_{fu}}\right) \times 100$$

The improvement factor appears to quickly approach an asymptote as a function of number of partitions (Table B1).

4. With nine partitions and a twelve second spin time, σ_f is nearly identical for both 20Hz and 99Hz input noise bandwidth; the **improvement factor** thus increases with increasing input bandwidth. The optimum partitioning has shorter subintervals and time constants at the beginning of the interval in the wide bandwidth case than in the narrow case, as might be expected.
5. With 20Hz input bandwidth and nine partitions, comparison of the 12 second spin time with the 99 second spin time indicates that not only does σ_f decrease with spin time for a given partitioning, the improvement factor increases as well.
6. While the exact nature of the relationship between improvement factor, input noise bandwidth, and spin time was not investigated, it is clear that using a variable time constant helps less as the noise bandwidth and spin time decrease. It is also clear that for the magnetometer, which has a noise bandwidth of about 20Hz and a spin time of 12 seconds, a noise reduction of 30% does not justify the additional hardware

Table B1

Numerical Results

Input Noise BW	Total Spin Time	Number of Partitions	σ_f	Improvement Over Single Partition
20Hz	12sec	1	.405	-
		4	.299	26%
		5	.295	27%
		6	.293	28%
		9	.290	29%
99	12	1	.444	-
		9	.291	35%
20	99	1	.159	-
		9	.102	36%

and associated problems. With a software implemented analyzer, on the other hand, these problems do not arise, and a rapid convergence scheme would be used.

References

Papoulis, A., **Probability, Random Variables and Stochastic Processes**, McGraw-Hill, *1965.*

Bruggemann, H., **Temporal Filtering Using Pixel Incrementing,** *SMPTE Journal,* August, 1980.

APPENDIX C

The adaptive filter configured with its output after the integrators will produce DC signal levels proportional to the amplitudes of the portion of the input signal with the same phase and frequency as the reference signals. This is simply demonstrated by considering Figure C-1.

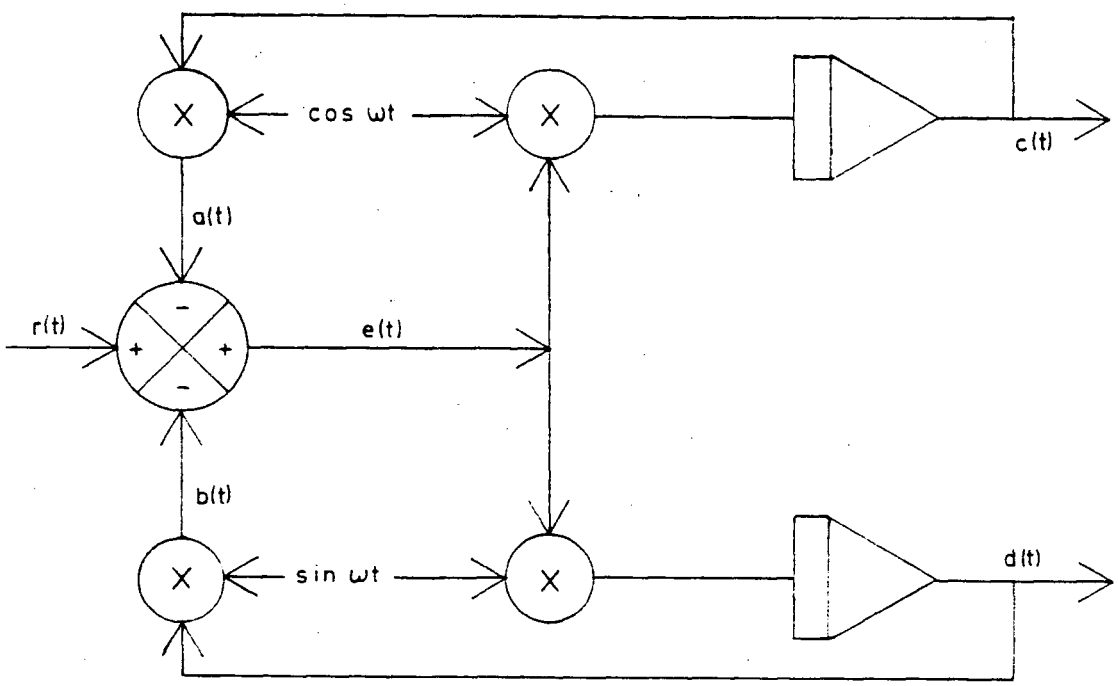

Figure C-1: Adaptive Filter

$$e(t) = r(t) - (a(t) + b(t)) \tag{C1}$$

$$b(t) = d(t) \cdot \sin \omega t$$

$$a(t) = c(t) \cdot \cos \omega t \tag{C2}$$

$$d'(t) = \frac{1}{RC} \cdot e(t) \cdot \sin \omega t$$

$$c'(t) = \frac{1}{RC} \cdot e(t) \cdot \cos \omega t \qquad (C3)$$

Eliminating $a(t)$, $b(t)$, and $e(t)$ from (C1), (C2), and (C3), and defining

$$\vec{D}(t) = [d(t), c(t)]^T$$

gives

$$\vec{D}'(t) + \frac{1}{RC} \begin{pmatrix} \sin^2 \omega t & \sin \omega t \cdot \cos \omega t \\ \cos \omega t \cdot \sin \omega t & \cos^2 \omega t \end{pmatrix} \vec{D}(t) = \frac{1}{RC} \begin{pmatrix} \sin \omega t \\ \cos \omega t \end{pmatrix} r(t) \qquad (C4)$$

For the case of an input signal of arbitrary phase and amplitude, and frequency ω,

$$r(t) = A \cos \omega t + B \sin \omega t \qquad (C5)$$

we may redefine

$$\vec{D}(t) = [d(t) - B, c(t) - A]^T \qquad (C6)$$

and rewrite (C4) as

$$\vec{D}'(t) + \frac{1}{RC} \begin{pmatrix} \sin^2 \omega t & \sin \omega t \cdot \cos \omega t \\ \cos \omega t \cdot \sin \omega t & \cos^2 \omega t \end{pmatrix} \vec{D}(t) = 0 \qquad (C7)$$

By inspection, a stable solution is

$$\vec{D}(t) = 0 \qquad (C8)$$

or

$$c(t) = A$$

$$d(t) = R \tag{C9}$$

as advertised.