

Signal processing techniques for on-line partial discharge monitoring of high voltage electrical machines.

JAMIESON, J.W.

1994

The author of this thesis retains the right to be identified as such on any occasion in which content from this thesis is referenced or re-used. The licence under which this thesis is distributed applies to the text and any original images only – re-use of any third-party content must still be cleared with the original copyright holder.

SIGNAL PROCESSING TECHNIQUES
FOR ON-LINE PARTIAL DISCHARGE
MONITORING OF HIGH VOLTAGE
ELECTRICAL MACHINES

A thesis submitted in partial fulfilment for the degree of Doctor of Philosophy.

JOHN W. JAMIESON

BEng. (Hons.) in Electronic and Electrical Engineering

March 26th, 1994

**The Robert Gordon University, Aberdeen in collaboration with B.P. Petroleum
Development Ltd.**

ABSTRACT

Signal Processing Techniques for On-Line Partial Discharge Monitoring of High Voltage Electrical Machines

John W Jamieson March 1994

Large A.C. electrical machines are found everywhere in modern industry. Traditionally, electrical plant has had a high reliability, primarily caused by the complexity of electrical design leading to generous over-rating. However, modern rotating electrical machines are now designed to tighter margins. In addition to this, there is a strong economic incentive to continue running older machines, with a relatively good thermal efficiency, up to and beyond their original design lifetime. It is not surprising, therefore, modern industry experiences a small but significant number of faults in their rotating plant, which, of course, eventually leads to failures and subsequent loss of availability.

One cause of machine failure is insulation breakdown of the stator winding. This can happen gradually and can be characterised by an increase in partial discharge activity over time. This activity can be measured using a well documented non-invasive technique. However, current analyses relies on experts visually interpreting results and making judgements based upon their knowledge and experience.

This work describes a completely new approach to the analysis of partial discharge data:

- A high speed analogue and digital electronic system was specified, designed and developed that is capable of directly digitising the extremely short duration partial discharges.
- Modular software routines were designed to control the electronics and display the digitised information for analysis.
- Additional software analysis routines were developed to process the discharge information in the time domain.

- A completely new technique was devised that allows the removal of interference signals from the monitored data that could otherwise prevent accurate analysis of the partial discharge severity.
- Different methods are investigated in order to measure the frequency content of partial discharges. This had never been attempted before, primarily due to the lack of suitable technology.
- Finally, a suitable frequency analysis algorithm was developed and applied to experimental data in order to investigate the feasibility of using the frequency content of partial discharges to ascertain their inception point within the stator winding.

DECLARATION

I hereby declare that this thesis is a record of work undertaken by myself, that it has not been the subject of any previous application for degree, and that all sources of information have been duly acknowledged

In the course of research, the following were included in an approved programme of advance studies:

- 1 week Advanced Signal Processing Course held at Southampton University, June 1988
- Attendance and presentation of a paper to the University Power Engineering Conference at Queen's University Belfast, September 1989
- 1 week trip to the Technical University of Berlin on an information exchange programme, May 1990

John W Jamieson

March 1994

ACKNOWLEDGEMENTS

I wish to acknowledge the following people and organisations, to whom I am indebted and without whom this project and thesis would not have been possible; Mr Stewart J Chalmers for his support and guidance as my main project supervisor and for his helpful comments and continuous encouragement to complete this thesis; Mr David G Edwards for his significant contribution of knowledge as project supervisor; BP Petroleum Development Ltd, in particular Prof I D Stewart, for funding and support of this project; and finally, Miss Lynn Ledingham for her understanding of the many late nights and weekends spent preparing this thesis.

OBJECTIVES

The research project described here is an addition to a well established and comprehensive programme of research, at The Robert Gordon University, related to on-line condition monitoring of large electrical machines. The work commenced in 1980 and, in recent years, has been funded by research contracts with industrial organisations, including BP Petroleum Development Ltd. (now BP Exploration Aberdeen)

The potential benefits of on-line condition monitoring are of greatest importance in high cost continuous process industries, such as oil production and electrical power generation. Techniques developed at the Robert Gordon University are already being applied in these industries, in particular, to the detection of incipient fault conditions in rotors of induction motors, and in high voltage stator windings of both induction motors and synchronous generators.

The work described here is directed towards non-invasive measurement and analysis of partial discharge activity within the stator of high voltage electrical machines. To that end, this investigation is primarily concerned with the development of hardware and software techniques that are capable of processing the semi-random partial discharge pulse patterns and identifying the characteristics which have most relevance to the diagnosis of insulation condition.

Table of Contents

ABSTRACT

DECLARATION

ACKNOWLEDGEMENTS

OBJECTIVES

CHAPTER 1 INTRODUCTION.....	1
1.1 Scope of the Thesis.....	2
1.2 Electrical Machine Condition Monitoring.....	3
1.2.1 Fault Mechanisms in 3-Phase Induction Motors.....	4
1.2.2 Condition Monitoring Strategies.....	6
1.3 Condition Monitoring of Electrical Insulation.....	6
1.3.1 Causes of Insulation Degradation.....	7
1.3.2 Review of On-Line Diagnostic Test Techniques.....	9
1.3.3 Partial Discharge Characteristics.....	10
1.4 Insulation Monitoring Using Rogowski Coils.....	13
1.4.1 The Rogowski Coil.....	13
1.4.2 Monitoring Configuration.....	14
1.4.3 System Calibration.....	16
1.4.4 Interference Signals.....	16
1.5 Summary.....	17
CHAPTER 2 SINGLE-PHASE PARTIAL DISCHARGE HARDWARE.....	18
2.1 Introduction.....	18
2.1.1 Technical Constraints.....	18
2.1.2 Operational Considerations.....	19
2.2 System Hardware.....	20
2.2.1 The Rogowski Coil.....	21
2.2.2 The High-Gain, Low-Noise Amplifier.....	21
2.2.3 The Anti-Alias Filter.....	27

2.2.4	Analogue to Digital Converter	33
2.2.5	Hardware Controller & Data Storage Module.....	35
2.2.6	Host Computer System.....	48
2.3	Summary	49
 CHAPTER 3 SINGLE-PHASE DISCHARGE MONITORING SOFTWARE		50
3.1	Introduction	50
3.2	Data Acquisition and Storage	52
3.2.1	Controlling the Data Acquisition Hardware.....	52
3.2.2	Storing the Raw Partial Discharge Data.....	52
3.2.3	Retrieving Partial Discharge Data Files	53
3.3	Waveform Display	53
3.4	Partial Discharge Pulse Analysis.....	55
3.5	Summary	58
 CHAPTER 4 THREE-PHASE PARTIAL DISCHARGE HARDWARE.....		59
4.1	Introduction	59
4.2	System Hardware.....	59
4.2.1	Three Phase Amplifier Configurations.....	62
4.2.2	Hardware Controller.....	63
4.3	Summary	72
 CHAPTER 5 THREE-PHASE MONITORING SOFTWARE		73
5.1	Introduction	73
5.2	General Display and Pulse Analysis.....	75
5.2.1	General Display Functions	75
5.2.2	Three-Phase Pulse Analysis Functions.....	77
5.3	Noise Rejection.....	81
5.3.1	Autocorrelation.....	82
5.3.2	Cross-Correlation.....	85
5.3.3	Three Phase Statistical Scheme.....	87
5.4	Application of Pulse and Noise Analysis.....	91
5.5	Frequency Analysis.....	98
5.5.1	Fast Fourier Transform.....	99

5.5.2 Maximum Entropy Method.....	117
5.6 Summary	136
CHAPTER 6 ANALYSIS OF SPECTRAL DATA.....	137
6.1 Introduction.....	137
6.1.1 Three Dimensional Display Options.....	137
6.2 Microsoft Excel as a Data Analysis Tool.....	142
6.2.1 Spreadsheet Introduction	142
6.2.2 Data Analysis Using Microsoft Excel.....	144
6.3 A Rudimentary Partial Discharge Location Classifier.....	161
6.4 Summary	166
CHAPTER 7 SYSTEM APPLICATION AND FUTURE DEVELOPMENT	167
7.1 Major Achievements	167
7.2 Monitoring Setup.....	168
7.3 Time-Domain Processing	169
7.3.1 Visual Analysis.....	170
7.3.2 Pulse Processing.....	170
7.4 Frequency-Domain Processing.....	171
7.5 Future Development.....	173
7.5.1 Existing System Specification	173
7.5.2 Future Hardware.....	173
7.5.3 Future Software and Processing.....	176
CHAPTER 8 CONCLUSIONS.....	181
REFERENCES.....	184
APPENDIX	

CHAPTER 1

INTRODUCTION

The research project described here is an addition to a well established and comprehensive programme of research, at The Robert Gordon University, related to on-line condition monitoring of large electrical machines. [1,2,3] The work commenced in 1980 and, in recent years, has been funded by research contracts with industrial organisations, including BP Petroleum Development Ltd (now BP Exploration Aberdeen).

The potential benefits of on-line condition monitoring are of greatest importance in high cost continuous process industries, such as oil production and electrical power generation. Techniques developed at The Robert Gordon University are already being applied in these industries, [4,5,6] in particular, to the detection of incipient fault conditions in rotors of induction motors, and in high voltage stator windings of both induction motors and synchronous generators.

The work described here is directed towards non-invasive measurement and analysis of partial discharge activity within the stator of high voltage electrical machines. To that end, this investigation is primarily concerned with the research and development of hardware and software techniques that are capable of processing the semi-random partial discharge pulse patterns and identifying the characteristics which have most relevance to the diagnosis of insulation condition.

A survey of published literature shows that there are various techniques and instruments that successfully monitor partial discharge activity. However, high-speed digitising of partial discharges and the subsequent application of digital signal processing techniques, particularly high-resolution spectral estimation methods, are relatively un-researched, especially in the area of on-line, non-invasive machine condition monitoring.

1.1 Scope of the Thesis

Chapter 1 introduces the topic of electrical machine condition monitoring with particular emphasis placed upon partial discharge mechanisms, their causes and consequences.

Chapter 2 deals with the investigation into the requirements of a practical partial discharge monitoring system, and the resulting hardware development. A substantial amount of hardware development had to be undertaken prior to the commencement of the basic research as suitable commercial hardware was not available. Design methodologies are detailed for all aspects of the prototype system.

In Chapter 3 the software hierarchy is outlined and the rudimentary display and storage functions described. In addition to this, signal processing algorithms, consisting of various pulse analysis routines, used in preliminary studies into the practicality of partial discharge processing, are detailed.

Chapter 4 follows the evolution of the monitoring hardware into a system capable of monitoring all three phases of a machine simultaneously.

Chapter 5 describes the software design process for the monitoring system, including the investigation and eventual implementation of a technique to remove periodic interference, commonly superimposed on the partial discharge data. This chapter also details the exploration into the frequency domain of partial discharge signals.

Chapter 6 deals with the analysis of the results obtained from the spectral data of partial discharges and describes how a rudimentary partial discharge location classifier was formulated.

Chapter 7 collates all of the previous information by demonstrating how the instrumentation system would be applied in real-life situations, and further, outlines possible enhancements to the basic design and the direction that future research in this area may take. Finally, Chapter 8 discusses the project's achievements and makes some general conclusions.

1.2 Electrical Machine Condition Monitoring

Large A.C. electrical machines are found everywhere in modern industry; from pumps to generators, their applications are endless. Traditionally, electrical plant has had a high reliability, primarily caused by design parameter over-rating. However, modern rotating electrical machines are now designed to tighter margins as economics, and subsequently profit, take on greater significance in a very competitive marketplace. In addition to this, there is a strong economic incentive to continue running older machines, with a relatively good thermal efficiency, up to and beyond their original design lifetime.

It is not surprising, therefore, that in a climate of such harsh economic reality, modern industry experiences a small but significant number of faults in their rotating plant, which, of course, eventually leads to failures and subsequent loss of availability.

The financial costs of such losses in availability of machines in-service commonly represents a significant percentage of the capital cost of the machine. This is especially so in high-cost operational environments, such as offshore oil and gas production, where an entire platform may have to shut down if a particularly crucial machine fails. For this reason, on-line condition monitoring of electrical machines has taken on considerable importance, allowing operators to more accurately schedule maintenance and provide significant pointers regarding imminent machine failure.

The benefits that condition monitoring can bring to any particular machine depend on:

- The amount of warning that can be given about an impending failure.
- The value of the machine itself.
- The total cost of the machine failing, including:
 - loss of production due to the failure
 - direct cost of repair or replacement
 - whether standby capacity or a spare machine is available
- The cost and reliability of the monitoring system itself.

The amount of warning that can be given about an impending failure is very much a function of the monitoring strategy employed, as well as the time taken for the fault to develop.

e.g. If a monitoring scheme could only detect a fault one second before the machine fails, then this scheme would not be of much use to the operators. But, if a scheme could progressively monitor the development of a fault over a number of weeks or months, it would be of considerable value to the operators. However, some faults can develop very rapidly and even continuous monitoring may not detect them quickly enough to prevent breakdown.

Condition monitoring of an electrical machine will only yield benefits if:

- The monitoring system is reliable and can give sufficient warning of a failure in most cases.
- The machine can be taken out of service and repaired at a cost substantially less than the likely cost of repairs following failure.
- The results of monitoring produce significantly less of a production loss due to maintenance scheduling, as opposed to unexpected failures.

1.2.1 Fault Mechanisms in 3-Phase Induction Motors

The most prevalent machine found in modern industry today is the induction motor. Found in all sizes, the induction motor is used in household and light industrial applications, as well as heavy industry and power generation. Consequently, it is important to be familiar with the fault mechanisms and some of the monitoring techniques associated with induction machines.

Large induction motors are often used in harsh or hostile environments, such as mines or offshore oil exploration and production platforms. Especially in places such as offshore platforms, these machines are exposed to diverse amounts of contamination and are often operated quite harshly (direct on-line starting is prevalent whereby a machine is started at full load resulting in high mechanical and thermal stresses)

which, not surprisingly, can lead to faults not commonly experienced in machines used in more moderate applications.

The manifestations of faults can be either electrical, mechanical or a combination of both, and include such faults as cracked or broken rotor bars, eccentricity problems, bearing faults and winding problems.

Rotor bar breakages are quite a common fault found in motors, caused when the conducting bars in the squirrel cage rotor break away, or just crack, at their junction with the conducting end-ring of the rotor. This may happen for a number of reasons, such as casting problems during manufacturing producing weaknesses, or excessive vibration caused by operational factors. When a bar cracks it produces a high resistance joint on the end-ring, causing overheating which may eventually lead to the bar completely breaking. This, in turn, leads to further mechanical damage to the motor, caused by the resulting torque fluctuations and corresponding change in vibration patterns.

Eccentricity is when the air gap between the rotor and stator is non-uniform and comes in two forms: static eccentricity, whereby the air gap profile is constant (i.e. the minimum air gap always appears at the same point, for example: a rotor/stator misalignment), and dynamic eccentricity where the minimum air gap rotates with the rotor (e.g. bent rotor shaft). The problems caused by eccentricity are two-fold: firstly, in severe cases, the rotor may actually make contact with the stator, causing a "rub" and damaging both the stator and rotor. Secondly, the non-uniformity of the air gap, and thus the magnetic flux, may cause abnormal mechanical stresses and vibration patterns, resulting, among others, excessive bearing wear which may in turn lead to yet greater levels of eccentricity.

Stator winding faults are very commonly (30-40%) the cause of machine failure. They may develop as a result of insulation contamination, conductor vibration, voltage transients/ impulses due to welding, or insulation degradation due to internal discharges. If the insulation breaks down, then a short circuit can locally occur between coils, causing even more damage. Severe insulation degradation, if undetected, may propagate through the stator in a very short time.

Most faults, however, develop gradually (e.g. as the insulation material ages or contamination increases) so periodic monitoring can help prevent catastrophic failure of the motor, as well as providing guidance in scheduling maintenance.

1.2.2 Condition Monitoring Strategies

Two common methods of monitoring induction motors are vibration and stator current monitoring.

Vibration monitoring utilises a number of transducers, usually fixed around the bearings of the motor, to monitor the varying patterns of vibration. This can provide an indication of different types of faults, as the vibration is caused by the changes in flux pattern of the air gap producing mechanical vibrations or acoustic noise. This monitoring method has been employed successfully in a number of industries, but has the inherent disadvantage of having to affix transducers to the machines being monitored, often in quite inaccessible places.

The current drawn by an ideal motor should have a single component at the supply frequency. Any changes in load will cause speed fluctuation and modulate the amplitude of the current, producing sidebands around the fundamental when viewed in the frequency domain. Defects in the rotor will generate further sidebands which can be used as the basis of a monitoring strategy [1]. This technique has the advantage that no direct access to the motor is required as the phase current can be monitored by a simple clip-on current transformer placed around the phase wire at the control room end.

Winding faults are monitored using different techniques, and as they are of particular importance to this project, will be dealt with in more detail.

1.3 Condition Monitoring of Electrical Insulation

It is known that 30-40% of all large motor failures are attributable to the breakdown of the high-voltage stator insulation. Irrespective of the incidence of insulation faults, it is clearly advantageous to periodically monitor the state of the insulation with a

view to identifying a problem before it becomes a disaster, or to allow repair work to be scheduled in a way to minimise loss of production.

Traditionally, diagnostic testing of machine insulation has been done off-line [8], whereby the machine is disconnected from the normal supply and various non-destructive tests performed using techniques such as DLA (Dielectric Loss Analyser) or simply using a Megger. This system of monitoring has inherent disadvantages, not the least of which is that these tests can only be carried out during routine maintenance shutdown periods.

On-line diagnostic techniques are now becoming more popular because they can be applied whilst the machine is running in normal service. Also, these tests can be applied as often as is considered necessary, and the resulting measurement characteristics reflect the motor's true operating role.

1.3.1 Causes of Insulation Degradation

The mechanisms that cause the electrical insulation to break down within large machines can be split into four broad categories: thermal stress, electrical stress, mechanical stress and environmental stress.

Any increase in the temperature of a machine will cause thermal ageing of the insulation. At high temperatures the insulation will break down very rapidly: for every 10°C above its rated limit, the insulation life is halved [9]. Thermal degradation results in delamination, cracking, embrittlement or depolymerisation. This leads to further wear of the winding, speeding up the breakdown process.

Electric stress across a gaseous inclusion causes ionisation and there is a current flow or partial discharge through the ionised gas. Progressively, the ionisation process changes the chemical properties of the surrounding resin and channels are cut into the insulation base. These discharges occur external to or within small voids in the coil insulation structure and are generated as part of the normal stator winding characteristic. Figure 1.1 shows the discharge process within a small gas filled void in the insulation structure.

As the electric stress builds up across the insulation, following the high voltage supply, a corresponding increase in stress builds up across the void. Upon reaching the breakdown threshold of the gas inside the void, a discharge occurs across it, temporarily relieving the electric stress. If the supply voltage is still rising, the stress then begins to build up again until the whole discharge process repeats. The entire process repeats each half cycle.

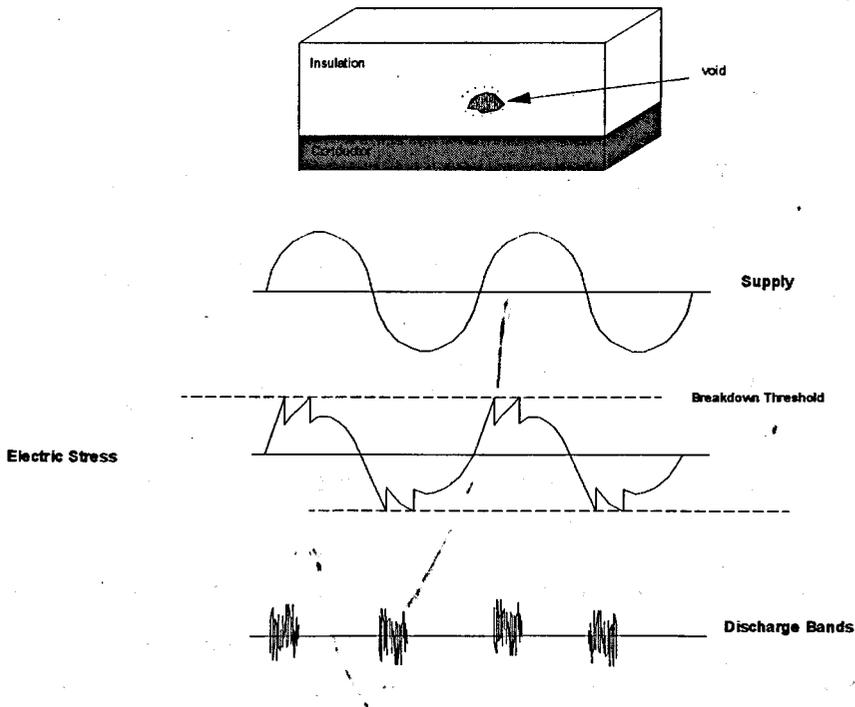


Figure 1.1 Partial Discharge Mechanism

The erosion effects due to electric stress may be accelerated by transient pulsing from switching surges and other system disturbances, such as direct-on-line starting. Discharge erosion, especially slot discharge erosion whereby a surface current of 'corona' flows between the outside of the insulation and the core laminations, can affect the fit of the coil in the slot by damaging the winding insulation or core packing, leading to accelerated erosion due to increased mechanical stresses.

Mechanical stresses produce fretting between the copper, insulation and core and can result from any number of causes:

- Differential thermal expansion and contraction.
- Severe direct-on-line starting duty.
- Stator core or tooth vibration.
- Electromagnetic excitation.
- Inadequate bracing or packing of the slots.

Contamination of the end-windings from substances in the operational environment, such as water and salt, also contribute to the breakdown process. Contamination of the insulation can change its electrical properties which, in turn, can result in increased partial discharge activity and more rapid degradation [5]. This environmental stress factor is of particular importance for machines operated in harsh conditions, as in offshore oil and gas production platforms.

1.3.2 Review of On-Line Diagnostic Test Techniques

Certain measurable quantities can directly correlate with insulation damage. These include mechanical vibration, gaseous products and partial discharges [10]. All three quantities are, to some degree, measurable during normal operation.

Mechanical vibration may be caused when a bar or coil loosens within a stator slot. Measuring the magnitude and rate of increase of such vibrations can provide an indication of the severity of the damage and can be a pointer to ultimate insulation failure. Vibrations can be measured by means of accelerometers attached to the stator frame. However this technique is not normally used as it will also pick up vibrations due to other causes, such as the rotor faults discussed earlier.

Certain dielectric's, when exposed to heat or partial discharges, evolve various gaseous by-products that can be identified and quantified using gas chromatography or spectroscopy. Additionally, discharges occurring in air external to the insulation, produce ozone which can also be detected. The resulting level of evolved gases provide a direct measure of the damage occurring.

Partial discharges are measurable using a variety of techniques[7,9,10,12]. These include wrapping a high frequency current transformer around the neutral connector, detecting pulses developed across a blocking capacitor which is attached to the terminals of the machine and detecting the radio frequency energy, generated when a partial discharge occurs, by placing a RF aerial inside the machine.

In addition to this, there are a number of off-line techniques used to measure partial discharge activity, but these have inherent disadvantages [5] and are not considered further:

- Tests can only be carried out during routine maintenance shutdown periods which may be 2 years or more apart and therefore may not detect rapidly developing faults.
- Test equipment is expensive and bulky.
- Whole phase winding is stressed to the same potential so the test is not representative of normal service.

The on-line technique adopted in this project uses Rogowski coils, essentially high frequency current transformers, which are connected to the phase windings of the machine and detect the high frequency current pulses, caused by partial discharge activity, as they propagate out of the stator winding. Before this technique is reviewed further, it is useful to examine the characteristics of partial discharges.

1.3.3 Partial Discharge Characteristics

When a partial discharge occurs in the insulation structure of a high voltage stator winding, a short-duration current pulse is generated in the HV conductor adjacent to the discharge site. This current pulse effectively replaces the energy lost to the insulation structure when the discharge occurs. The shape of the current pulse reflects the time sequence of the partial discharge 'sparking' process. i.e. it has a very sharp rise time and a slower fall time. the resulting electrical disturbances travel away from the discharge site, along the winding conductors, at a speed which is a characteristic of the winding. It has been shown that the stator winding behaves like a transmission line to high frequency events, with frequency-dependant attenuation [11]. During the propagation process, the shape of the pulse changes, due to the

frequency-dependant attenuation characteristic, such that on arrival at the machine terminals, where it can be measured, the shape of the pulse bears little resemblance to the original discharge pulse.

It has been shown [12] that the current pulse arriving at the terminals of the machine is, in fact, made up of two separate pulses, a low magnitude fast pulse and a higher magnitude slow pulse, Figure 1.2. Depending upon the inception site of the discharge, a variable time will elapse between detecting the fast and slow pulses. Also, the fast pulse has a lot higher frequency components than the slow pulse, therefore the frequency response of the detection system can dictate if either, or both, pulses are measured.

Figure 1.2 highlights the differences in measured partial discharges when they occur in different parts of the winding. The upper plot shows the resulting detected pulse when a pulse is injected three coils distant from the HV terminals, and the lower graph shows the same but for a discharge injected seven coils distant. Examining the two graphs, attenuation of the higher frequencies associated with the slow pulse can be noted as the pulse travels further through the winding. Also the difference in the propagation speeds can be highlighted by examining the time intervals between detecting the fast and slow pulses in both cases.

The propagation delay of the slow pulses is used as a basis for a location technique to identify discharge sites within a stator winding [3].

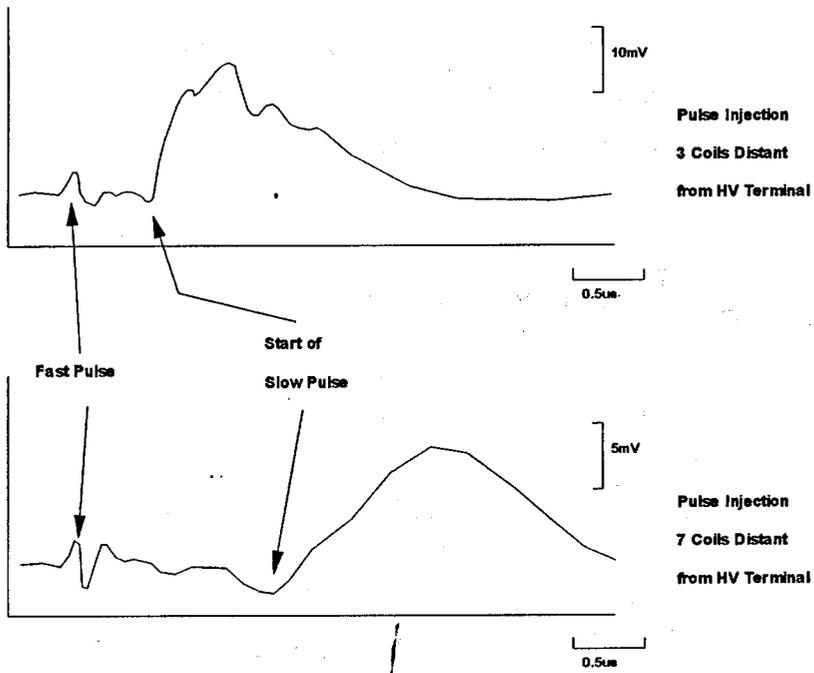


Figure 1.2 Typical Fast & Slow Propagating Pulses

Note: Amplitude information in Figure 1.2 is included for comparison purposes only as no information is given on the energy content of the injected pulse.

1.4 Insulation Monitoring Using Rogowski Coils

It is very desirable, from an operators point of view, that the basic transducer used to obtain on-line partial discharge information should be of a totally non-invasive nature thereby reducing installation costs. This means that ideally the coupling between the high voltage circuit and the monitoring system should be inductive, rather than capacitive.

1.4.1 The Rogowski Coil

A Rogowski coil is essentially an air-cored current transformer wound in a toroid. Such coils have been used for many years in a variety of roles ranging from the measurement of high-energy electron beam current pulses to the monitoring of power frequency shaft currents in large rotating machines.

At power frequencies (50-60Hz) the Rogowski coil has a very low inductive impedance, necessitating the use of integrating circuitry to measure power frequency primary currents. However, this low inductive impedance is clearly an attractive feature for partial discharge monitoring purposes, since the high voltage circuit will often be carrying a motor supply current of hundreds of amperes in the presence of relatively minute discharge current pulses. The higher frequency content of partial discharge pulses means that the Rogowski coil no longer needs integrating termination and can be designed to reject the power frequency components whilst passing the higher frequencies associated with partial discharge activity [5].

The voltage distribution in the phase winding of a machine is graded uniformly from the line terminal down to the star point. In practice, this means that most of the high energy discharges occur in the first few coils adjacent to the line-end. Therefore, the ideal location for the Rogowski coil for on-line monitoring is at the line-end terminals of the machine. It is possible to make measurements at the star point however, but the attenuation characteristics of the winding mean any monitored discharges would be severely attenuated due to the distance that they have to travel through the winding. It is also possible to make measurements at the switch gear end of the supply cable as the attenuation characteristics of the cable are less pronounced than those of the stator winding itself. However, a number of machines are typically

connected to the same high-voltage busbars via relatively short lengths of cable, so any measurements made could be contaminated by signals from different sources.

1.4.2 Monitoring Configuration

The typical monitoring arrangement adopted when using Rogowski coil transducers is shown in Figure 1.3.

A Rogowski coil is connected around each phase supply cable. The output from each coil is then fed into an amplifier to boost the signal level. Once amplified, the resulting partial discharge signals are displayed on an oscilloscope or special discharge detector, as well as recorded on magnetic tape for later analysis, or just as a permanent record.

The phase reference unit is used to produce a reference supply waveform with which the partial discharges, from any given phase, can be correlated. A typical display example is shown in Figure 1.4.

From Figure 1.4, it can be seen that the larger partial discharges normally occur in bands, 180° apart, at the points of maximum electric stress, referred to the supply voltage.

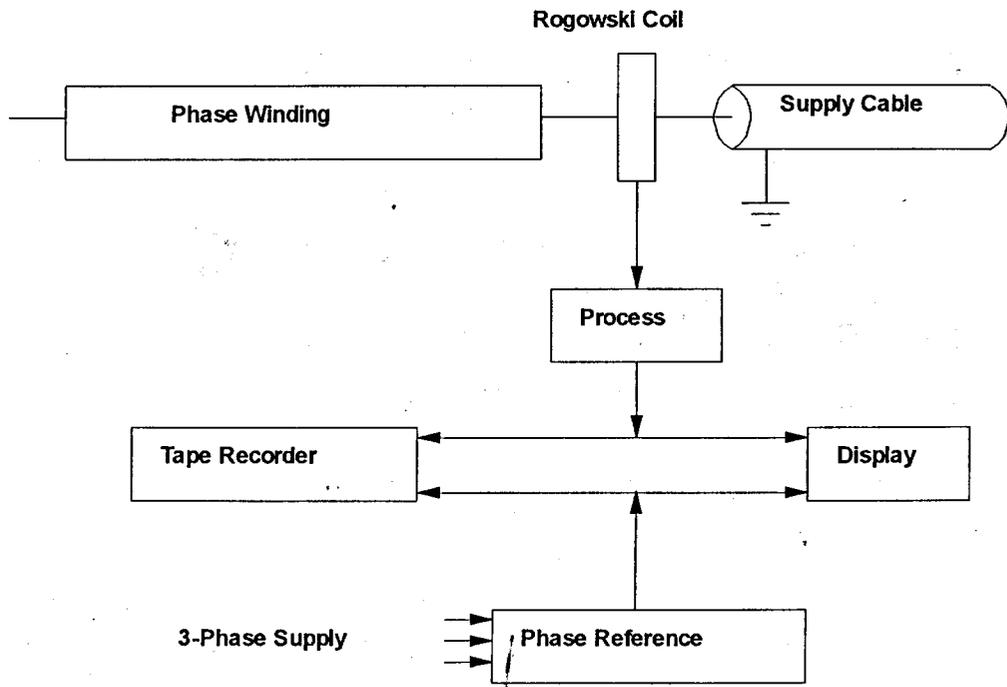


Figure 1.3 Typical Monitoring Arrangement

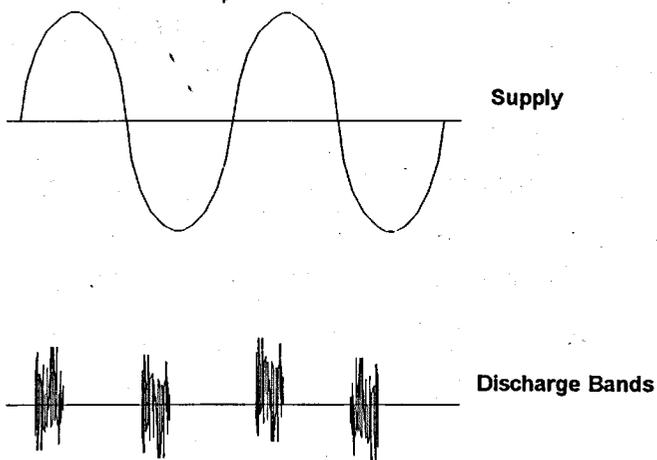


Figure 1.4 Display Example

1.4.3 System Calibration

The measured partial discharge current pulses need to be quantified in some way that pertains to their discharge energy. This is done by calibrating each machine under test by injecting a pulse of known charge content into the motor terminals. The simulated discharge is produced by means of a low voltage source capacitively coupled to the high voltage line. Once this calibration pulse has been detected and processed by the system, all subsequent detected partial discharge pulses can be compared to the calibration pulse and thus their charge content inferred in pico-coulombs.

Fortunately, each motor/cable combination needs only to be calibrated once, as the calibration process is most conveniently carried out when the machine is shut down.

1.4.4 Interference Signals

Typically, in an industrial situation, several high-voltage machines will be connected to the same busbars. Under some conditions this could lead to signals from one machine contaminating the results obtained from another. However, external signals travelling along the supply cable towards the machine under test will already have experienced significant attenuation due to energy being transferred onto multiple connections on the busbar [5]. In addition, the Rogowski coil is relatively insensitive to external interference signals in this application.

In some situations external interference signals do find their way into the monitoring system. These are typically attributable to high power switching applications, such as thyristor speed controllers for large motors. These controllers produce signals of very similar frequency content to partial discharges, but of a significantly greater amplitude so they tend to get propagated into the monitoring system and swamp any discharge activity that is occurring at that time.

These interference pulses have a regular repetition rate so can easily be distinguished from the random nature of the partial discharges. This means, that as far as an expert is concerned, the interference signals are no more than an inconvenience. However, they pose a significant problem when developing an automatic monitoring system.

1.5 Summary

Electrical machine condition monitoring is a very diverse and complex field. Various techniques have been outlined that are capable of monitoring specific faults as they develop. Such techniques are continually being adapted and improved as a competitive marketplace demands more accurate and reliable methods.

A non-invasive technique has been described, that is used to monitor partial discharge activity within the stators of high voltage electrical machines. This technique is used as the basis for the further development of hardware and software to process the discharge pulse patterns and to investigate means with which to extract further useful information from the available signals.

CHAPTER 2

SINGLE-PHASE PARTIAL DISCHARGE

HARDWARE

2.1 Introduction

The concept of a computer-based instrumentation system for the detection and analysis of partial discharges is by no means a new one. Austin and James [15], back in 1976, demonstrated the feasibility of such a system, limited only by the available technology. Tanaka and Okamoto [16,17,18], using superior technology, considerably advanced the capabilities of the computer-based measurement system. These systems invariably relied upon some form of pulse-stretching or similar detection techniques to capture the partial discharges. Technological advances now make it possible to directly digitise partial discharges and analyse their properties using the new generation of microcomputers.

The design of such a system, however, is not straightforward as there are a number of technical, as well as operational, constraints to be taken into consideration.

2.1.1 Technical Constraints

The non-invasive method of detecting partial discharge activity results in signals of very small energy (2000-8000pC), immersed in noise [5]. This places stringent design requirements upon the initial analogue processing stages, specifically the need for adequate amplification whilst maximising bandwidth and limiting noise.

The digitising process also has far-reaching design implications: most importantly, the speed and resolution of the analogue to digital converter has to be a compromise between the bandwidth and noise performance of the amplifier, the specification of the anti-alias filter, the subsequent memory requirements for storing the digitised discharges, and the information processing speed of the microcomputer.

2.1.2 Operational Considerations

Ultimately, any monitoring system developed would be operated in an industrial environment, therefore the relevant working practices and conditions have to be taken into account during the design process. Special consideration has to be given to the effect of electrical interference upon the system, potentially the most serious of which is the thyristor switching of DC motor speed controllers. These electrical drives produce periodic pulses of relatively large amplitude ($>100000\text{pC}$) with similar frequency content to partial discharges, thus they tend to be processed by monitoring instrumentation, swamping concurrent partial discharge activity, leading to difficult diagnoses.

Having outlined the important design constraints, the system hardware can now be described in detail.

2.2 System Hardware

A block diagram of the initial partial discharge monitoring system can be seen in Figure 2.1.

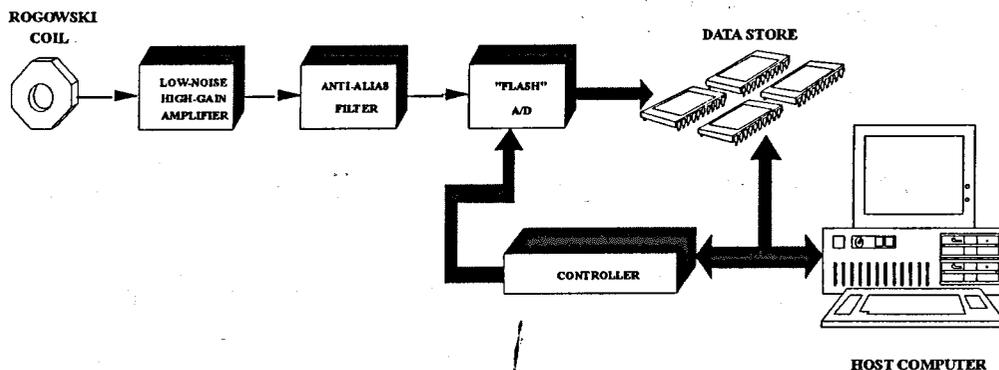


Figure 2.1 System Overview

This system was designed to monitor a single phase of a high voltage machine under test, and was built primarily as a development tool to investigate the feasibility of analysing digitised discharges.

The operation of the system is as follows:-

The basic transducer, a Rogowski Coil, is clamped around the supply cable to one phase of the machine under test. The transducer picks up any high frequency current pulses, caused by partial discharge activity, as they propagate out of the stator winding back towards the supply, and the resulting signals are passed to a low-noise, high-gain amplifier to boost their inherently low amplitudes. Once amplified, the partial discharge signals are applied to an anti-alias filter network to limit their high

frequency content in order to prevent sampling errors occurring when digitised. After the signals have been fully conditioned, they are digitised at a rate of 15MHz before being placed in a bank of Static RAM, ready to be processed by the host computer.

The system was designed to be as flexible as possible in order to accommodate future expansion as well as facilitating design and/or processing changes.

The following sections deal with the individual system components in more detail.

2.2.1 The Rogowski Coil

On-line condition monitoring applications ideally require non-invasive transducers to gather the relevant information. In the case of partial discharge monitoring, special Rogowski Coils have been developed[5], discussed earlier in Section 1.4.1.

2.2.2 The High-Gain, Low-Noise Amplifier

The partial discharge pulses, detected by the Rogowski Coil, are of several millivolts amplitude. In order to quantify the design requirements of the initial amplification stage, investigations had to be carried out into the exact nature of the signals that would ultimately be applied to the system.

The amplitude and shape of partial discharge pulses vary both with the size and origin of the discharges within the stator winding. In order to extract the maximum amount of information from such pulses, a high bandwidth system with good dynamic range and accompanying noise performance is required. Previous experiments using the Rogowski Coil measurement technique [5] have shown that the upper frequency limit need not be as high as initially expected. This led to the upper amplification frequency limit to be fixed at 4MHz, a reasonable compromise between frequency resolution and subsequent design difficulty.

In determining the necessary performance of the amplifier stage, laboratory experiments were conducted on a 6.6KV stator winding, the general set-up of which is shown in Figure 2.2.

Artificial discharges of known magnitude were applied to the line-end of the stator by means of a voltage generator (E) in series with a capacitor (C), with the terminating network constructed to mimic the surge impedance of the machine supply cable.

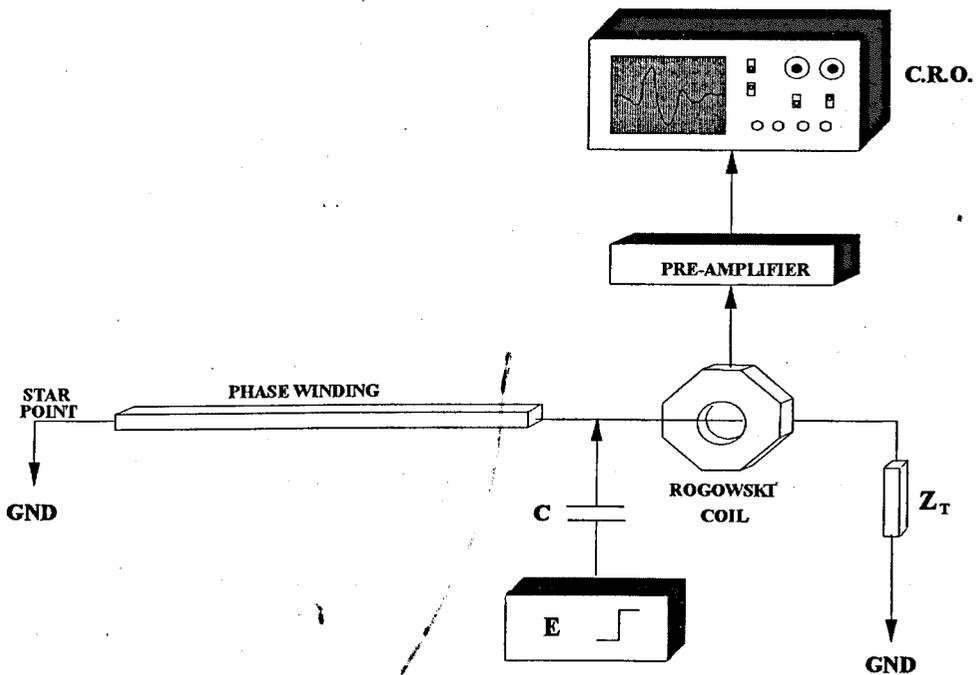


Figure 2.2 Experimental Set-up

The first amplifier design tested on the laboratory set-up used an ultra-high frequency operational amplifier AD5539, the circuit of which can be seen below in Figure 2.3.

The operational amplifier, AD5539, used in this pre-amplifier design, has a gain-bandwidth product of 1.4GHz, a unity gain bandwidth of 220MHz and a high slew rate of 600V/ μ s. The small signal open-loop gain and phase plots show that the device is stable with noise gains greater than 5, when used as a non-inverting amplifier. However, lead compensation was subsequently required to modify the frequency response of the amplifier to improve its stability.

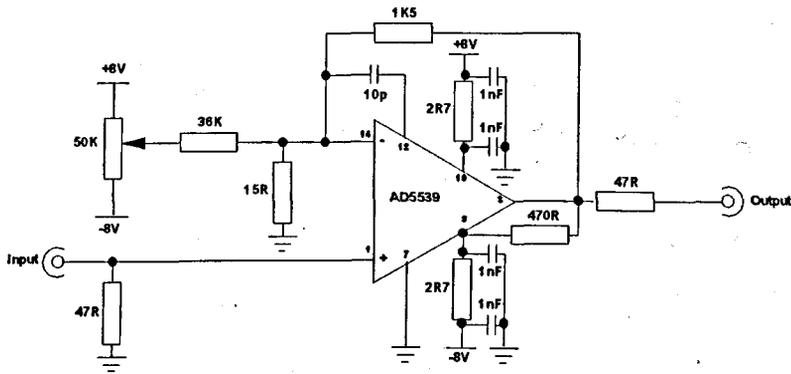


Figure 2.3 Amplifier Circuit Diagram

The pre-amplifier was constructed on a printed circuit board and tested using the arrangement of Figure 2.2. Simulated discharges of 1000pC in magnitude were injected into the line end of the stator winding, using the relationship $Q=CV$ to determine the size of the simulated discharge. The relatively high surge impedance of the stator winding means that most of the injected pulse is transmitted to the termination unit, which merely acts like a length of feeder cable. Thus the simulated discharges are detected by the Rogowski Coil and applied to the pre-amplifier circuit.

Initially, a capacitor of 100pF was used with applied voltage steps of 10V, resulting in charge injections of 1000pC into the stator winding. The output from the pre-amplifier was measured to be approximately 0.48V pk-pk with 10mV of background noise.

With charge injections of 10-15pC the resulting output voltage from the pre-amplifier was found to be in the region of 30mV. However the output amplitudes for these differing charge injections were not in proportion, therefore 100000pC charges were

injected to verify that a problem existed. The output measured in this case was 1.5V pk-pk, not saturated as expected. The Rogowski Coil was then disconnected and the experiment repeated, with the results indicating that the majority of the signals being applied to the pre-amplifier were not originating from the transducer, but instead were directly attributable to pick-up.

The amplifier was then placed in a screened box, to help alleviate the interference problem. However, this only served to make the amplifier unstable, probably attributable to the poor layout of the pcb. It was therefore decided to build a new amplifier version with additional features incorporated to alleviate the instability problems.

The second pre-amplifier design can be seen in Figure 2.4.

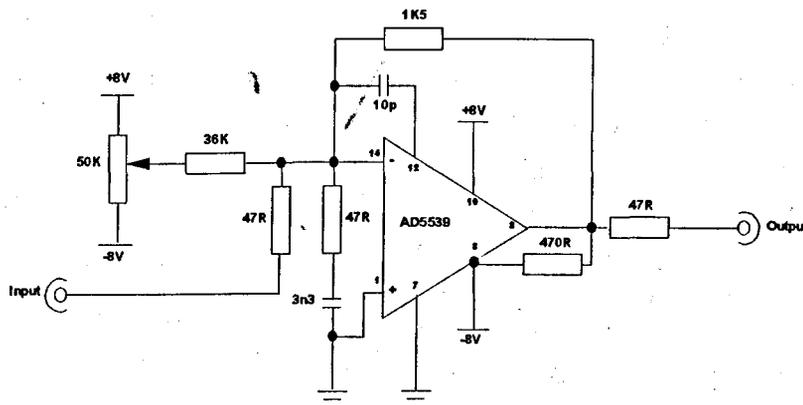


Figure 2.4 Second Amplifier Design

The pcb of version 2.0 has various options incorporated into its design, namely the ability to configure the op-amp in non-inverting or inverting mode, the option of driving the input cable screen from the feedback network, facilities for both lag and lead compensation, as well as the possibility of signal DC blocking.

The RC lag network (Figure 2.4) was designed to introduce a pole at approximately 1MHz in the open loop response. This pole does not affect the closed loop response until about 10MHz, well above the frequency range of interest, thus improving amplifier stability without compromising the desired response. In addition to the lag network, lead compensation was introduced in order to fully stabilise the amplifier response by giving an extra 90° phase margin at 10MHz.

The amplifier configuration of Figure 2.4 was evaluated experimentally and was found to have a gain of about 40 from DC to 14MHz, before falling off sharply. The original test arrangement was then used to see if any anomalies still existed in the amplifier design, but again, pick-up prevented true gauging of the amplifier performance along with the Rogowski Coil. Further evaluation of the amplifier design led to the conclusion that the AD5539 operational amplifier was not ideally suited for a pre-amplifier stage as its noise performance could lead to problems if a high gain system was needed. It was subsequently decided to design a low-noise "front-end" using discrete transistors, to produce a stage with an overall gain of about 20, so reducing the need for such high noise performance in subsequent stages.

A block diagram of the envisaged system can be seen in Figure 2.5.

The criteria behind this configuration of amplifier relies on the principle that the overall noise performance is predominantly dependent on the first stage. In order to achieve better noise performance discrete transistor amplifier configurations were investigated.

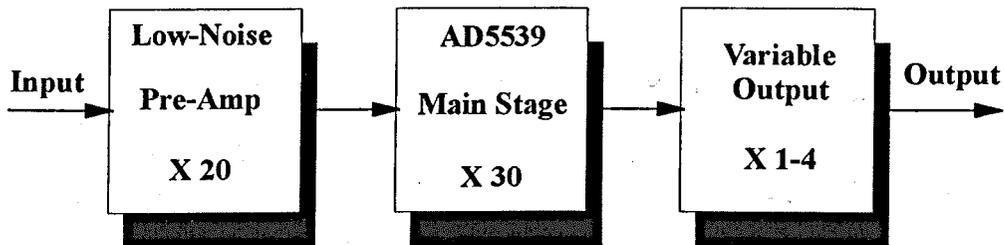


Figure 2.5 Block Diagram of the Overall Amplifier

Many different configurations of transistor amplifier were tested with results indicating that a FET input stage was most desirable, providing high input impedance, coupled to further bipolar transistor stages for the main gain of the amplifier. Further investigations led to the conclusion that designing a suitable pre-amplifier would be too costly in time, for what was not the main area of the project. Therefore it was decided that a specialist pre-amplifier would be bought at a later stage in the project, allowing work to continue on the next stages of the analogue circuitry.

The next step in the amplifier design process was to develop the main amplifier stage, based on an op-amp, to produce a gain of about 30.

At this time a new operational amplifier became available, from Analog Devices, the AD840 wideband, fast settling op amp. This new device offered a gain-bandwidth product of 400MHz and was stable at gains of 10 or more. This meant, in practice, that no compensating circuitry was necessary as the required gain was 30.

A complete amplifier system was constructed, incorporating the discrete input stage and subsequent AD840 op-amp stage. The resulting amplifier was encouraging, but still subject to high levels of noise, presumably due to its power supply and pcb layout. It was subsequently decided to "buy in" a suitable amplifier and concentrate on the other aspects of the system hardware.

2.2.3 The Anti-Alias Filter

It can be shown by Fourier methods that the spectrum of a sampled signal is an infinite series of repeats of its continuous time spectrum, Figure 2.6. If the sampling frequency is at least twice the highest frequency component of the original, then these "copies" of the original are distinct, and the information contained in the original signal is not corrupted by the sampling process. If this condition is not met, then the low frequency components of one copy mingle with the high frequency components of the adjacent lower copy, Figure 2.7.

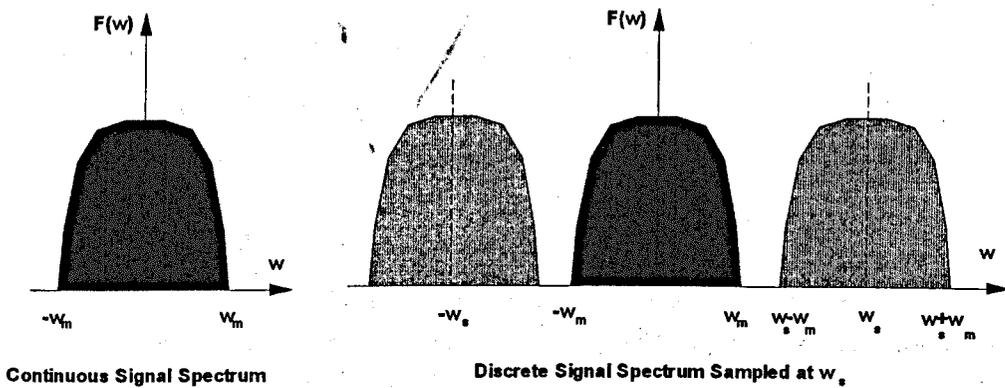


Figure 2.6 Signal Spectrum Comparison

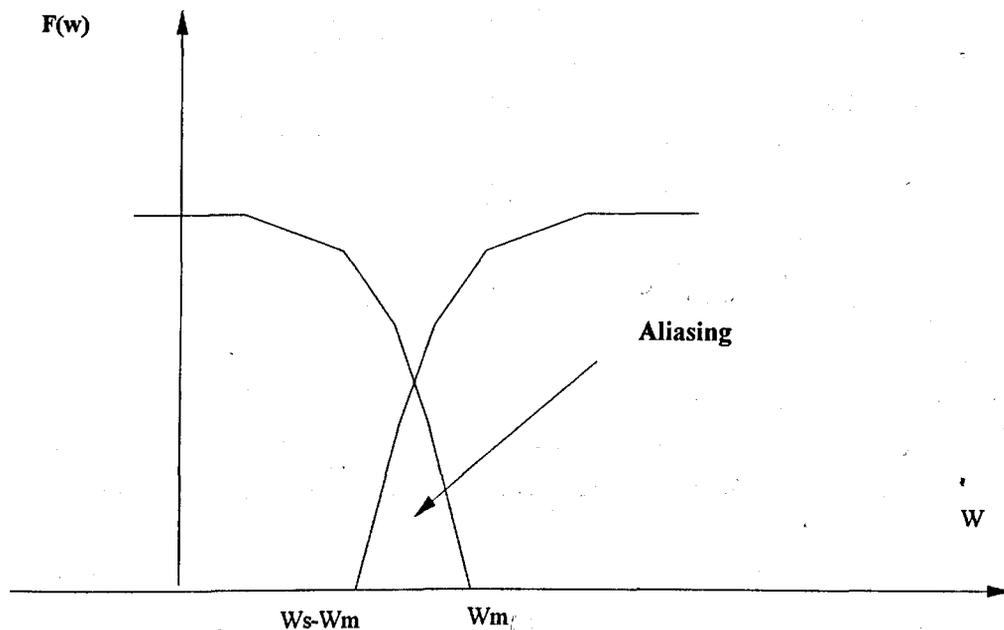


Figure 2.7 Aliasing Effect

This condition of high and low frequency components intermingling is called aliasing. In effect, a signal that is sampled at too low a frequency has an apparent spectrum which is not that of the original, due to high and low frequency components being intermixed.

Any analogue signal which is sampled prior to digitising is liable to aliasing unless the sampling frequency is chosen high enough to prevent it, or the relevant high frequency components of the signal are attenuated to below some appropriate level.

This means, in practice, that aliased components must be kept below the noise threshold of the signal channel. This is achieved by using anti-alias filters to limit the frequency content of the signal prior to sampling, and by specifying a sampling frequency which is high enough to create enough bandwidth for the filter attenuation slope to bring the aliased components down to the desired level.

Anti-alias filters are commonly implemented using cascaded active filter networks, designed using the traditional active filter polynomials.

In this case, a number of factors influenced the design of the anti-alias filter network:-

- The sampling frequency had been dictated, by other constraints, to be 15MHz.
- The signal to noise ratio of the system was approximately 48dB (8-bit analogue to digital converter).
- A large signal bandwidth was desired in order to maximise the amount of usable, extractable information on the partial discharge activity, without compromising the noise or phase integrity of the signal with a large filter network.

Certain of the common filter polynomial approximations are designed for optimum amplitude characteristics, with less importance placed on phase behaviour. The Butterworth and Chebychev polynomials fall into this category. The Bessel polynomial approximations are more appropriate for anti-alias filter design as they lead to linear phase response filters. Unfortunately, Bessel filters have an inadequate rolloff characteristic for this application, so a compromise had to be reached. This came in the guise of Thompson-Butterworth filters. This class of filters are a compromise between the classic Bessel and Butterworth designs, which give improved attenuation characteristics whilst maintaining a linear phase response.

The particular circuit configuration employed in this type of filter is the standard Unity-Gain, Sallen-Key, second-order low pass section, shown in Figure 2.8.

A number of alternative design possibilities exist for this configuration, with Figure 2.9 showing the various attenuation characteristics for 6th-order Thompson-Butterworth filters.

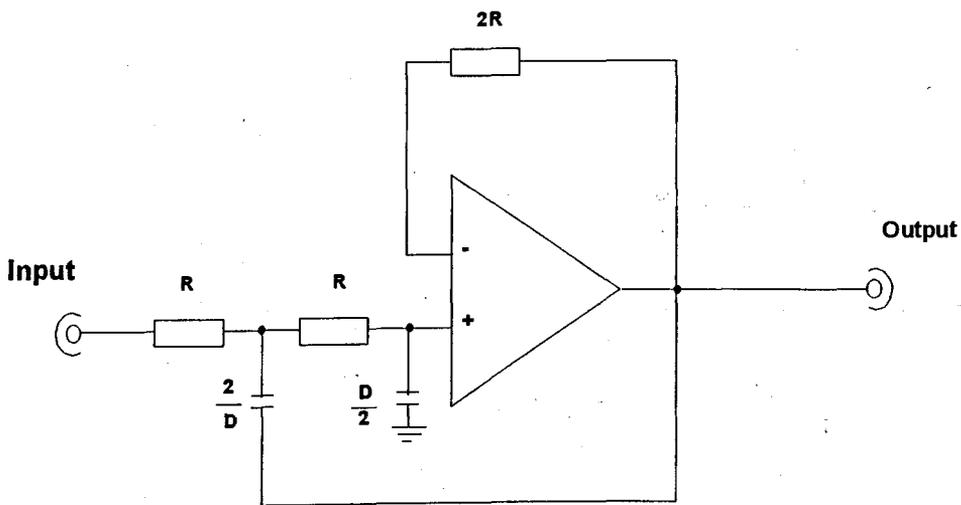


Figure 2.8 Unity-Gain, Sallen-Key Low Pass Filter

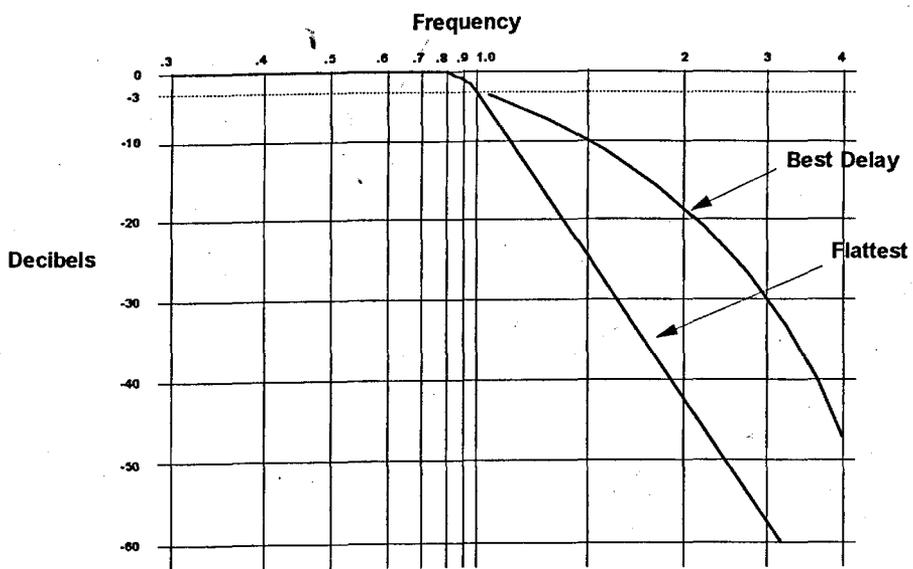


Figure 2.9 Thompson-Butterworth Filter Characteristics

A Compromise between the Best Delay and Flattest implementations was chosen, with an upper signal frequency limit of 2.5MHz and section design values of:-

<u>First Section</u>		<u>Second Section</u>		<u>Third Section</u>	
<i>Freq.</i>	<i>Damping</i>	<i>Freq.</i>	<i>Damping</i>	<i>Freq.</i>	<i>Damping</i>
1.268	1.945	1.301	1.521	1.382	0.711

This gives approximately 48dB attenuation at 7.5MHz (3-times f and $f_s/2$).

The component values for each stage are calculated as follows:-

$$C_1 = C \times \frac{D}{2}$$

$$C_2 = \frac{4}{D^2} \times C$$

$$R \times C = \frac{1}{2\pi \times f_n \times f_c}$$

f_n - Normalised frequency

f_c - Cut-off frequency

D - Damping

1st Stage

$$f_n = 1.268$$

$$D = 1.945$$

$$R \times C = \frac{1}{2\pi \times 2.5 \times 10^6 \times 1.268}$$

Given $\frac{D}{2} = 0.9725$ and $C_1 = C \times 0.9725$

Now an iterative process is used to find the best fit for preferred values of R, C₁ and C₂

let

$$C_1 = 27 \text{ pF} \quad \text{and} \quad C = \frac{27 \times 10^{-12}}{0.9725}$$

$$C_2 = \frac{4}{D^2} \times C_1 = 27.76 \times 10^{-12}$$

$$\approx 57 \text{ pF}$$

$$R = \frac{1}{2\pi \times 5 \times 10^6 \times 1.268 \times 27.76 \times 10^{-12}}$$

$$= 1\text{K}8\Omega$$

The component values for stages 2 and 3 are calculated in a similar way, giving the anti-alias filter circuit of Figure 2.10, with its corresponding frequency response characteristic depicted in Figure 2.11.

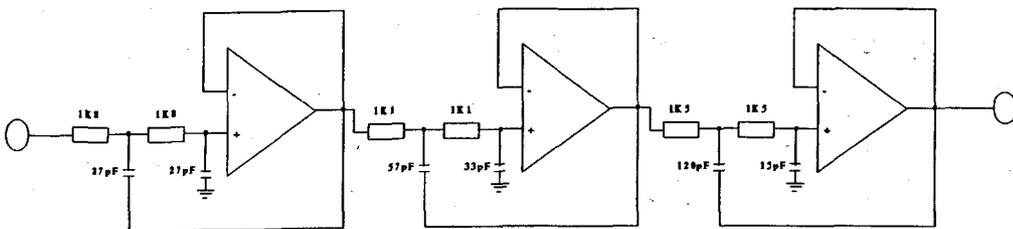


Figure 2.10 Anti-Alias Filter Circuit

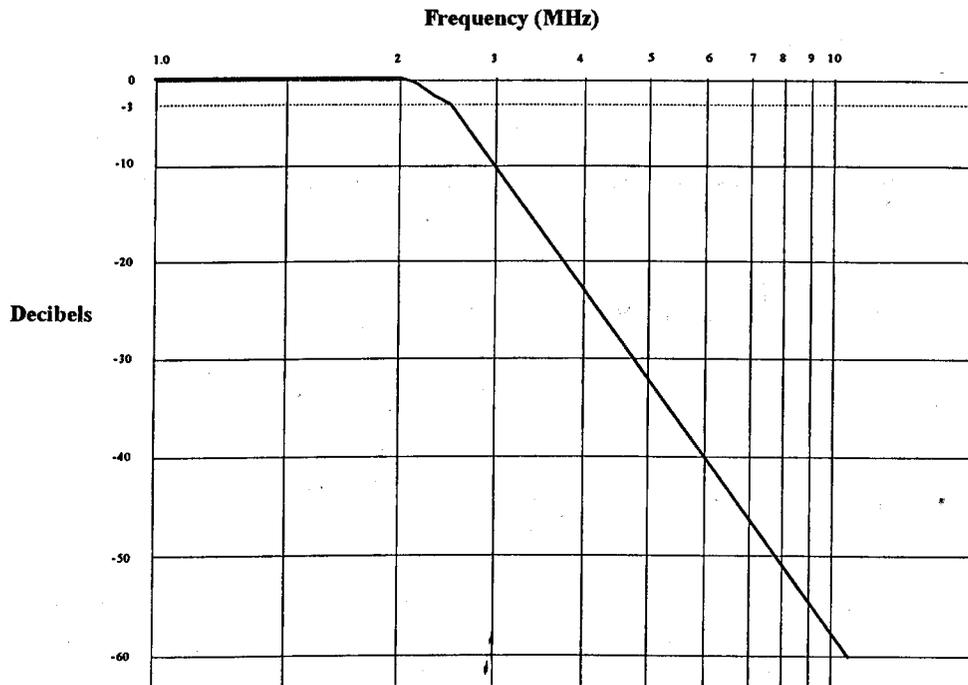


Figure 2.11 Anti-Alias Filter Frequency Response

2.2.4 Analogue to Digital Converter

The specification of the analogue to digital converter (ADC), more than any other single component, influences the overall system design, because the digitising speed dictates the necessary performance and size of the subsequent data storage module, as well as the bandwidth of the preceding amplifier and anti-alias filter. The resolution of the ADC, again, has a direct effect on the size of the data store, in addition to placing constraints on the system noise performance.

These decisions regarding the analogue to digital converter have important ramifications on both the cost and overall performance of the prototype system, so initially, a balance was reached between cost, performance and design difficulty, so ensuring that the prototype had adequate performance with the option of upgrading at a later stage.

The particular device chosen for the prototype system is a Micro Power Systems MP7684, CMOS, 8-bit, "flash" analogue to digital converter. The converter can operate at speeds up to 20MHz without the need for any sample and hold circuitry, thus simplifying the design. The ADC requires a unipolar analogue input, necessitating design of a circuit to add a DC offset of $\frac{1}{2}V_{ref}$ to the analogue input signal.

The circuit employed for the analogue to digital converter can be seen in Figure 2.12, along with the offset generator, which is incorporated into the anti-alias filter network.

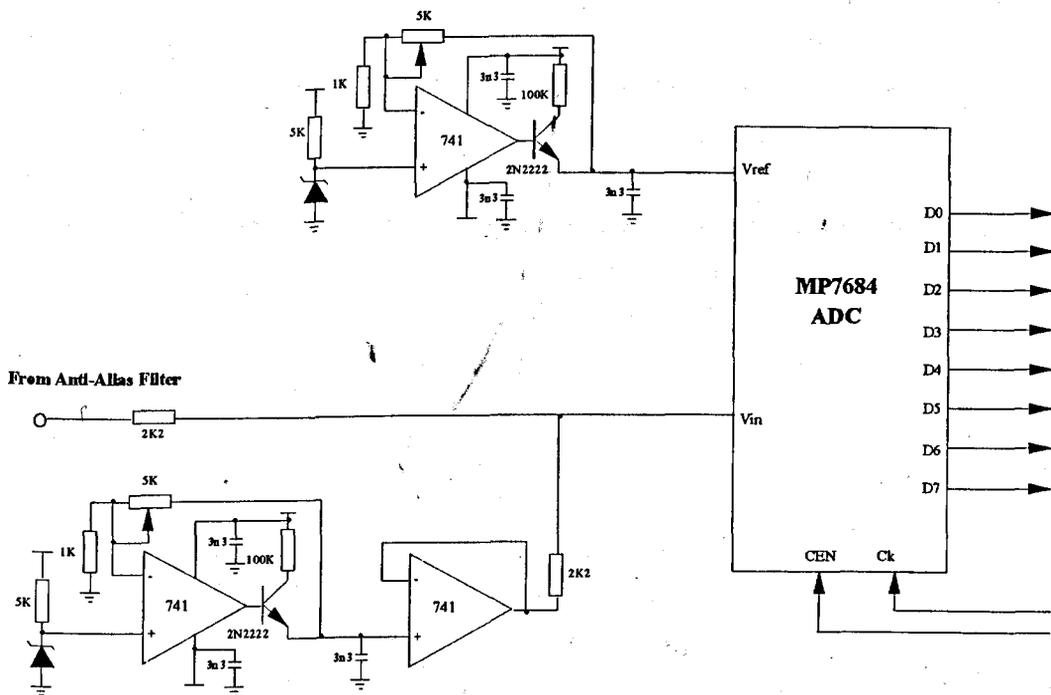


Figure 2.12 Analogue to Digital Converter Circuit

2.2.5 Hardware Controller & Data Storage Module

The roles of the hardware controller and primary storage modules are to oversee and control the sampling process, as well as providing a temporary storage facility for the digitised signal.

A data storage module is necessary due to the high sample rate of the system; a computer that is capable of a data transfer rate of 15MBytes/sec would be required if a hardware storage module were not to be used.

The initial design problems revolved around choosing the size of the storage facility. This had to be a trade-off between cost, speed and the amount of data that it had to hold. The main criteria influencing this decision was the need to provide enough data samples to produce meaningful results and facilitate removal of superimposed thyristor noise pulses, mentioned previously in section 2.1.2.

If the latter requirement is examined, bearing in mind that the noise is periodic in nature, but can occur anywhere in the mains cycle, then in order to uniquely identify the noise pulses a number of consecutive 50Hz mains cycles need to be captured. Assuming that three complete cycles are needed to unambiguously ascertain the position of, and consequently eliminate the thyristor interference, then it is a simple matter to estimate the required size of the data store.

Period of 3 x 50Hz waveforms	= 0.06 seconds
Number of samples in 3 cycles	= period x sample rate
	= 0.06 x 15x10 ⁶
	= 900 000 samples

This means that a large amount of Static RAM had to be employed to cope with the amount of data being produced by the analogue to digital converter. Real-time data compression methods were considered in order to reduce the memory overhead, however a minimum figure could not be arrived at for the success of any scheme, as it would be dependent upon the level of partial discharge activity at any given time. Also, the control circuitry to compress data at 15MBytes/second would be extremely complex.

The method chosen was to build a bank of 880KBytes (901120 bytes) of Static RAM to act as a temporary data store. In order to make this cost-effective, slow Static RAM had to be used, and therefore a novel architecture had to be conceived if the data conversion rate was to be maintained (Figure 2.13).

Each memory board is accessed once every four samples, in a "round robin" fashion, effectively quadrupling the memory bandwidth. Once the memory is full, the data can be transferred to the host computer for processing.

The entire sampling and subsequent transferring of data is controlled, at the hardware level, by dedicated control logic. This takes the form of an ASM (Algorithmic State Machine), and is implemented in a Programmable Logic Device.

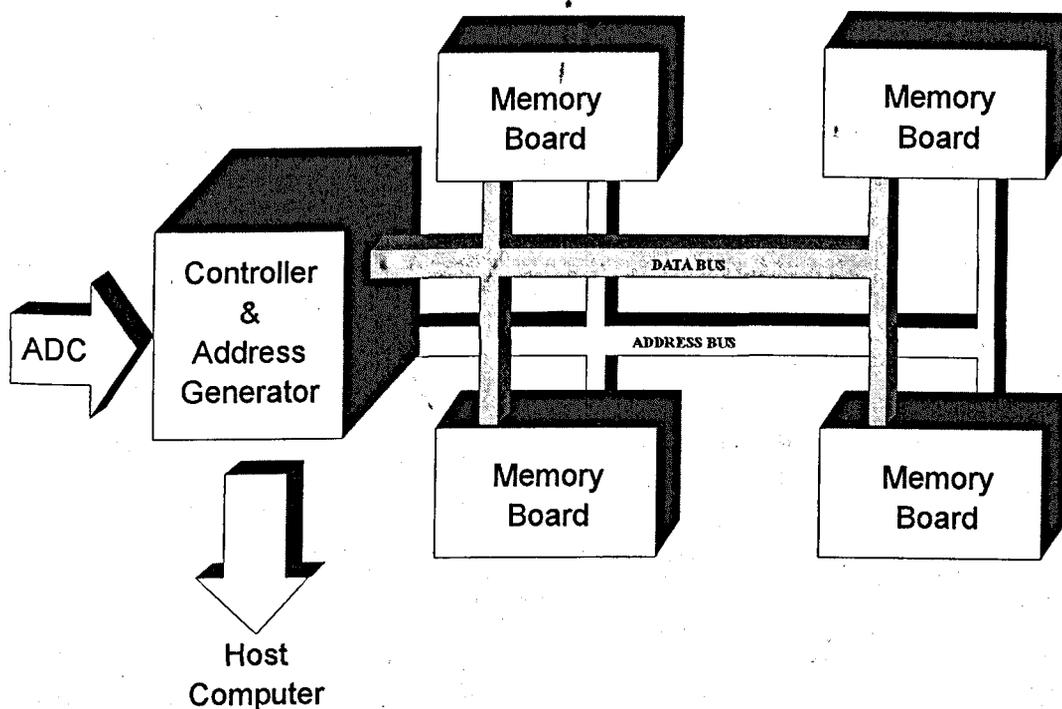


Figure 2.13 Memory Architecture

Use of an ASM, facilitates an algorithmic description of the controlling sequence, and hence, allows simple translation to the logic functions required to implement it. A flowchart of the controlling sequence can be seen in Figure 2.14.

The flowchart itself is quite meaningless until the functions of the terms used are explained.

- The outputs of the state machine are as follows:-
- The memory address generator can be cleared by setting its reset line low (CCLR in the ASM chart).
- The analogue to digital converter is activated by setting the signal ADCEN = 1.
- The computer is informed of a completed operation by the signal FINISH = 1.

There are four inputs to the controller:-

- The hardware is informed of a data sample request by using the S input.
- Transfer of the captured data is initiated using the R input.
- Z_C allows the controller to synchronise the sampling process to the 50Hz mains cycle by using a zero-cross detector, connected to the reference supply.
- The F input indicates that the memory is full, or that all of the data has been transferred.

Initially, the controller starts in the idle state (0). The address generator is clear and inactive. On reception of a 'sample' signal, S, it checks for initial synchronism with the reference 50Hz mains waveform; the zero-cross detector signal, Z_C, has to go HIGH then LOW for synchronism to be achieved. If Z_C is HIGH then the system proceeds to state 1, with the address generator still kept inactive. If Z_C subsequently goes LOW, synchronism with the 50Hz mains reference has been achieved and the data capture process can begin.

Data capture begins when the controller enters state 2. The analogue to digital converter is enabled and the address generator started (by removing the reset signal). Sampling proceeds until the hardware indicates that the memory is full, F, and the system then enters state 3, informing the computer that the operation is complete

(FINISH), and disabling the address generator. It stays in this state until the computer recognises the completed operation by removing the 'sample' command, S. The system then reverts back to the idle state (0).

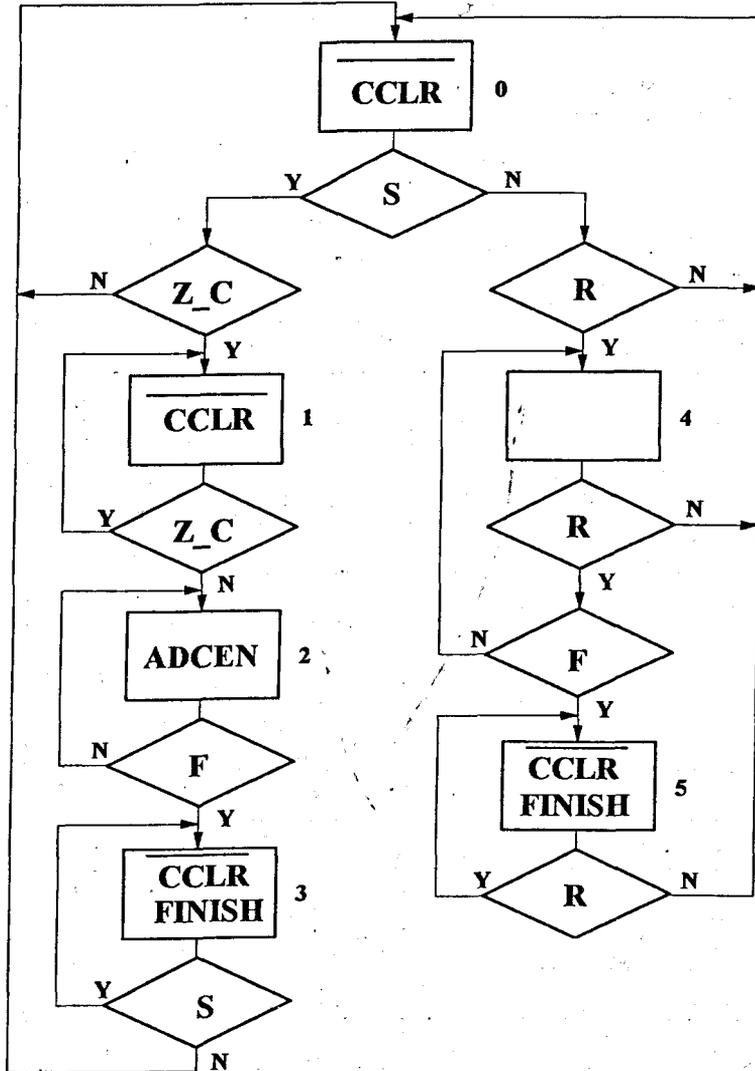


Figure 2.14 Controller Flowchart

The read cycle is very similar in operation, but without the synchronising aspect. When in the idle state (0), if the sample signal is LOW and the read signal, R, is set HIGH then the controller enters data transfer mode by proceeding to state 4. Here,

the address generator is activated and the transfer of data then proceeds until the read signal, R, is removed, or the hardware indicates that all of the data has been transferred (F). If the read signal is removed at any time, the controller simply reverts back to the idle state (0). If all of the data has been transferred, then the controller tells the computer that the operation has been completed (FINISH) and waits for the read signal to be removed before reverting back to the idle state.

The state transition table for the controller ASM chart can be seen in Table 2.1, and by allocating binary addresses to each state, a truth table can be found, Table 2.2.

Present State	Inputs				Next State	Outputs		
	S	R	Z C	F		CCLR	ADCEN	FINISH
0	0	0	X	X	0	0	0	0
0	1	X	0	X	0	0	0	0
0	1	X	1	X	1	0	0	0
0	0	1	X	X	4	0	0	0
1	X	X	0	X	2	0	0	0
1	X	X	1	X	1	0	0	0
2	X	X	X	0	2	1	1	0
2	X	X	X	1	3	1	1	0
3	0	X	X	X	0	0	0	1
3	1	X	X	X	3	0	0	1
4	X	0	X	X	0	1	0	0
4	X	1	X	0	4	1	0	0
4	X	1	X	1	5	1	0	0
5	X	0	X	X	0	0	0	1
5	X	1	X	X	5	0	0	1

X - Don't Care

Table 2.1 ASM State Transition Table

Present State			Inputs				Next State			Outputs		
Q ₂	Q ₁	Q ₀	S	R	Z_C	F	Q ₂₊	Q ₁₊	Q ₀₊	CCLR	ADCEN	FINIS H
0	0	0	0	0	X	X	0	0	0	0	0	0
0	0	0	1	X	0	X	0	0	0	0	0	0
0	0	0	1	X	1	X	0	0	1	0	0	0
0	0	0	0	1	X	X	1	0	0	0	0	0
0	0	1	X	X	0	X	0	1	0	0	0	0
0	0	1	X	X	1	X	0	0	1	0	0	0
0	1	0	X	X	X	0	0	1	0	1	1	0
0	1	0	X	X	X	1	0	1	1	1	1	0
0	1	1	0	X	X	X	0	0	0	0	0	1
0	1	1	1	X	X	X	0	1	1	0	0	1
1	0	0	X	0	X	X	0	0	0	1	0	0
1	0	0	X	1	X	0	1	0	0	1	0	0
1	0	0	X	1	X	1	1	0	1	1	0	0
1	0	1	X	0	X	X	0	0	0	0	0	1
1	0	1	X	1	X	X	1	0	1	0	0	1

X - Don't Care

Table 2.2 ASM Truth Table

Using the truth table, logic expressions can be found for the next state variables and the outputs, in terms of the present state variables and the inputs.

$$D_2 = \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot \overline{S} \cdot R + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot R + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot R$$

$$D_1 = \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot \overline{Z} + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} + \overline{Q_2} \cdot Q_1 \cdot Q_0 \cdot S$$

$$D_0 = \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot S \cdot Z + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot Z + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot F + \overline{Q_2} \cdot Q_1 \cdot Q_0 \cdot S + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot R \cdot F + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot R$$

$$CCLR = \overline{Q_2} \cdot Q_1 \cdot \overline{Q_0} + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0}$$

$$ADCEN = \overline{Q_2} \cdot Q_1 \cdot \overline{Q_0}$$

$$FINISH = \overline{Q_2} \cdot Q_1 \cdot Q_0 + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0$$

Using these logic expressions, the entire controller can now be implemented with flip-flops and NAND gates. The flip-flops hold the present state variables, whilst the logic determines what the next state will be, when the system is next clocked, according to the present state and the pattern of the inputs. The outputs are merely determined from the present state variables(i.e. a Moore machine). A diagram of the ASM controller can be seen in Figure 2.15.

The ASM controller is implemented in a Electrically Programmable Logic Device (EPLD), along with some other logic functions, whilst the time-critical functions, such as the address generator, are implemented using discrete 74F series devices. A schematic of the full contents of the EPLD can be seen in Figure 2.16, with Figure 2.17 showing the configuration of the 20-bit address generator.

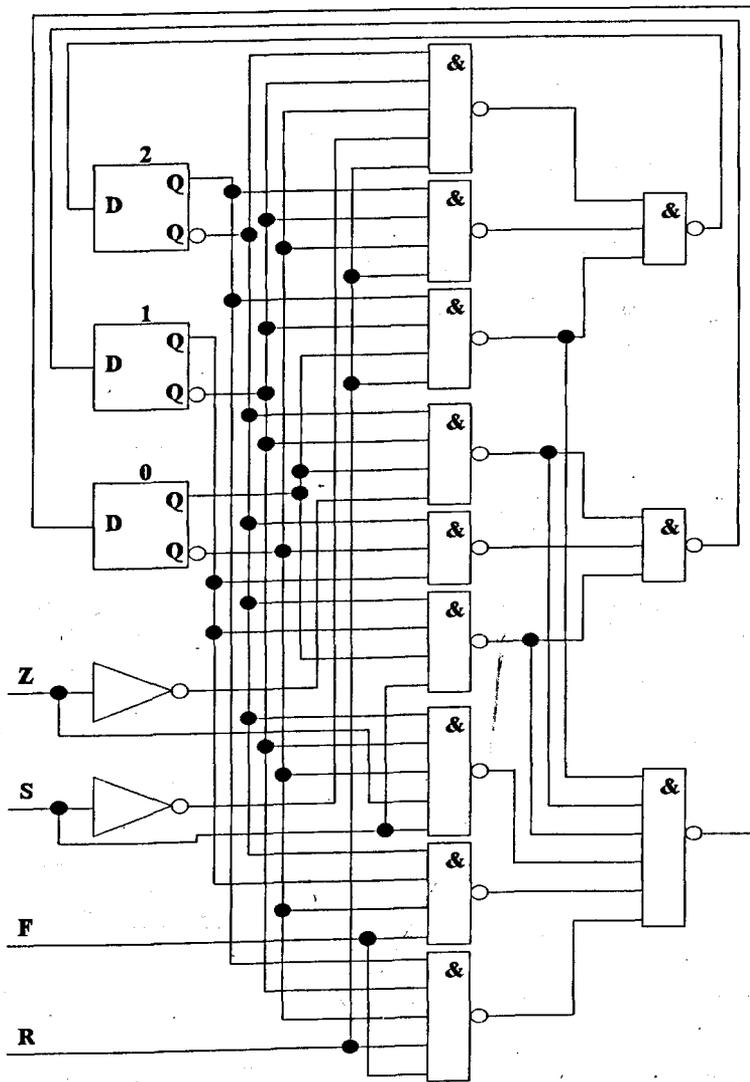


Figure 2.15 ASM Controller

The ASM controller is initialised to state 0 during power-on by attaching a simple power-on-reset network, consisting of a resistor and a capacitor, to the RESET line of the flip-flops.

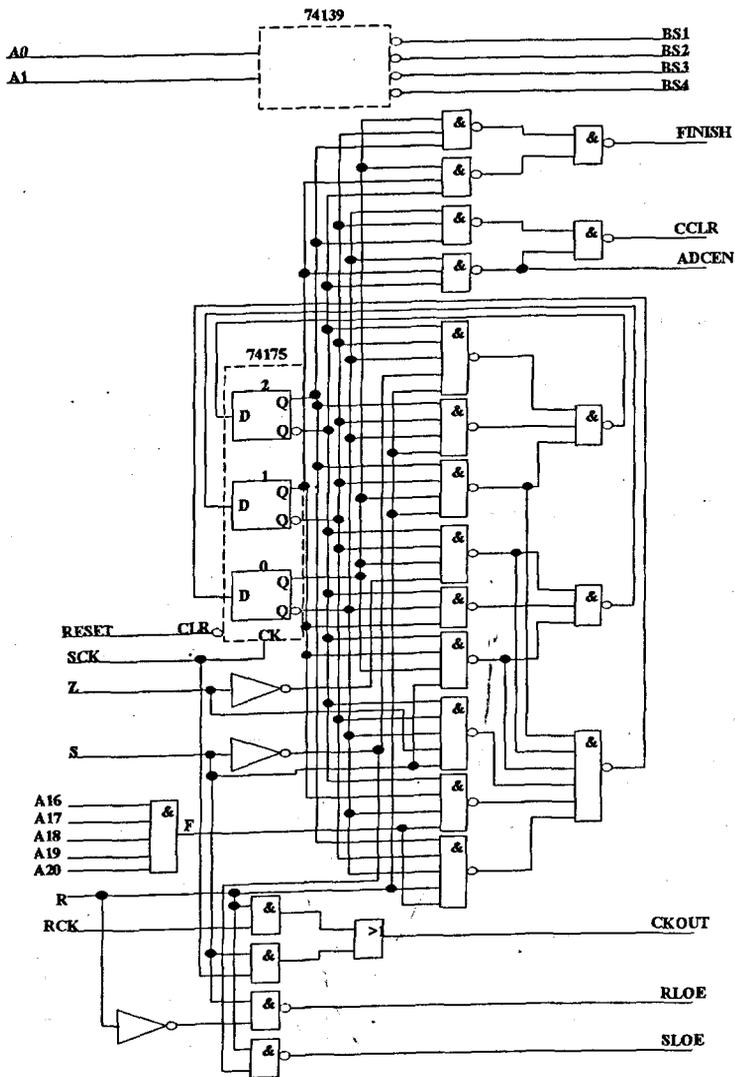


Figure 2.16 EPLD Contents

All of the logic designs were extensively simulated using Apollo DN3000 workstations in order to find any unforeseen problems in the design, so ensuring minimum changes during testing of the constructed hardware.

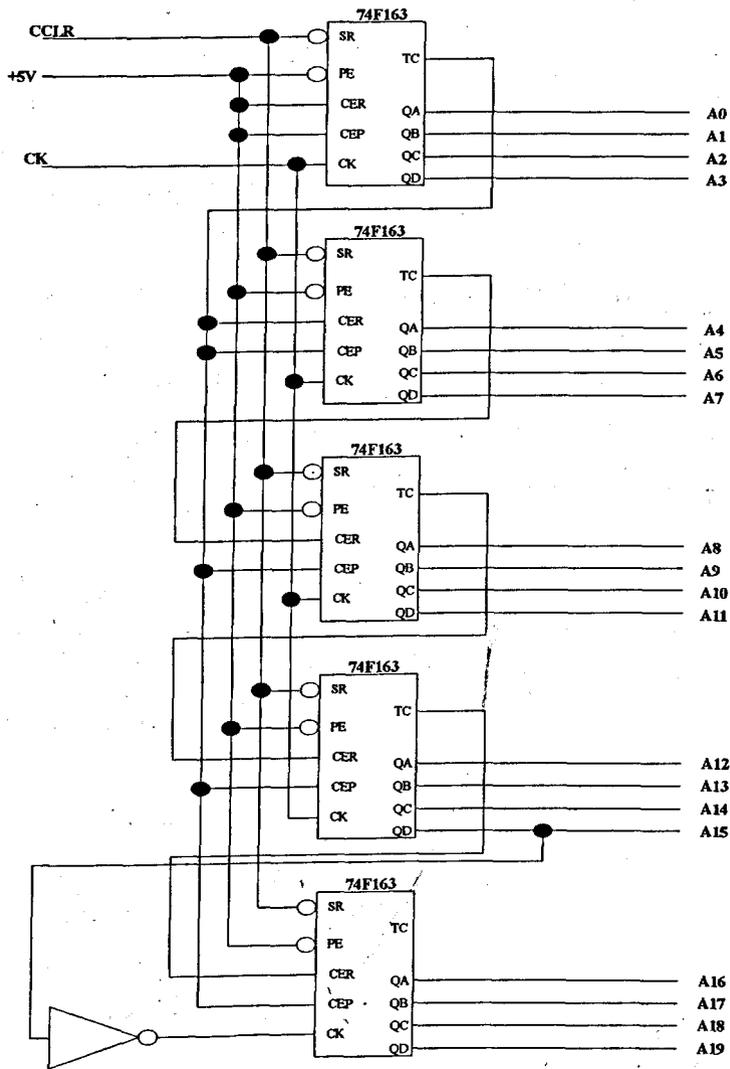


Figure 2.17 Address Generator

The upper 18 address bits are applied to each memory board, and latched when each board is enabled. The memory boards are each enabled in turn by decoding the lower 2 address bits, using a 2 to 4 decoder (implemented on the EPLD).

The five highest address bits are applied to a NAND gate to produce the F input for the ASM, indicating that the external memory boards are full.

The data acquisition and transfer cycles have their own clocks associated with each process; the sampling operation is carried out using a 15MHz clock, whilst data transfer to the host computer is carried out using a software generated clock signal that allows the computer to transfer data at its own rate. The ASM controller always uses the 15MHz hardware clock.

From the EPLD schematic, Figure 2.16, it can be seen that each clock signal can only be routed through to the output if it is accompanied by its appropriate enable signal, S or R. This ensures that the data acquisition clock is applied to the external hardware (i.e. the memory and address generator) only when a 'sample' operation is ongoing, and the transfer clock can only be applied during a 'read' operation.

Two enable signals for the data latches, of the analogue to digital converter and the computer interface, also need to be derived from the S and R inputs to ensure that there are no data collisions on the bus during an operation. Figure 2.16 also shows the simple logic implemented to achieve this.

Each memory board is basically the same. Each board consists of 8, 32K x 8-bit Static RAM ICs, various data and address latches and an EPLD to implement the miscellaneous logic functions necessary to generate the controlling signals. A schematic of a memory board can be seen in Figure 2.18, with the contents of the EPLD shown in Figure 2.19.

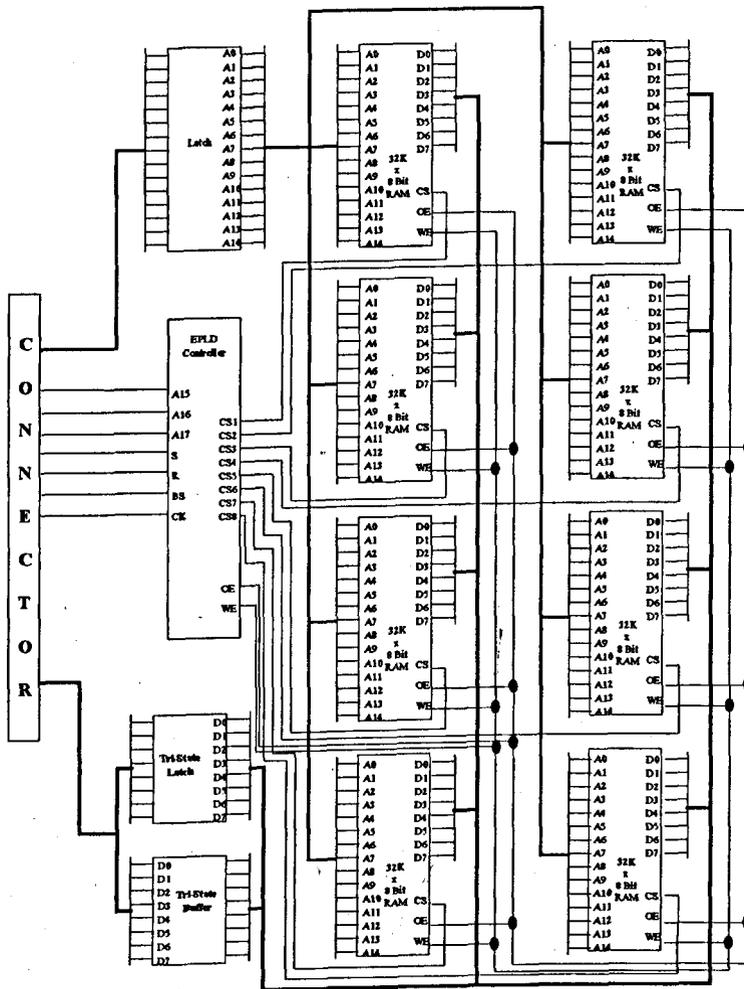


Figure 2.18 Memory Board Schematic

Each memory board only stores one sample out of four, thus both the data and the address need to be latched in order to give the memory time to complete its write cycle. Also during the 'read' cycle, only one memory board at a time can place data on the common bus, therefore a tri-state buffer (74F241) is required for use during the read cycle.

All of this means that various enable signals need to be generated. This is taken care of in the EPLD, Figure 2.19.

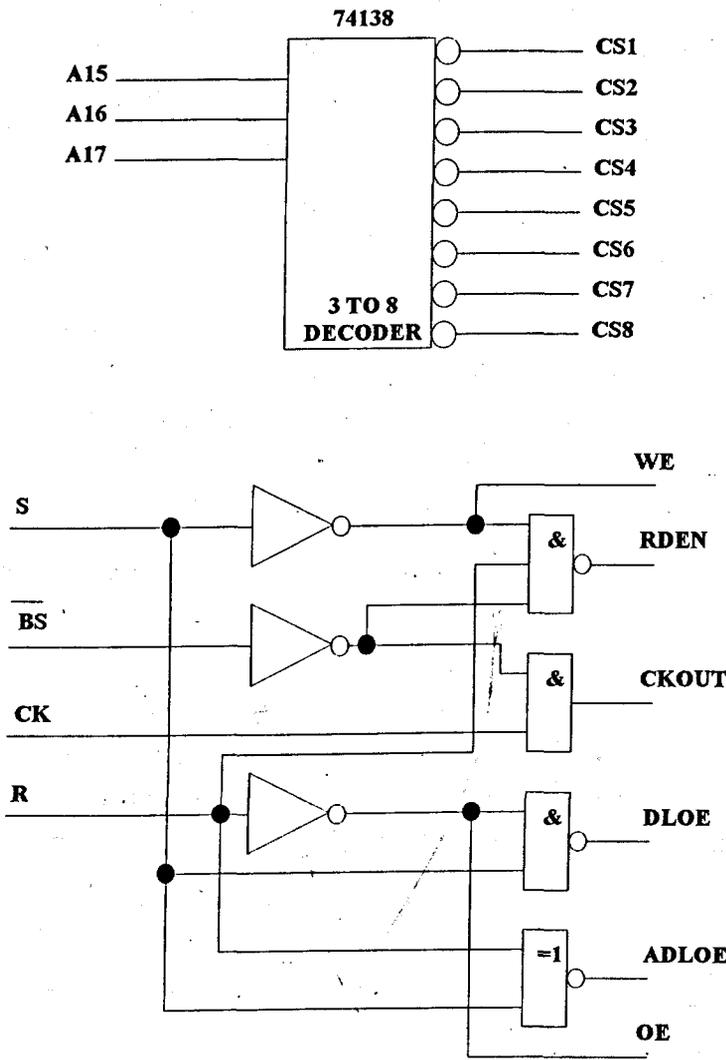


Figure 2.19 EPLD Contents

The clock used to latch both the address and the data is derived from the board enable signal and whatever clock signal is applied to the memory board from the controller. The address latch output enable signal (ADLOE) is simply the XOR of the 'read' and 'sample' signals; thus the address latch is active whenever one of the signals is active. The data latch output enable (DLOE), used during the data acquisition process, is active when the read signal (R) is LOW, with the sample signal (S) set HIGH. The enable signal for the tri-state buffer (RDEN), used during data transfer, is derived from the sample, read and board select signals.

The individual RAM IC enable signals are produced when the three highest address bits, supplied to the memory boards, are decoded using a 3 to 8 line decoder. The RAM write enable and output enable signals are simply the read and sample signals inverted (WE and OE are active LOW).

2.2.6 Host Computer System

The host computer system is the central component of the system in that it provides all of the operator interaction, controls data acquisition, processes the partial discharge information and displays the results for interpretation by the user.

A decision was made, early in the project, to use the IBM family of personal computers as the processing platform because of the large amount of extra hardware and software available for them.

The computer used during the development of the single-phase partial discharge monitoring system is an IBM AT type machine with EGA graphics and 1MByte of memory. The external hardware is interfaced to the computer using a digital I/O card, which occupies one expansion slot of the PC and provides 48 programmable input/output lines organised into 6 ports.

2.3 Summary

Investigations were carried out into the exact requirements of a new computer-based partial discharge monitoring system. This included experiments with a laboratory stator winding to determine the necessary performance of the amplification stages, as well as setting system parameters such as ADC sample rate and subsequent memory storage capacity. These parameters, once set, allowed detailed design of the various system modules to proceed and eventually combined into a coherent prototype.

The next stage in the prototype development involves examining the software design process necessary to operate the hardware and implement the data processing functions.

CHAPTER 3

SINGLE-PHASE DISCHARGE MONITORING SOFTWARE

3.1 Introduction

A comprehensive software project was undertaken to satisfy the initial objectives of the system. These objectives were:-

- Provide a user-friendly environment that could be used to test the existing system, and which could expand throughout the duration of the project.
- Incorporate different processing algorithms to extract useful information from the raw data.
- Provide the facility to display the results in graphical form for interpretation by the user, and provide easy operator interaction.

To accommodate all of these, the expanding menu system, shown in Figure 3.1, was devised.

The system can be split into three broad sections, each of which deal with a different aspect of the discharge monitor:

- Data acquisition and storage
- Waveform display
- Partial discharge pulse analysis

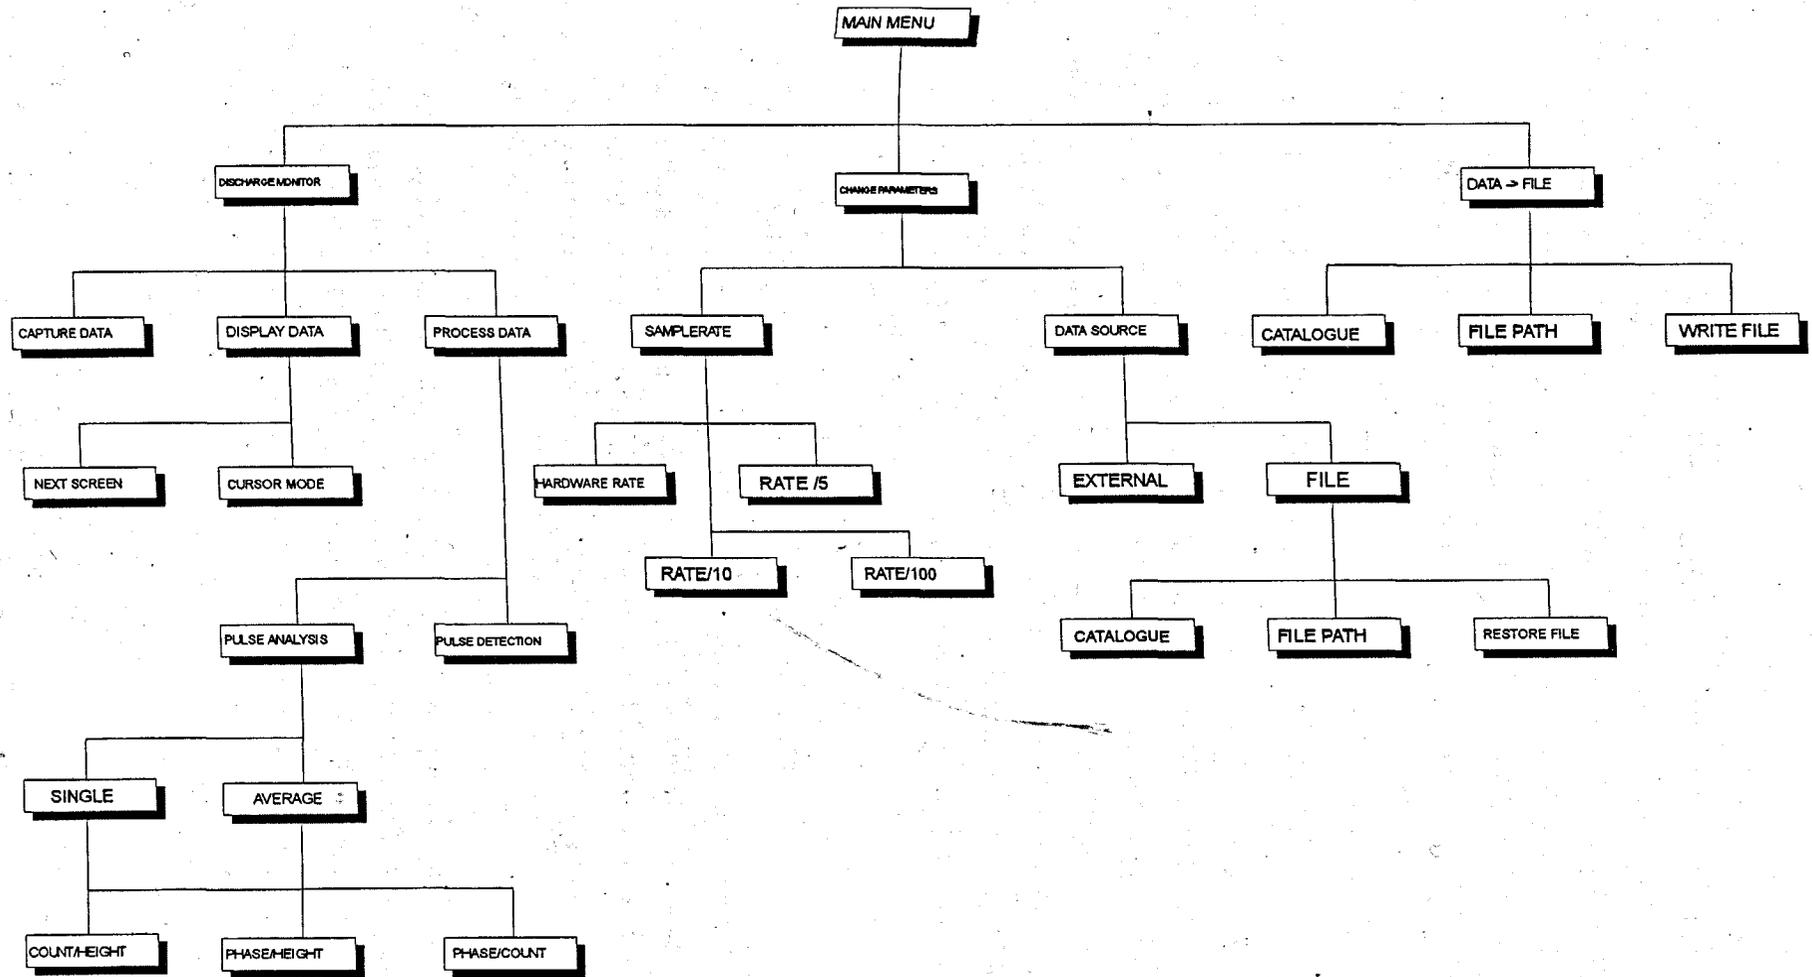


Figure 3.1 Menu Hierarchy

3.2 Data Acquisition and Storage

Before analysis of partial discharge activity can occur, the raw data must be made available to the host computer. In addition to this any instrumentation system must have a facility to permanently record the information for later processing or for comparison with subsequent measurements. This may take the form of a paper log, such as a chart or graph, or a direct record of the raw data, stored on magnetic tape or disk.

3.2.1 Controlling the Data Acquisition Hardware

Control of the custom-built hardware is achieved using a programmable interface card which plugs into an expansion slot in the computer. This card then effectively appears in the memory map of the computer and can be written to and read from as if it were standard memory locations. The card itself has 48 programmable input/output lines, organised into 6 ports, each of which has a number of addresses associated with them. One port was configured to read the data from the external hardware, whilst another was used to control the data acquisition process by activating the appropriate control signals outlined in the previous chapter.

3.2.2 Storing the Raw Partial Discharge Data

A facility had to be incorporated into the software to allow the user to permanently store the acquired data on computer disk. This would enable the operator to process the raw data at a later date, or use the recorded information to make a comparative analysis once more measurements have been taken or new processing routines added to the system.

The newly acquired data, still resident in the hardware, is read into the computer a byte at a time and then placed in a file on disk, the name of which is supplied by the operator. Each file stored uses approximately 900000 bytes of disk space, however, this space can be reduced by using standard file compression programs.

3.2.3 Retrieving Partial Discharge Data Files

Obviously if the facility is available to store acquired data in files, then a corresponding retrieval facility must also be provided. Here, the user has the option of switching between the data acquisition hardware and a previously stored file, for the primary data source. This means that they can acquire (and process) data using the external hardware, with that data being resident in the external hardware, and make direct comparisons with data stored in a file.

3.3 Waveform Display

The first step in the analysis process is to display the acquired data in graphical form. A typical display can be seen in Figure 3.2 (typical scale of microseconds and millivolts).

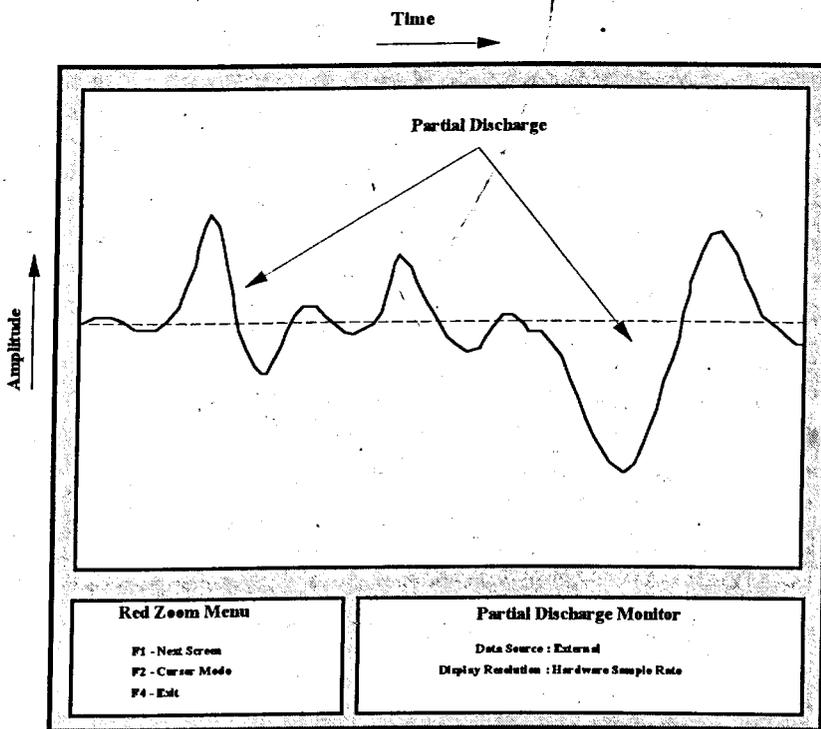


Figure 3.2 Typical Waveform display

The operator has certain parameters that he can change to alter the display. One parameter controls the effective displayed sample rate, by averaging a number of samples before plotting. This facility puts a larger portion of the acquired data record on the screen at one time allowing the user to quickly scan through the data record. The available choices are 5, 10 and 100 averages per displayed sample. Another parameter allows him to control the source of the data displayed on the screen, file or direct.

In addition to the standard display routine, the user can call up a cursor that tracks along the displayed data record, giving amplitude and time information, Figure 3.3.

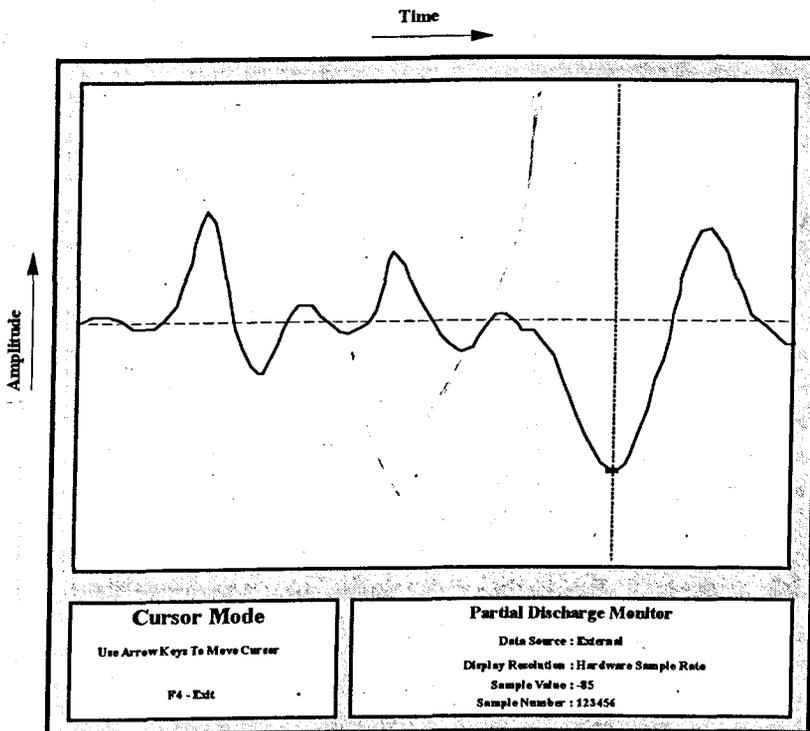


Figure 3.3 Tracking Cursor Display

3.4 Partial Discharge Pulse Analysis

Three different pulse analysis schemes were researched in order to gain information about the general characteristics of partial discharge activity within the machine under test. The first scheme to be developed was Pulse/Height Analysis. This entailed designing a software routine to identify, measure and count individual partial discharge pulses within the 900 000 data samples, and ultimately produce a plot of the results on the computer screen for operator interpretation. Figure 3.4 shows an typical Pulse Height Analysis plot for a machine with a large number of small discharges and a few large ones. Using this analysis option, a skilled operator can get a graphical indication of the severity of discharges, as well as their repetition rate for any given machine. Note that the range of pulse heights can only vary between -128 and +127, due to the 8-bit digitising hardware.

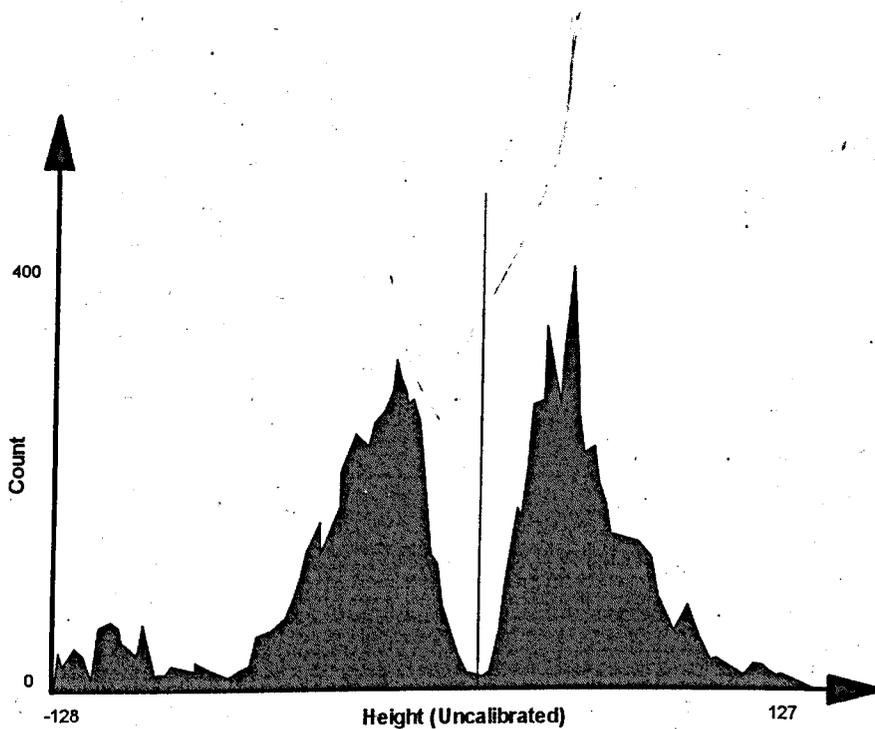


Figure 3.4 Pulse Height Analysis

The second analysis scheme developed was Phase/Height Analysis. This scheme involved identifying and measuring the height and overall phase of individual discharges within the period of the 50Hz supply fundamental. A typical plot of this can be seen in Figure 3.5, and shows the distribution of maximum positive and negative discharge pulses throughout the 20ms cycle period. Using this processing option, a skilled operator can see where the large discharges occur in the supply cycle, and can therefore spot any anomalous behaviour that may be indicative of some unusual discharge phenomena.

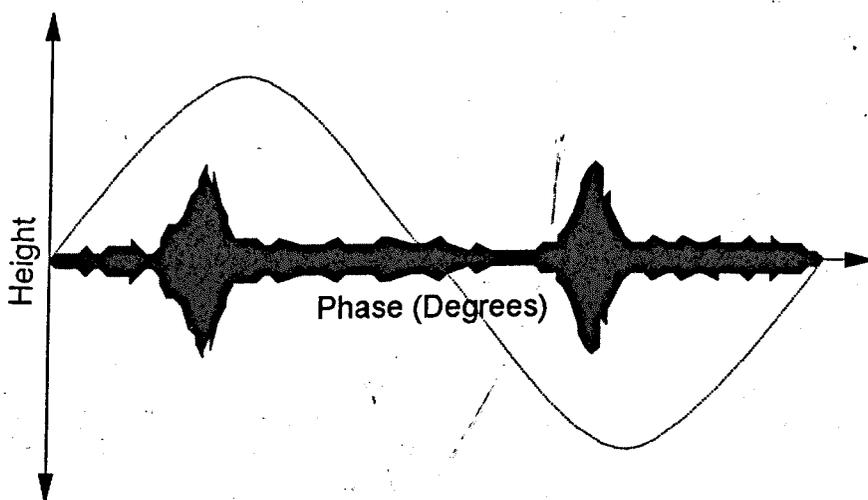


Figure 3.5 Phase/Height Analysis

The final analysis scheme developed was Phase/Count Analysis. This scheme involved splitting the 50Hz mains cycle into 180 phase "windows" of 2° angle, then counting and displaying graphically the number of discharges that occur in each. A typical plot of this can be seen in Figure 3.6, and shows the distribution of partial discharge activity throughout the 20ms cycle period. Using this processing option, a skilled operator can see where discharge activity is occurring within the mains cycle, and can therefore spot any anomalous discharge behaviour. Because this processing

option does not take into account the amplitude of the discharges, it is more useful in identifying repetitive low-amplitude discharges which may otherwise have been missed.

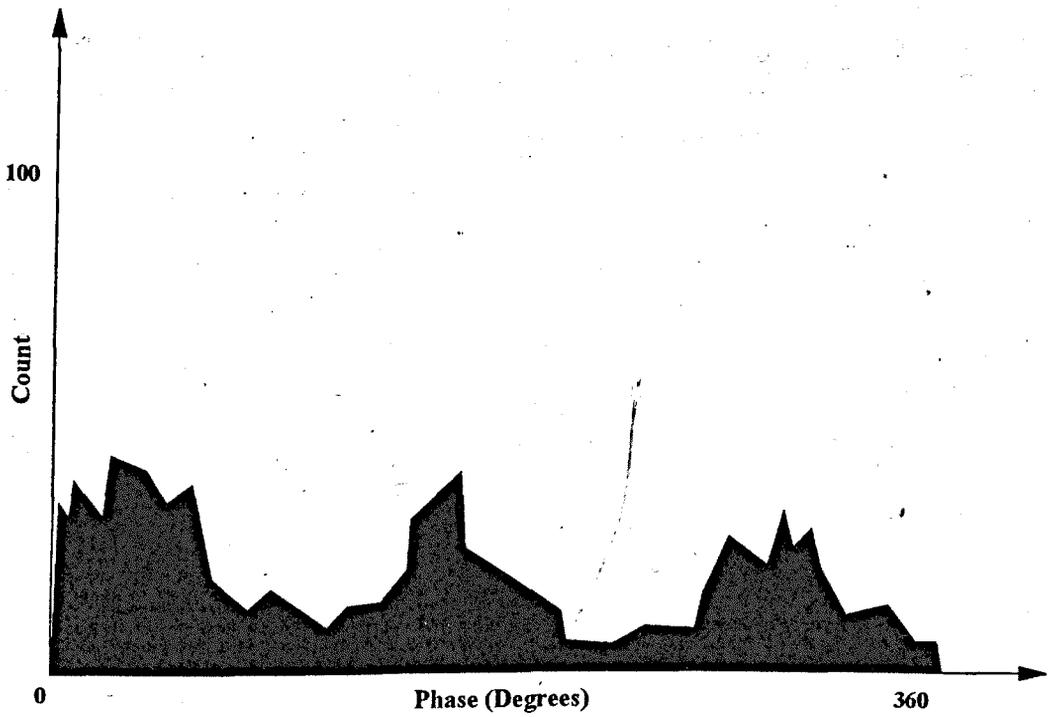


Figure 3.6 Phase/Count Analysis

3.5 Summary

An expanding menu system was developed that allows new features to be easily incorporated into its structure, as well as providing a basic user-friendly environment to interact with the operator.

The facilities described encompass data storage and retrieval, waveform display and rudimentary pulse processing algorithms. All of these facilities combined produce a prototype system that is capable of elementary discharge monitoring, for a single phase of a high voltage electrical machine.

CHAPTER 4

THREE-PHASE PARTIAL DISCHARGE

HARDWARE

4.1 Introduction

Having demonstrated the feasibility of a high speed data acquisition system for partial discharge processing, the next logical step was to enhance the system so that it could analyse the discharge properties of all three phases of a machine simultaneously. A three-phase monitoring system has a number of inherent advantages over a single-phase approach, such as the ability to detect different types of discharge phenomena, and a greater potential for noise rejection. A block diagram of the three-phase system can be seen in Figure 4.1.

4.2 System Hardware

The new system hardware differs considerably from its single phase counterpart in that the earlier system used four banks of discrete RAM Integrated Circuits to hold the discharge samples. These banks of memory were written to in a "round-robin" fashion, effectively quadrupling their bandwidth, making real-time data acquisition possible. With advances in memory technology, and corresponding price reductions, 256K fast static RAM ICs became available and have completely replaced the banks of discrete memory. This means that the data samples for one complete mains cycle can be stored in a single IC with no added complications, and a correspondingly large reduction in physical size. Each phase has its own data acquisition pcb with two additional PCBs, one controlling the process and interfacing to the computer, with the other containing the three phase amplifiers. Figure 4.2 shows how the system PCBs are arranged.

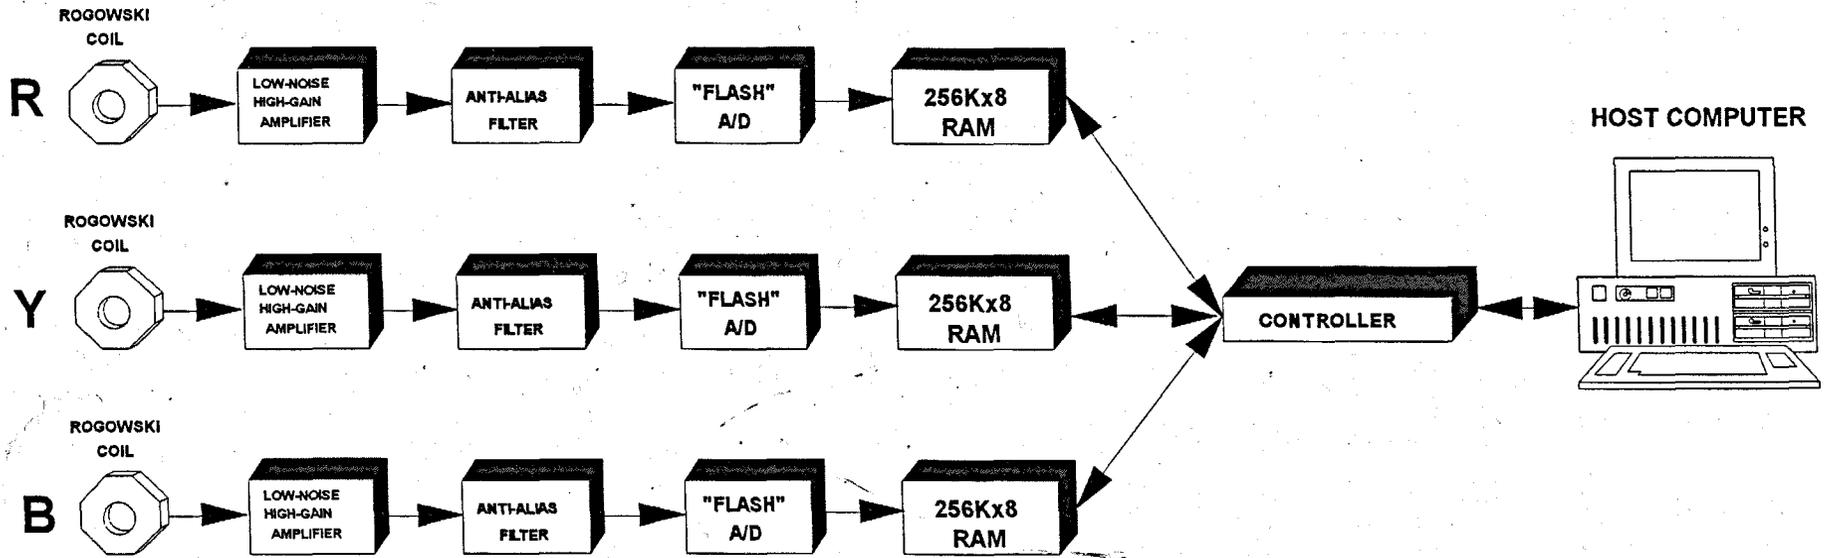


Figure 4.1 Three-Phase Monitoring System

A separate pcb is used for the amplifiers to enable them to be easily interchanged, depending on the amplification requirements - on line machine measurement, taped data analysis or laboratory machine experimental measurements.

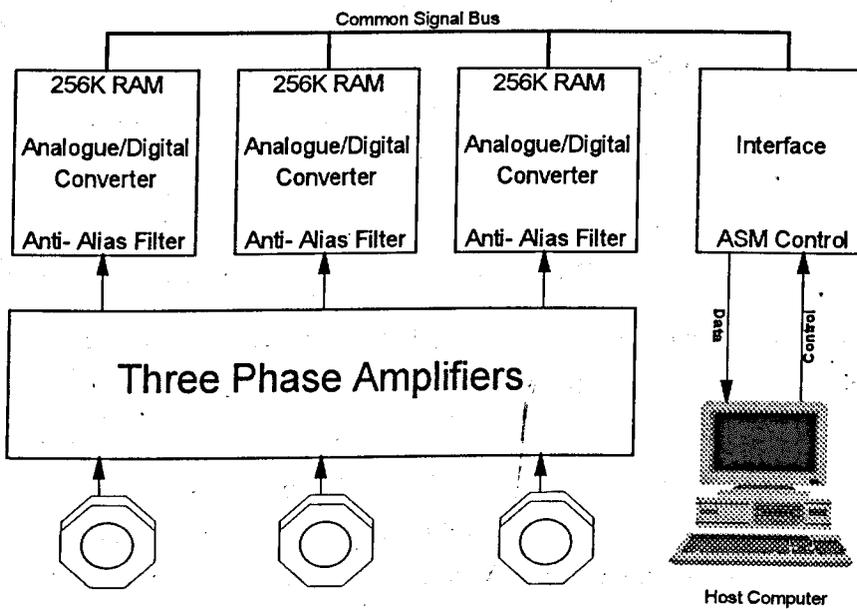


Figure 4.2 PCB Arrangement

The sample rate of the digitising process had to be reduced so that a complete digitised supply cycle (20ms) could fit into 256K of memory. It was accordingly decided to reduce the sample rate from 15MHz to 10MHz, which necessitated a routine alteration of the analogue processing stages. In addition to this, the change in architecture in the three-phase system demanded a radical re-design of the control logic. More stringent requirements were also placed on the control circuitry due to the faster memory write cycle (the write cycle had to be completed every sampling interval, as opposed to every 4 intervals in the single-phase system).

4.2.1 Three Phase Amplifier Configurations

A number of three-phase amplifier configurations were needed by the system; namely, a medium gain, low noise amplifier for use during experiments on the laboratory-based stator winding, a low gain amplifier for analysis of taped data recorded from on-line machines, and finally a high gain, low noise amplifier to use for obtaining data directly from on-line machines. Only the first two amplifier configurations needed to be developed because on-line measurement opportunities were never available.

Figure 4.3 shows the amplifier configuration used on one phase of the system for experiments on the laboratory machine; the other two phases are exactly the same. The gain of this amplifier stage is approximately 100. The reason that the gain does not need to be very high is that the simulated discharge pulses that are applied to the laboratory stator are of high amplitude.

Each phase monitor has its own dedicated analogue to digital converter, so accordingly each phase needs a dedicated anti-alias filter network. The filter networks employed are similar to the one used in the single-phase hardware, but with different component values to reflect the lowering of the system sample rate. It will serve no useful purpose to re-iterate the design process used for the anti-alias filter networks, suffice it to say that each was designed for a cut-off frequency of 1.5MHz.

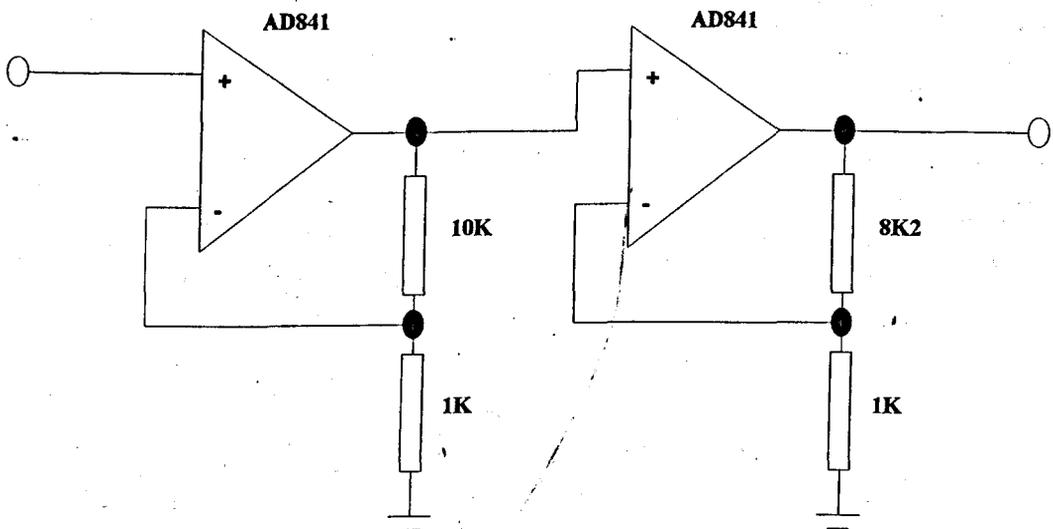


Figure 4.3 Amplifier Configuration

4.2.2 Hardware Controller

The hardware controller for the system commands each single-phase data acquisition module simultaneously and provides the interface to the host computer. Figure 4.4 shows a block diagram of the controller.

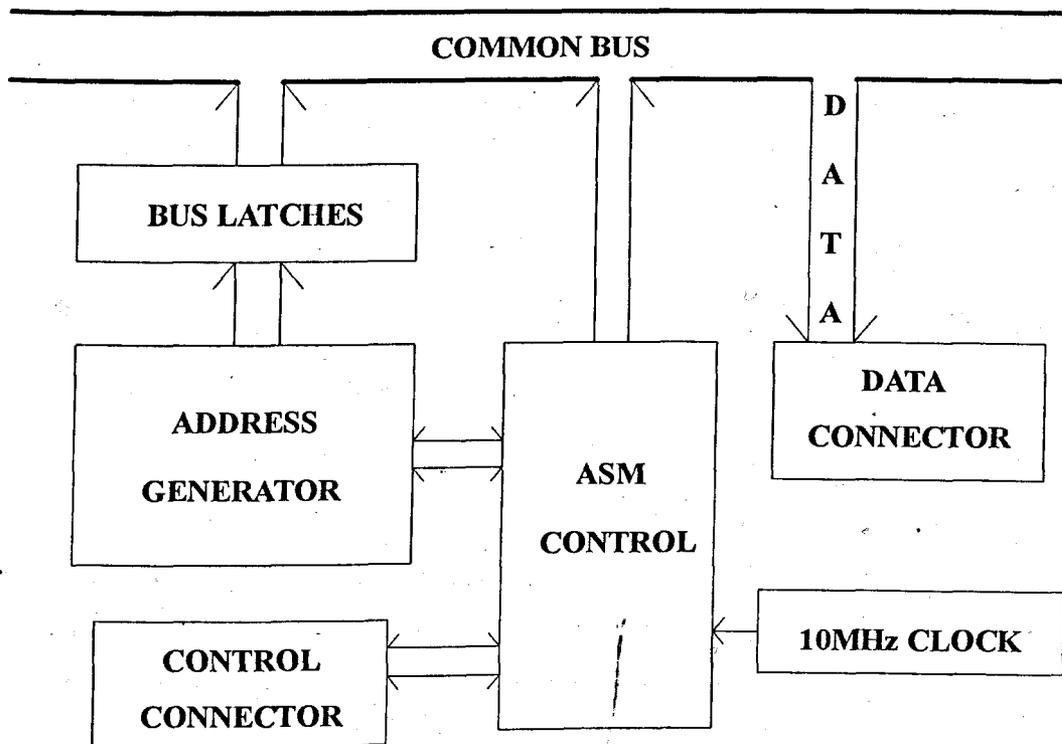


Figure 4.4 Controller Block Diagram

There are three connections to the hardware controller pcb, one to receive command signals from the computer, another to send the captured three-phase data back to the computer and finally a connection to the common bus system connecting all of the hardware together.

In addition to the usual latches and buffers required for interfacing purposes, the majority of the discrete logic devices used in the hardware controller are employed in the address generator. The address generator, as its name suggests, has to provide valid address information for the individual banks of RAM used to hold the digitised data. Each phase module has its own A/D converter and associated RAM, therefore each address generated can be applied to all three banks of memory simultaneously. This means that the address generator needs only to be capable of producing 256K addresses - 18 bits. This is achieved using five 74F163 binary counter ICs.

Again, central to the controlling function is an algorithmic state machine, along the same lines as before but with more output signals, which are used by discrete logic devices in the data acquisition process. The ASM flowchart for the system can be seen in figure 4.5.

The outputs of the state machine are as follows:

- The memory address generator can be cleared by setting its reset line low (ACLR in the ASM chart).
- The analogue to digital converters are all activated when the ADC signal is set high.
- The internal 10MHz data sampling clock is gated through to the address generator and acquisition PCBs when SEN is set high.
- The clock generated by the host computer, used to transfer the data, is gated through to the address generator and memory banks when REN is high.
- The data latches used to hold the output from the A/D converters are enabled when SL is low.
- Data can be read out of the memory banks when the OE line is set low.
- The data latches used to hold the output data from the memory are enabled when RL is low.
- The memory ICs are enabled with a low CS signal.
- The computer is informed of a completed operation when the signal FINISH is set high.

The inputs to the controller are the same as before:

- The hardware is informed of a data sample request using the S input.
- Transfer of the captured data is initiated using the R input.
- The sampling process is synchronised to the 50Hz supply fundamental by using a zero-cross detector which provides the Z_C signal.
- The F input indicates that the memory is full, or that all of the data has been transferred.

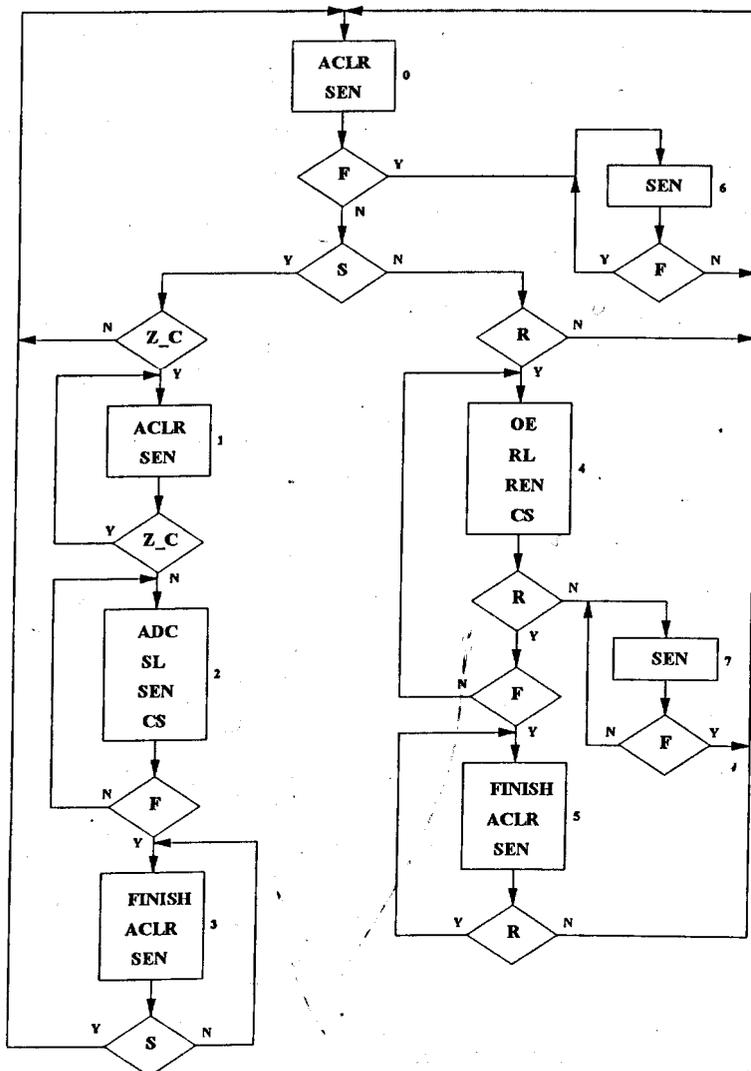


Figure 4.5 ASM Flowchart

Before the operation of the ASM chart is discussed, a small problem with the address generator has to be highlighted. This problem only came to light during testing of the single-phase system and whilst not serious, proved inconvenient. The problem arose with the particular configuration of binary counters needed to produce a synchronous count of 20 bits using a 15MHz clock. The configuration, shown in Figure 4.6, required that the IC used to provide the 4 most significant address bits was not connected to the system clock, but rather the MSB of the next lower counter. This meant, in practice, that when the reset signal was applied to the address generator, the four lower counters reset but the counter producing the 4 most significant address

lines did not, as there was no transition on its clock input. This failure to reset the four highest address bits meant that the signal indicating that the memory was full, F, would still be set, as it is derived from the most significant address bits. Thus under certain circumstances the controller would be fooled into thinking that the memory was full or that all of the data had been transferred to the host computer.

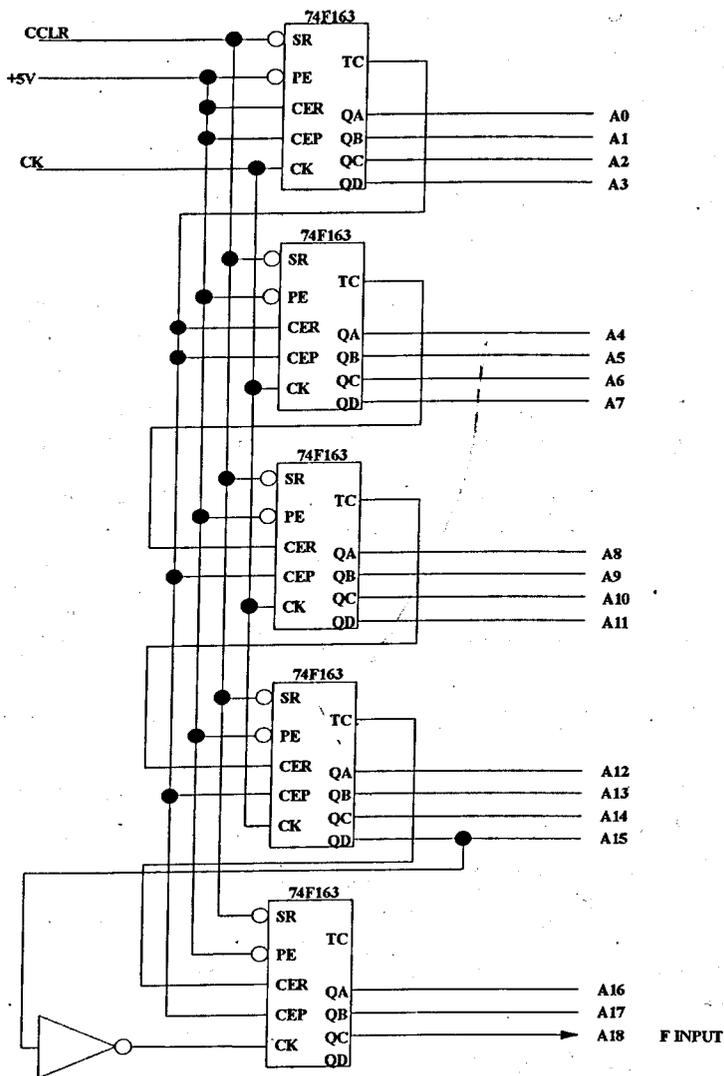


Figure 4.6 Address Generator

The three-phase system requires 18 bit addresses, so five counters still needed to be used, with the same problem resulting. The solution to this problem is implemented in the ASM chart and uses the fact that the F input is attached to the 19th bit of the

address generator (only 18 bits are needed, but the generator is capable of producing 20). The ASM checks for address bit 19 being set, and if it is set then the address generator is clocked until a transition occurs to reset it. This prevents the ASM receiving false information about the state of the memory.

Initially, the controller starts in the idle state (0). The address generator should be clear and inactive. If not, then it is made clear and inactive. On reception of a 'sample' signal, S, it checks for initial synchronism with the reference 50Hz mains waveform: the zero-cross detector signal, Z_C, has to go HIGH then LOW for synchronism to be achieved. If Z_C is HIGH then the system proceeds to state 1, with the address generator still kept inactive. If Z_C subsequently goes LOW, synchronism with the 50Hz mains reference has been achieved and the data capture process can begin.

Data capture begins when the controller enters state 2. The analogue to digital converters are enabled, along with their associated latches and memory, the address generator is also started (by removing the reset signal). Sampling proceeds until the hardware indicates that the memory is full, F, and the system then enters state 3, informing the computer that the operation is complete (FINISH), and disabling the address generator. It stays in this state until the computer recognises the completed operation by removing the 'sample' command, S. The system then reverts back to the idle state (0).

The read cycle is very similar in operation, but without the synchronising aspect. When in the idle state (0), if the sample signal is LOW and the read signal, R, is set HIGH then the controller enters data transfer mode by proceeding to state 4: here, the address generator is activated, also each bank of memory and associated latches are enabled and the computer-generated clock gated through. The transfer of data then proceeds until the read signal, R, is removed, or the hardware indicates that all of the data has been transferred (F). If the read signal is removed at any time, the controller simply reverts back to the idle state (0), after resetting the address generator. If all of the data has been transferred, then the controller tells the computer that the operation has been completed (FINISH) and waits for the read signal to be removed before reverting back to the idle state.

The truth table for the algorithmic state machine can be seen in Table 4.1.

Using the truth table, logic expressions can be found for the next state variables and the outputs, in terms of the present state variables and the inputs.

$$D_2 = \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot F + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot R \cdot S \cdot \overline{F} + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot R + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot F + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot \overline{F}$$

$$D_1 = \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot F + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot \overline{Z} + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot S + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot R + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot F + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot \overline{F}$$

$$D_0 = \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot S \cdot Z \cdot \overline{F} + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot Z + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot F + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot S + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot R + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot R \cdot F + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot R + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot \overline{F}$$

$$= \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot S \cdot Z \cdot \overline{F} + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot Z + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot F + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot S + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot R + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} \cdot F + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot R + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 \cdot \overline{F}$$

$$ACLR = \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0$$

$$= \overline{Q_2} \cdot \overline{Q_1} + \overline{Q_2} \cdot \overline{Q_0} + \overline{Q_1} \cdot \overline{Q_0}$$

$$ADC = \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0}$$

$$REN = \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0}$$

$$SEN = \overline{REN}$$

$$\overline{SL} = \overline{ADC}$$

$$\overline{OE} = \overline{REN}$$

$$\overline{RL} = \overline{REN}$$

$$\overline{CS} = \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 + \overline{Q_2} \cdot \overline{Q_1} \cdot \overline{Q_0} + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0$$

$$= \overline{Q_2} \cdot \overline{Q_1} + \overline{Q_2} \cdot \overline{Q_1} + \overline{Q_0}$$

$$FINISH = \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0 + \overline{Q_2} \cdot \overline{Q_1} \cdot Q_0$$

Using these logic expressions, the entire controller can now be implemented with flip-flops and NAND gates. The flip-flops hold the present state variables, whilst the logic determines what the next state will be, when the system is next clocked, according to the present state and the pattern of the inputs. The outputs are merely determined from the present state variables

PRESENT STATE			INPUTS				NEXT STATE			OUTPUTS								
A2	A1	A0	R	S	Z_C	F	A2+	A1+	A0+	ACLR	ADC	SEN	REN	SL	OE	RL	CS	FINISH
0	0	0	X	X	X	1	1	1	0	0	0	1	0	1	1	1	1	0
0	0	0	0	0	X	0	0	0	0	0	0	1	0	1	1	1	1	0
0	0	0	X	1	0	0	0	0	0	0	0	1	0	1	1	1	1	0
0	0	0	X	1	1	0	0	0	1	0	0	1	0	1	1	1	1	0
0	0	0	1	0	X	0	1	0	0	0	0	1	0	1	1	1	1	0
0	0	1	X	X	0	X	0	1	0	0	0	1	0	1	1	1	1	0
0	0	1	X	X	1	X	0	0	1	0	0	1	0	1	1	1	1	0
0	1	0	X	X	X	0	0	1	0	1	1	1	0	0	1	1	0	0
0	1	0	X	X	X	1	0	1	1	1	1	1	0	0	1	1	0	0
0	1	1	X	0	X	X	0	0	0	0	0	1	0	1	1	1	1	1
0	1	1	X	1	X	X	0	1	1	0	0	1	0	1	1	1	1	1
1	0	0	0	X	X	X	1	1	1	1	0	0	1	1	0	0	0	0
1	0	0	1	X	X	0	1	0	0	1	0	0	1	1	0	0	0	0
1	0	0	1	X	X	1	1	0	1	1	0	0	1	1	0	0	0	0
1	0	1	0	X	X	X	0	0	0	0	0	1	0	1	1	1	1	1
1	0	1	1	X	X	X	1	0	1	0	0	1	0	1	1	1	1	1
1	1	0	X	X	X	0	0	0	0	1	0	1	0	1	1	1	1	0
1	1	0	X	X	X	1	1	1	0	1	0	1	0	1	1	1	1	0
1	1	1	X	X	X	0	1	1	1	1	0	1	0	1	1	1	1	0
1	1	1	X	X	X	1	0	0	0	1	0	1	0	1	1	1	1	0

Table 4.1 ASM Truth Table

The logic expressions produced for the next state variables could be further simplified. However, there are two reasons for not simplifying them further: Firstly, there are no real advantages in simplifying the expressions manually, as when they are implemented in the EPLD, the design software automatically reduces the logic expressions. The second, and main, reason for not further simplifying the logic expressions is one of future design change, in that if any changes were made to the functionality of the ASM, the implementation of these changes in the EPLD could be accomplished easily, simply by moving a few NAND gate connections on the EPLD CAD software.

A diagram of the ASM controller can be seen in Figure 4.7.

The ASM controller is implemented in a Electrically Programmable Logic Device (EPLD), along with some other logic functions, whilst the time-critical functions, such as the address generator, are implemented using discrete 74F series devices.

In addition to the four inputs to the ASM, the EPLD has two input clock lines: the hardware clock, SCK, used by the ASM logic and the rest of the system when acquiring data, and the computer-generated clock, RCK, used when transferring data to the host computer for analysis. There is also a reset line to initialise the system when first switched on; this is externally connected to a power-on-reset circuit.

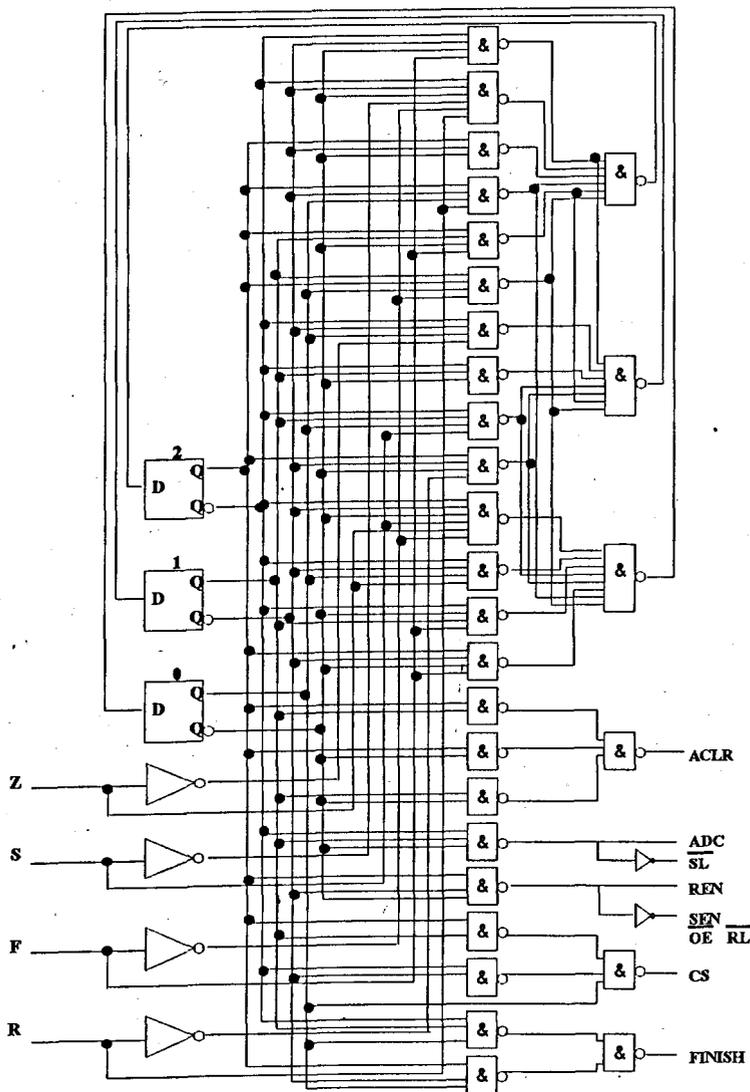


Figure 4.7 ASM Controller

4.3 Summary

Enhancing the prototype system to produce one capable of simultaneously monitoring all three phases of a machine under test was no easy task: the memory architecture had to be completely re-designed, with accompanying ramifications for the controlling hardware. Lessons learned during the development of the initial system were put to good use here, resulting in the new system having no discernible hardware bugs.

CHAPTER 5

THREE-PHASE MONITORING SOFTWARE

5.1 Introduction

Considerable changes had to be made to the existing software to enable it to fully utilise the new three-phase monitoring system. In addition to this, the host computer's graphics capability was upgraded from EGA to VGA, the higher resolution allowing more information to be displayed on the screen at one time.

The basic hierarchy of the software remained the same however, most notably, all of the display functions had to be updated in order to show the three phase information.

An expanding modular-based software hierarchy was developed early in the project, to allow easy upgrading of individual routines as well as facilitating the continual addition of new processing and display options. A diagram of the hierarchy can be seen in Figure 5.1.

From the diagram it can be seen that the hierarchy is in three main sections; a section that actually performs all of the processing and display functions, a section that allows the user to alter certain important parameters in the processing and display functions, and a final section that allows the captured partial discharge data to be stored as a file for later use.

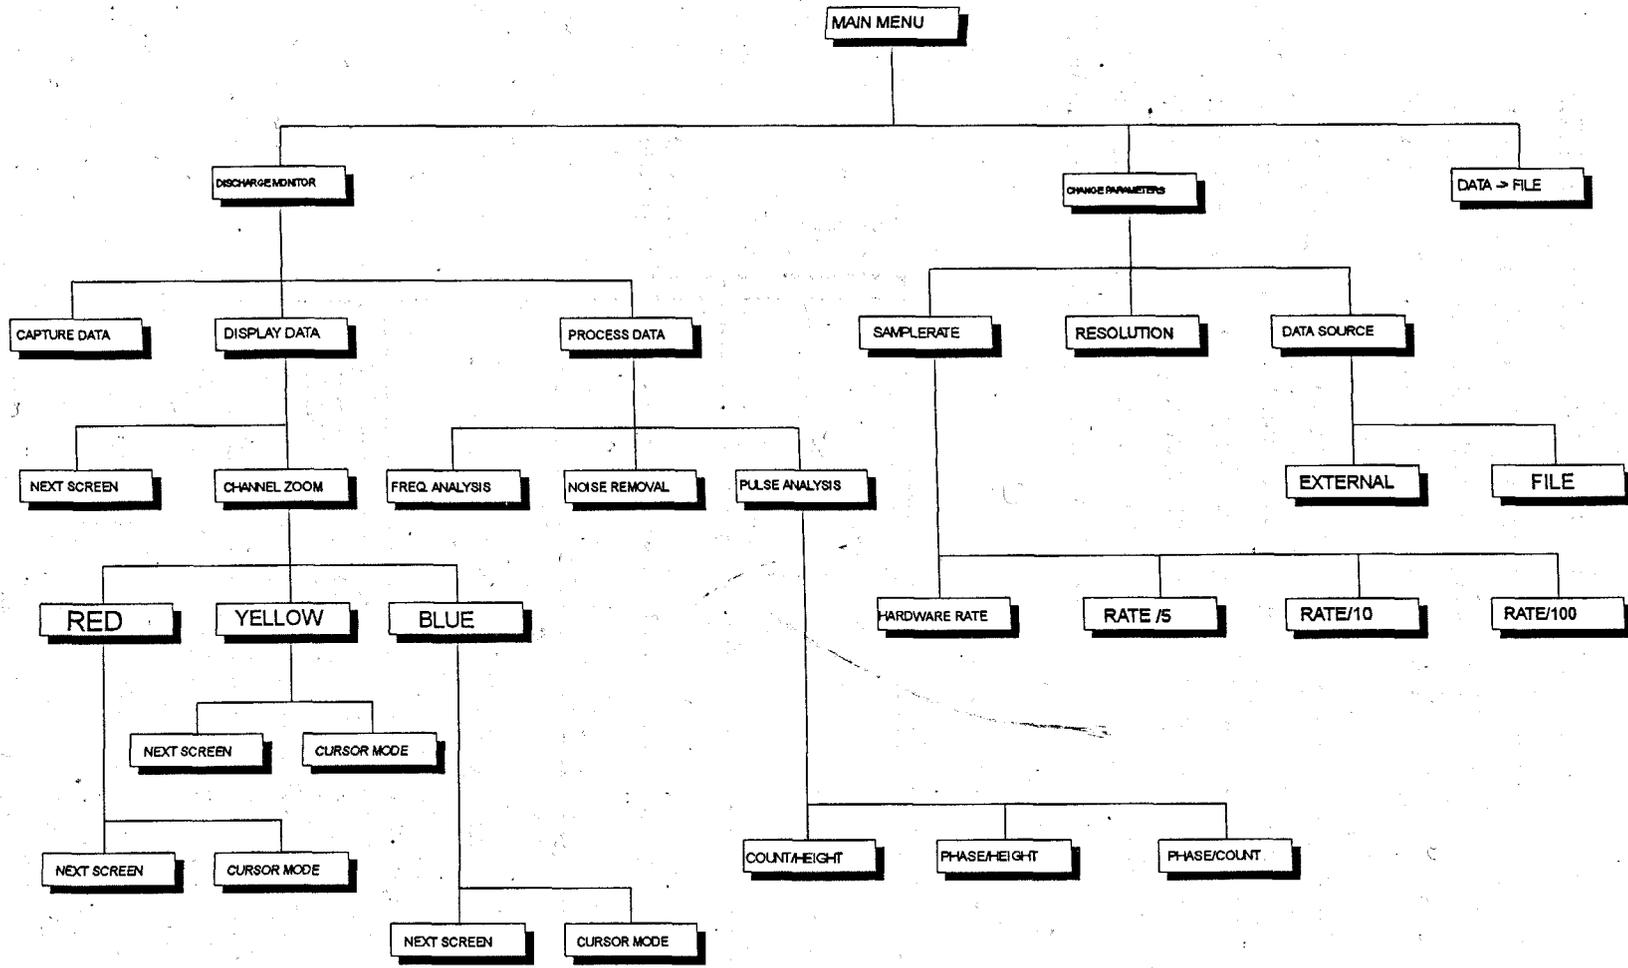


Figure 5.1 Menu Hierarchy

5.2 General Display and Pulse Analysis

5.2.1 General Display Functions

There are a number of display options that the operator can choose from, once the discharge information has been captured by the hardware or retrieved from a file. Figures 5.2, 5.3 and 5.4 show the sort of options available.

Firstly the operator can view the "raw" data for each phase on the screen, Figure 5.2. Only a small portion of the digitised information can be viewed at any one time, but using the facility to alter the display resolution ("sample rate"), the operator can obtain a compressed picture of the data. With this display option a skilled operator can scan the data, directly comparing all three phases visually, to get an overall idea of the sort of discharges that are occurring.

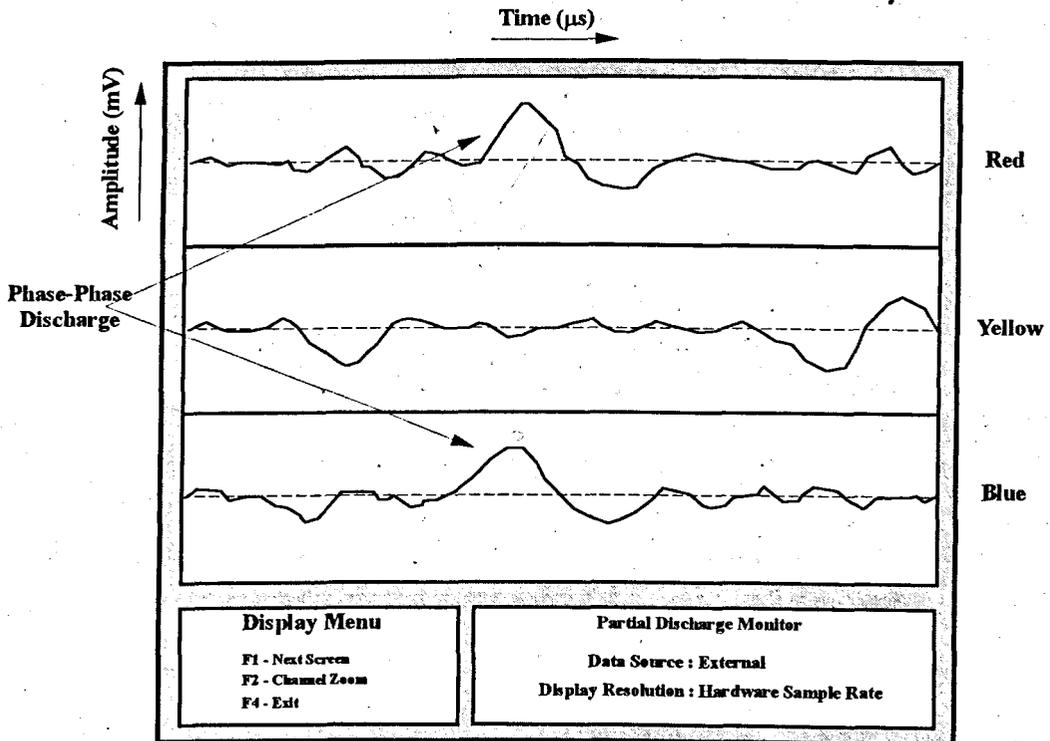


Figure 5.2 Three-Phase Data Display

At any time during the process of scanning the data, the user can opt to zoom in on a specific phase, Figure 5.3, receiving enhanced display resolution and the further option of a tracking cursor.

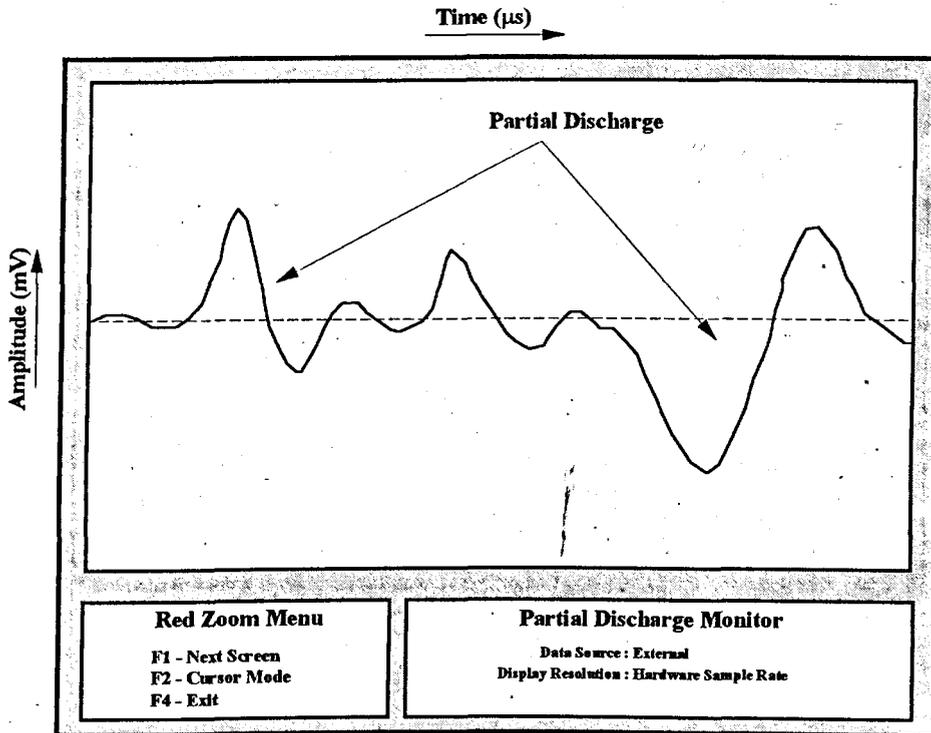


Figure 5.3 Single Phase Zoom

Upon zooming in on a specific phase, the operator can optionally invoke a tracking cursor, Figure 5.4. This allows a movable cursor to track across the waveform on the screen, giving amplitude and timing information about each data sample in turn. The additional information given to the operator would allow him to quantify relative discharge amplitudes and get an idea of where they occur within the complete data record.

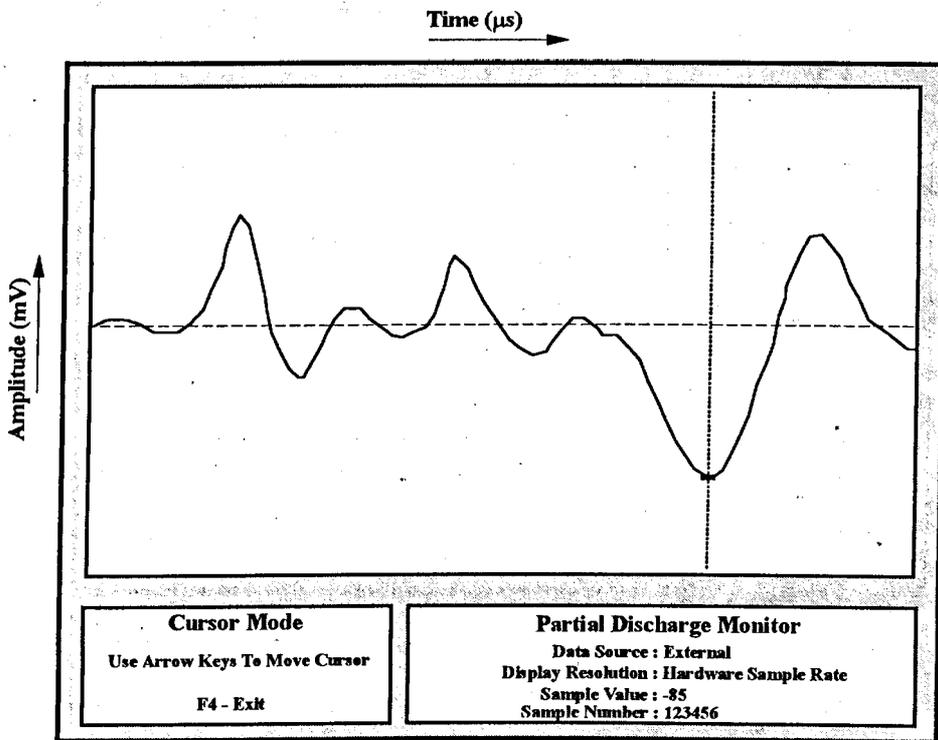


Figure 5.4 Tracking Cursor Display

5.2.2 Three-Phase Pulse Analysis Functions

Pulse Analysis, within the menu hierarchy, incorporates the three pulse analysis algorithms explained earlier: Pulse/Height Analysis, Phase/Height Analysis and Phase/Count Analysis. When applied to three phase data, these routines simultaneously display the corresponding analysis for all three phases, to allow operator comparison between phases. Figure 5.5 shows an example of the sort of results an operator might obtain using the Pulse/Height Analysis option.

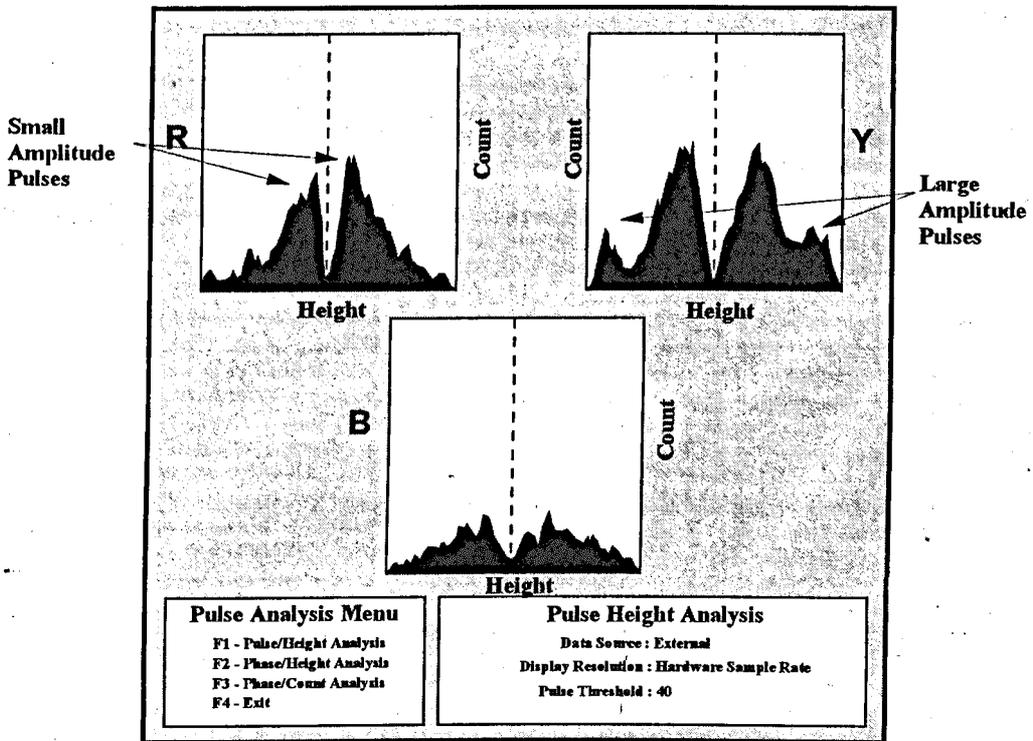


Figure 5.5 Pulse/Height Analysis

Here, individual partial discharge pulses are first identified, their height measured and a count kept of the number of pulses that occur at that height. Note that the amplitudes of the discharge pulses can be positive or negative and this is reflected in the resulting graphs. From the diagram it can be seen that the Yellow phase has more large amplitude discharges, whilst the Red phase has a greater number of smaller discharges. This additional comparative information will undoubtedly assist an operator in his eventual diagnosis of machine insulation health.

In Phase/Height Analysis, each phase sample record is broken down into 180 sections, each section representing 2° phase of the reference 50Hz supply. Each section is then scanned, in turn, and the largest positive and negative discharge pulses identified. At the same time, the total number of discharges is also established; this is used in Phase/Count Analysis.

Using the Phase/Height Analysis option, Figure 5.6, a skilled operator could identify regions in the plot where common discharges are occurring. These could be indicative of phase-phase partial discharges, or areas where the discharge activity is swamped by high-frequency noise, emanating from equipment such as thyristor converters. (The subject of noise rejection is a complex one and will be covered in a later section.)

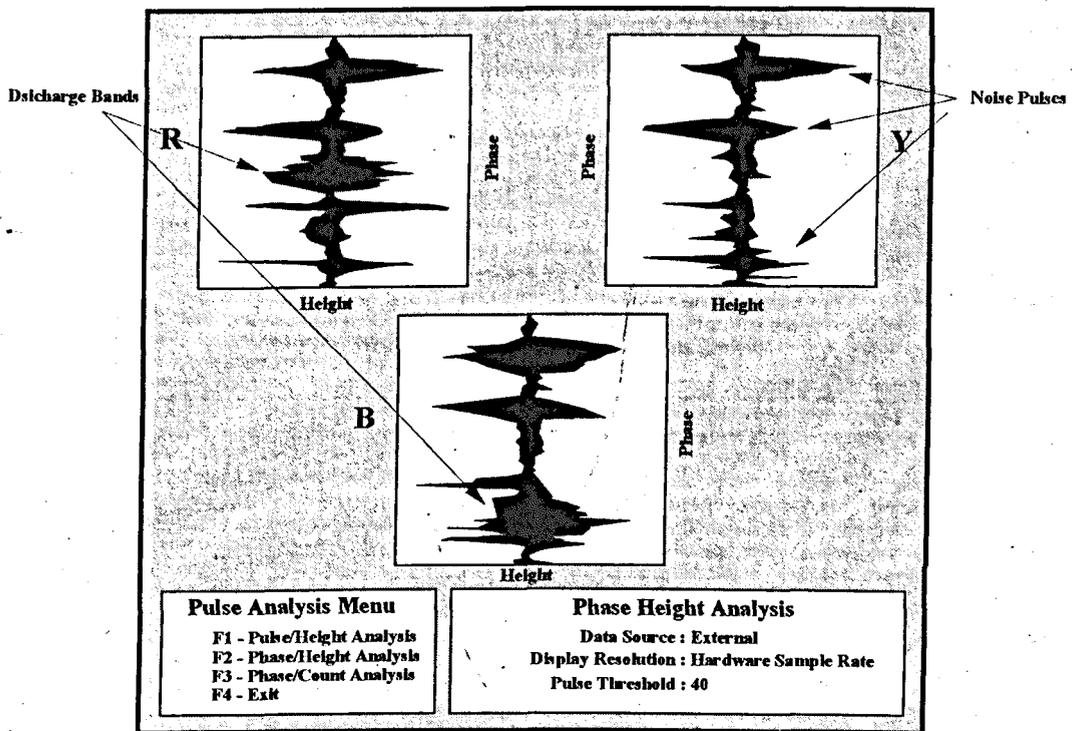


Figure 5.6 Phase/Height Analysis

The Phase/Count Analysis option, Figure 5.7, allows the operator to view the distribution of pulse repetition rates throughout the 50Hz supply period. This option, in itself, is not a particularly useful one, but does give an additional perspective and may detect some anomalous discharge behaviour that the other two analyses did not highlight.

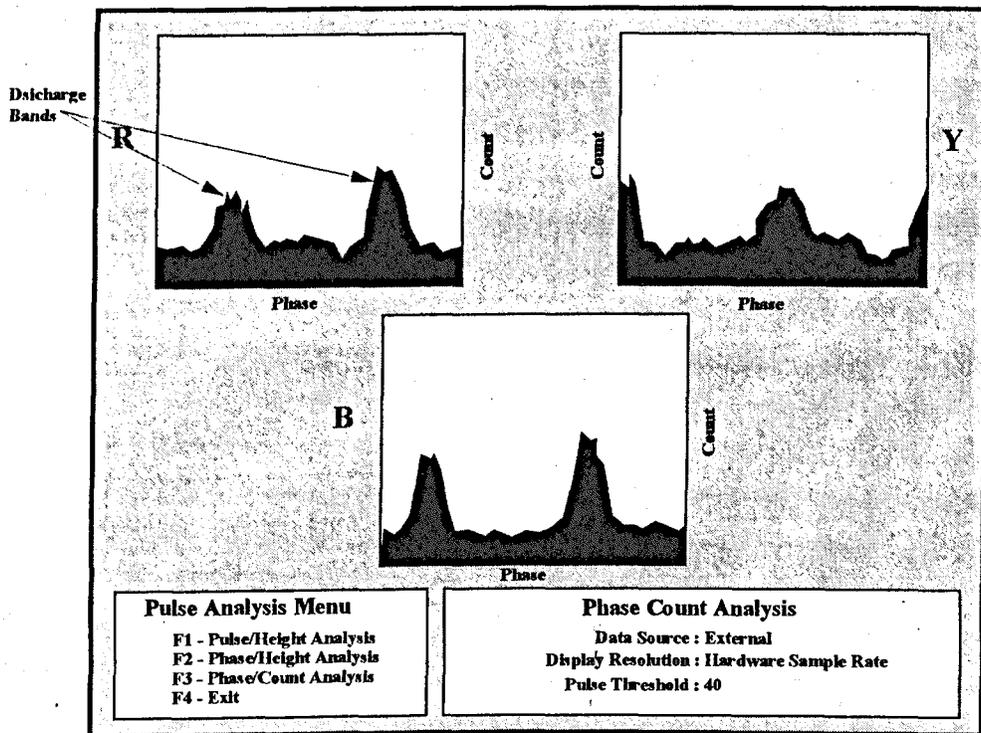


Figure 5.7 Phase/Count Analysis

5.3 Noise Rejection

Any practical monitoring system has to overcome the problem of superimposed noise generated by converter equipment. The switching on and off of thyristors used in such equipment cause high frequency energy pulses to be transmitted onto the mains supply. These noise pulses are periodic in nature and can number up to 24 in a single mains cycle, depending upon the type of converter equipment under use. Thyristor switching noise is of little significance to other devices connected to the mains, but can have a profound effect upon partial discharge monitoring equipment.

High frequency noise, such as that routinely generated by thyristor converters, has a similar frequency content to partial discharges and therefore cannot be removed by filtering. This noise, if processed, would undoubtedly affect the outcome of the analysis routines as, on the whole, thyristor noise has a greater energy content than partial discharges. This in turn could lead to erroneous conclusions about the insulation health of a machine under test.

Considering the periodic nature of this interference, Figure 5.8, several elimination possibilities exist: the simplest of these is subtraction of successive 50Hz cycles, whereby the periodic interference would be eliminated with the random partial discharges remaining unaffected. However as the partial discharges are not totally random in nature, but usually occur in the same general area of successive cycles, a subtraction algorithm would also alter the discharge patterns. A far more suitable scheme is to first locate the positions of the interference within the supply reference cycle and then "window" out the portions that are affected. In most cases this will have minimal effects on the discharges because the noise will have totally swamped any discharge activity in that section. This statement only holds true if the noise pulses do not coincide with the discharge bands. If this happens to be the case, there are no alternatives but to re-acquire a new set of data samples until the discharge bands are not concurrent with any noise pulses. (Converter pulse patterns can move within the mains cycle over time.)

In practice, a number of data sets would be acquired in order to minimise the possibility of missing discharges that occur less frequently or those that have been swamped by the imposed noise.

Locating the noise pulses in the cycle is by no means an easy matter as the most relevant signal processing algorithms, namely autocorrelation and cross correlation, are unsuitable due to the nature of the signals under scrutiny.

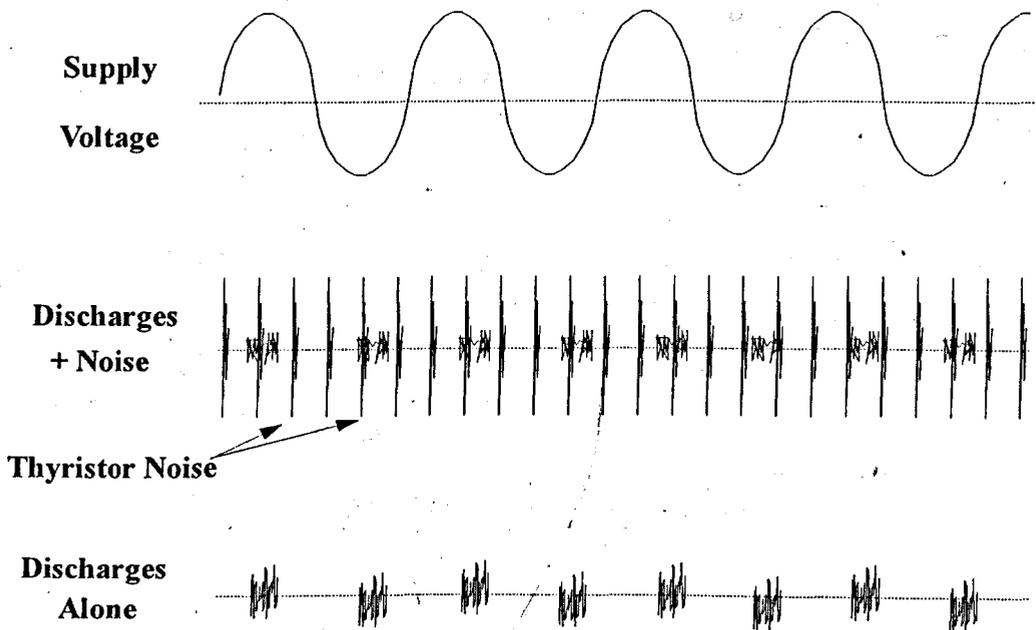


Figure 5.8 Thyristor Switching Noise

5.3.1 Autocorrelation

The autocorrelation function (ACF) of a signal waveform is an average measure of its time-domain properties, which makes it especially relevant for random signals. Formally the autocorrelation function is defined as

$$r_{xx}(\tau) = \lim_{T_0 \rightarrow \infty} \frac{1}{T_0} \int_{-\frac{T_0}{2}}^{\frac{T_0}{2}} f(t) \cdot f(t + \tau) \cdot dt$$

It is equal to the average product of the signal $f(t)$ and a time-shifted version of itself, and is a function of the imposed time-shift, τ . The above expression applies to a continuous signal of infinite duration. If it is used for a signal of finite length, such as a phase data record, the average product over a very long time interval T would tend to zero at all values of τ . It is necessary therefore to use a modified version of the autocorrelation function, generally called the 'finite ACF', which is defined as

$$r'_{xx}(\tau) = \int_{-\infty}^{\infty} f(t) \cdot f(t + \tau) \cdot dt$$

Here, as the signal is a sampled data one, the product of the signal and its shifted version only has non-zero values when the shift is equal to a multiple of the sampling interval T , therefore the ACF is defined as:

$$r_{xx}(k) = \lim_{N \rightarrow \infty} \frac{1}{(2N+1)} \sum_{m=-N}^N x_m \cdot x_{m+k}$$

Where x_m and x_{m+k} represent two sample values separated by kT seconds, and the summation has $(2N+1)$ terms with the integer parameter m taken between $-N$ and $+N$.

The ACF of a sinusoid, Figure 5.9, demonstrates the inherent problem in using this technique to identify the positions of periodic thyristor noise.

From the diagram it can be seen that maximum correlation's occur periodically: i.e. when peaks of the delayed signal occur simultaneously with a positive or negative peak in the original signal. This means that it is very useful for highlighting periodic signals immersed in random noise which, at first glance, is exactly what is required to determine the positions of the periodic thyristor noise pulses, Figure 5.8. However, bearing in mind the definition of the ACF, if an estimate is made of the result of an ACF on the sampled signal it is found that the ACF reaches a maximum value at a time shift of 0 (the signal exactly correlates with itself) and again at shifts equal to

multiples of the noise period, Figure 5.10. Therefore, essentially all the autocorrelation function discloses is the period of the thyristor pulses, not their positions within the reference supply cycle.

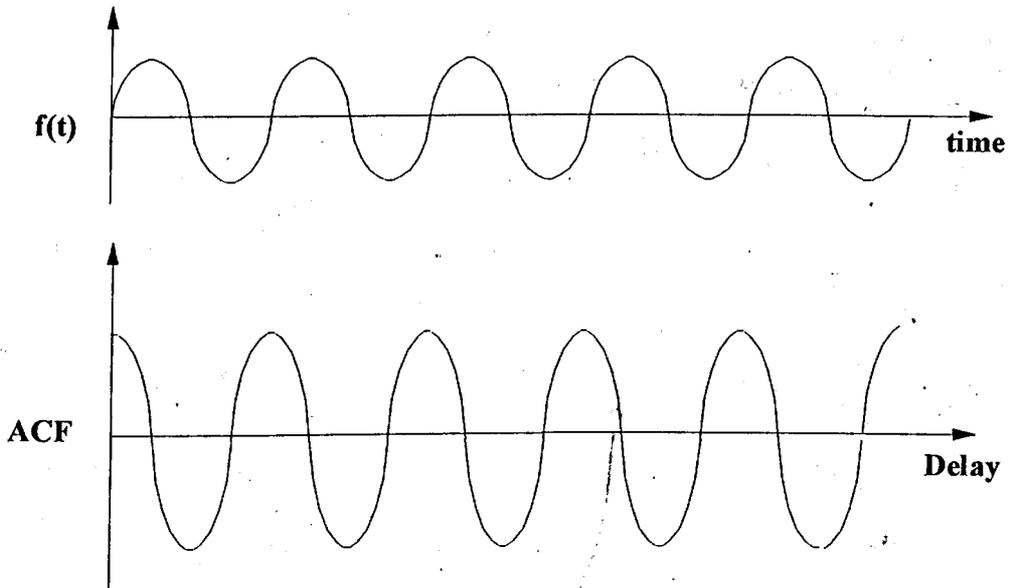


Figure 5.9 Autocorrelation of a Sinusoid

The larger peaks in the autocorrelation function in Figure 5.10 correspond to delays when the discharge pulses as well as the thyristor pulses correlate. i.e. the delay is equal to a multiple of the partial discharge bands, which are a function of the supply voltage.

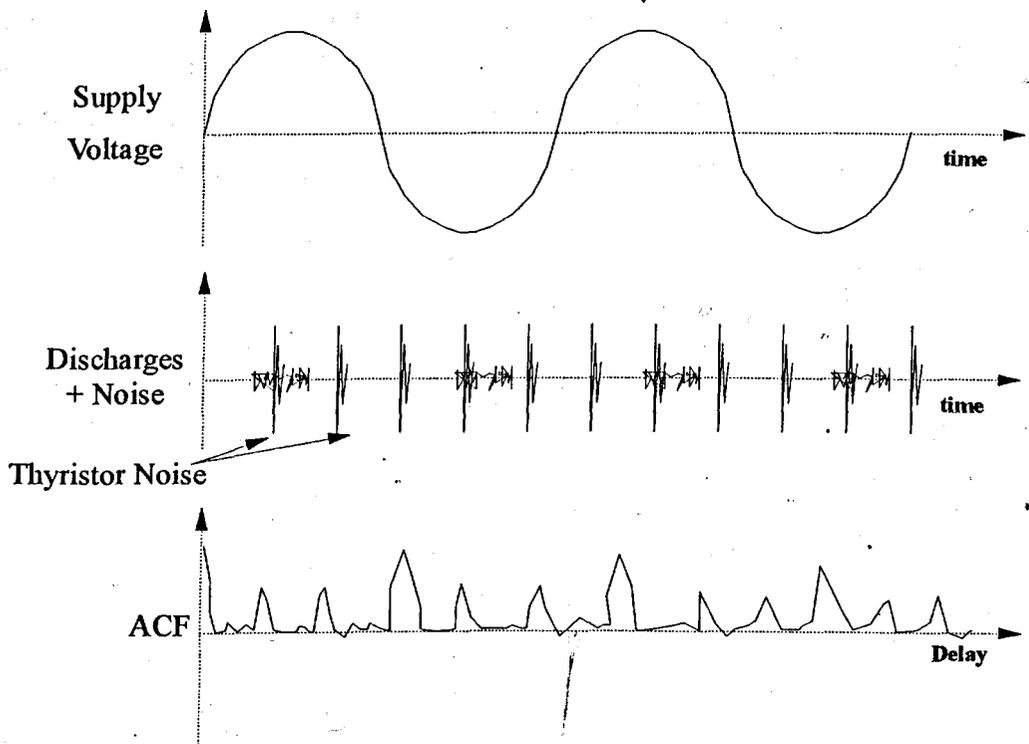


Figure 5.10 Autocorrelation of Discharges + Noise

5.3.2 Cross-Correlation

The cross-correlation function (CCF) describes the similarity between two signals in the same way as the ACF does for one. The CCF may be defined as:

$$r_{xy}(\tau) = \lim_{T_0 \rightarrow \infty} \frac{1}{T_0} \int_{\frac{T_0}{2}}^{\frac{T_0}{2} + \tau} f(t) \cdot f_2(t + \tau) \cdot dt$$

The corresponding definition for a sampled system is:

$$r_{xy}(k) = \lim_{N \rightarrow \infty} \frac{1}{(2N+1)} \sum_{m=-N}^N x_m \cdot y_{m+k}$$

In practice this means that the CCF is the average product of one signal and a time-shifted version of the second, Figure 5.11. Here, a random signal is cross-correlated with a time shifted version of itself. The resultant CCF plot gives a maximum when the delay of the CCF exactly equals the original time shift t_1 .

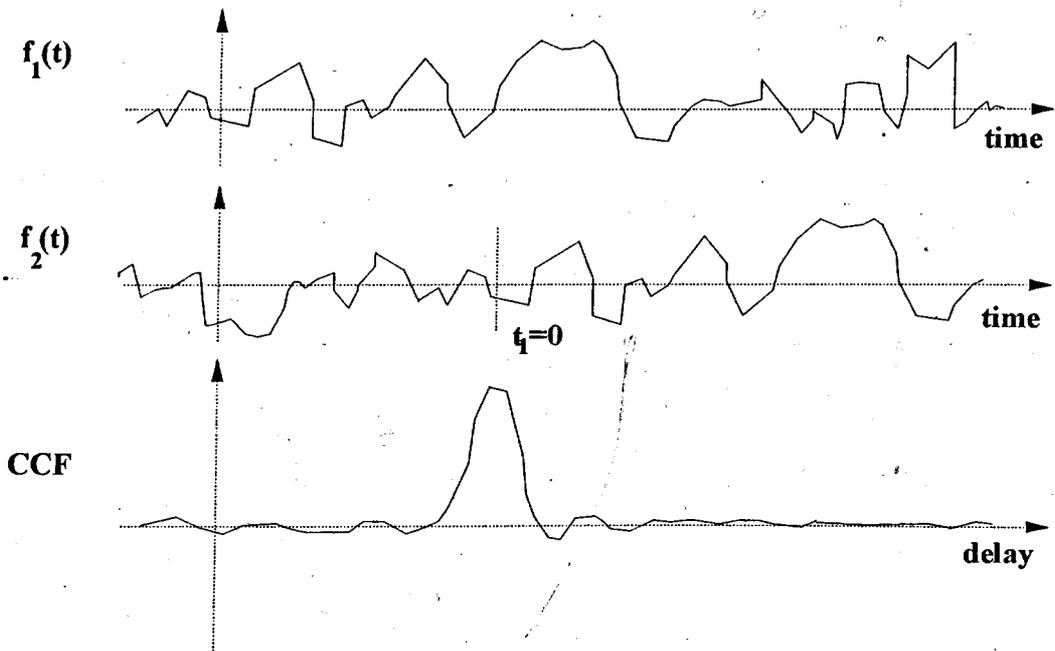


Figure 5.11 Cross Correlation of a Random Signal

In order to use this function to determine the positions of the thyristor pulses in the single phase sample record a signal comprising of a single thyristor pulse must be available to cross-correlate with the sample record, Figure 5.12. This solution is not a very practical one as there is no easy way to acquire, or simulate, a single thyristor noise pulse.

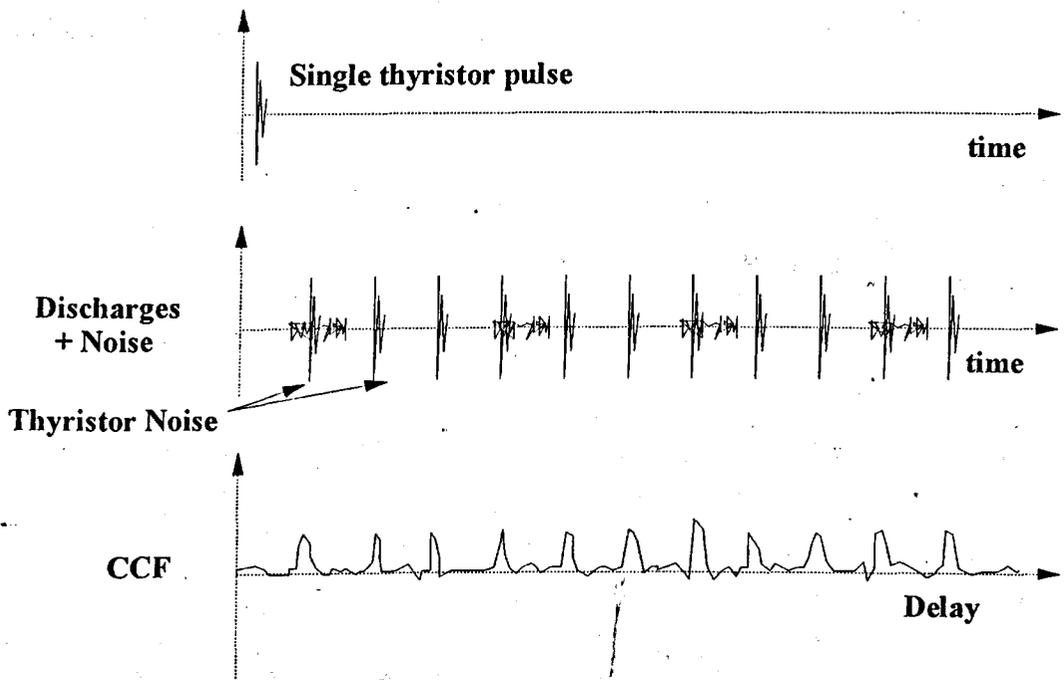


Figure 5.12 CCF of Single Noise Pulse and Sample Record

5.3.3 Three Phase Statistical Scheme

Thyristor interference is present on all three phases of the machine simultaneously: i.e. The noise pulses coincide almost exactly when digitised, Figure 5.13. This, coupled with the noise periodicity, means that, by using the three phase data, the positions of each noise pulse could be determined and all the data at that point ignored, leaving the remaining partial discharge information free to be processed normally.

The first problem in such a scheme is to somehow enhance the noise pulses to make them easier to identify. Partial discharges very rarely occur on three phases simultaneously, so if all three phases are added together any common information, such as the noise pulses, will be enhanced whilst everything else will tend to remain

constant. Having enhanced the noise pulses, a detection algorithm had to be developed that could discriminate between noise pulses and partial discharges.

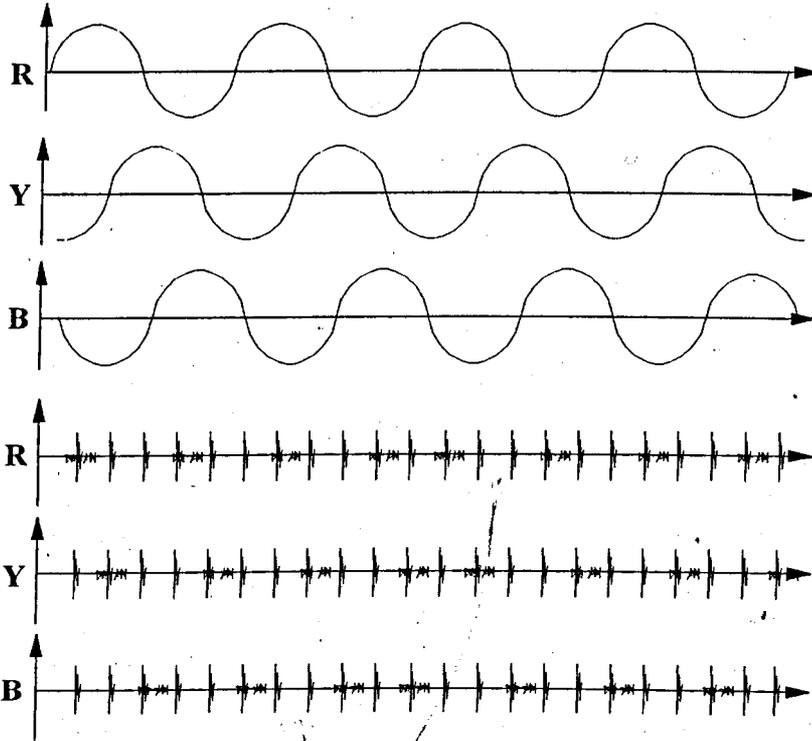


Figure 5.13 Three Phase Interference

Using rudimentary statistical analysis, once all three phases had been summed to produce a single common record, the mean level of the data samples is computed using:

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$$

where x_i is a sample of the summed record.

All of the pulses in the summed record that exceed the mean level by an operator adjustable factor are marked. Next, the distance between each marked pulse is

measured, simply by comparing sample numbers, and the mean distance calculated using:

$$\bar{d} = \frac{1}{k} \sum_{j=1}^k d_j$$

Where d_j is the distance between two identified pulses out of a total of k pulses.

The standard deviation of the distances is found using:

$$\sigma = \sqrt{\frac{\sum_{j=1}^k (d_j - \bar{d})^2}{(k-1)}}$$

The standard deviation is found in order to help discriminate between valid noise pulses and large-amplitude discharges, which otherwise could be identified as potential noise pulses.

The noise pulses are periodic in nature, so, starting from the largest detected pulse in the summed data record (it is assumed to be a noise pulse), the data is scanned using the mean and standard deviation of distance to pick out all other noise pulses. These marked pulses are then displayed on the screen for operator approval, Figure 5.14, before being removed from each data phase.

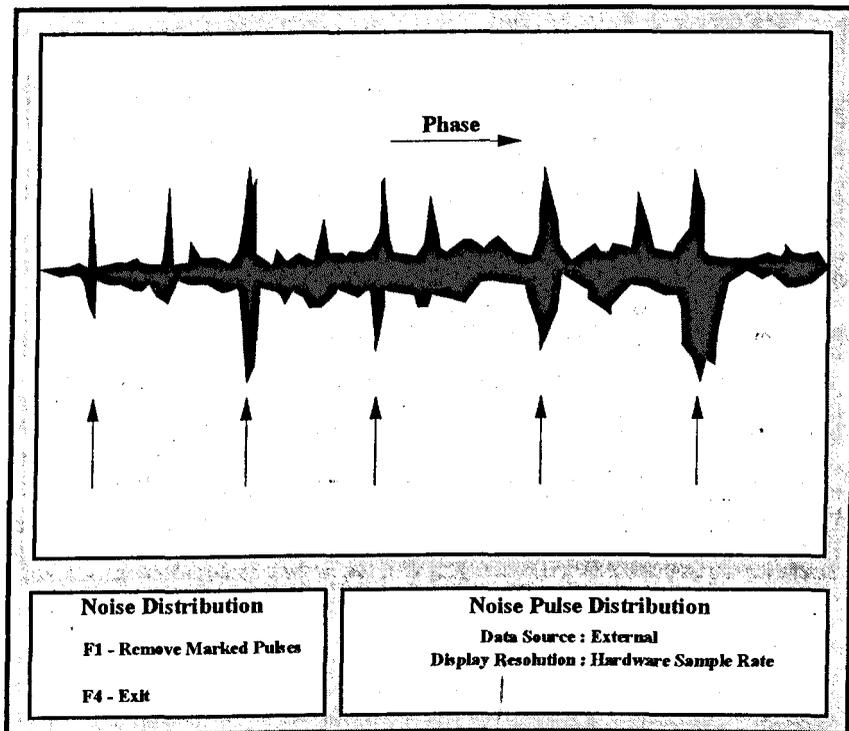


Figure 5.14 Noise Pulse Distribution

The removal process involves determining an optimum window size for each noise pulse that will remove it in its entirety without losing too much partial discharge information. This is computed along the same lines as the noise detection algorithm, whereby a threshold above the mean level of the summed signal is set. Each noise pulse is then scanned to find the optimum window size to completely remove it. The sample addresses that these windows correspond to are then used to remove the noise pulses from each phase record in turn (by setting the sample value at each address to 0).

5.4 Application of Pulse and Noise Analysis

So far, the discussion has been centred around the design and development of the relevant hardware and software routines necessary to implement a rudimentary partial discharge monitoring system. Having integrated these into a single package, the system could be applied to real partial discharge data.

A lot of partial discharge data, from many different machines, was available as analogue records on magnetic tape. Some of these were used to test the algorithms developed and to demonstrate the practical application of such a monitoring system. The data record used in the subsequent pages was obtained from site measurement on one particular 11KV machine operating on an offshore platform in the North Sea. This record was chosen because it demonstrates the common problem of thyristor converter interference.

The following diagrams are a direct representation of the computer screen during processing: a routine was developed to allow the image shown on the computer screen to be dumped as a bit-map to a printer. Unfortunately, no colour information can be reproduced so the full effect of the information on display is lost.

Figure 5.15 shows the three-phase pulse height analysis for the machine under test, prior to the thyristor interference removal. The absence of colour information necessitates some explanation about the plot: the top-left plot represents the red phase, the lower plot is the yellow and the top-right plot is the blue phase pulse/height analysis.

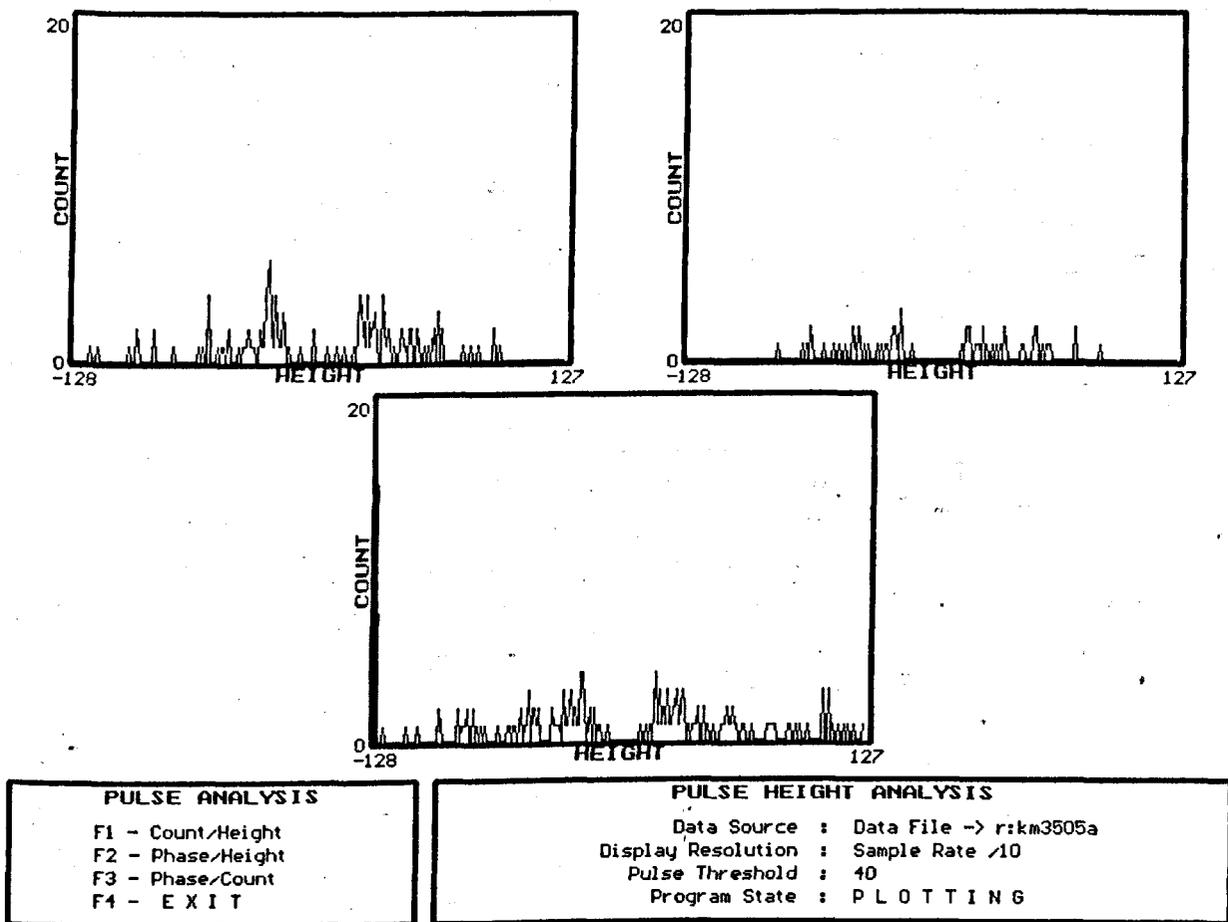


Figure 5.15 Pulse/Height Analysis

From the diagram it is not readily apparent that there is any interference on the partial discharge data. The plots show a relatively small amount of discharge activity, but with some higher discharges occurring: readings at the extremes of each box indicated pulses of large amplitude.

If the Phase/Height Analysis is performed, Figure 5.16, a different picture emerges.

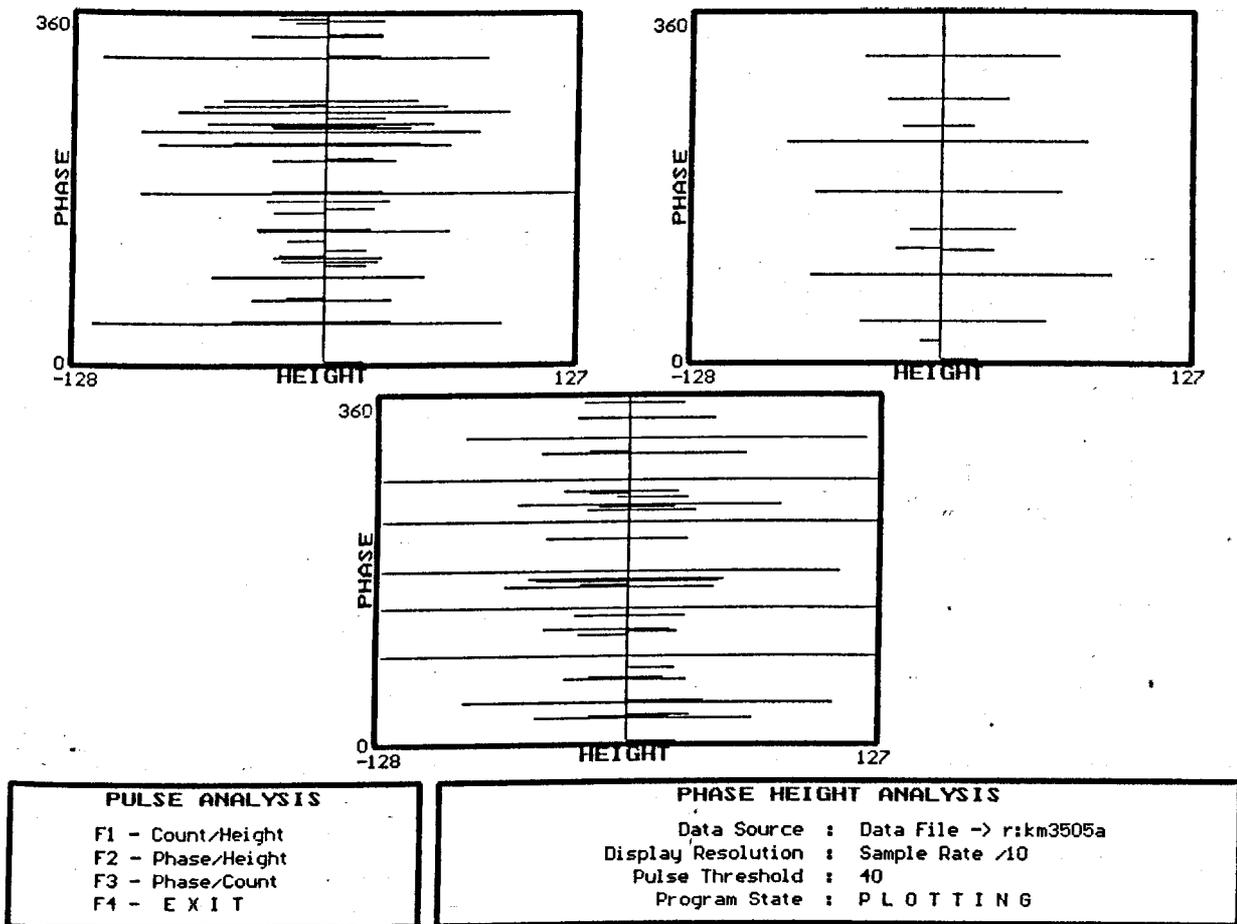


Figure 5.16 Phase/Height Analysis

This plot shows a completely different story to that of the pulse/height analysis. Here, each phase sample record is split into 2° phase windows and the maximum positive and negative pulses found and plotted within each window.

From the diagram, the periodic, large-amplitude interference pulses can be seen. These interference pulses, if not recognised and removed, would undoubtedly lead to an over-estimation of the partial discharge problems within the machine.

If the statistical noise removal scheme is now applied, Figure 5.17 results.

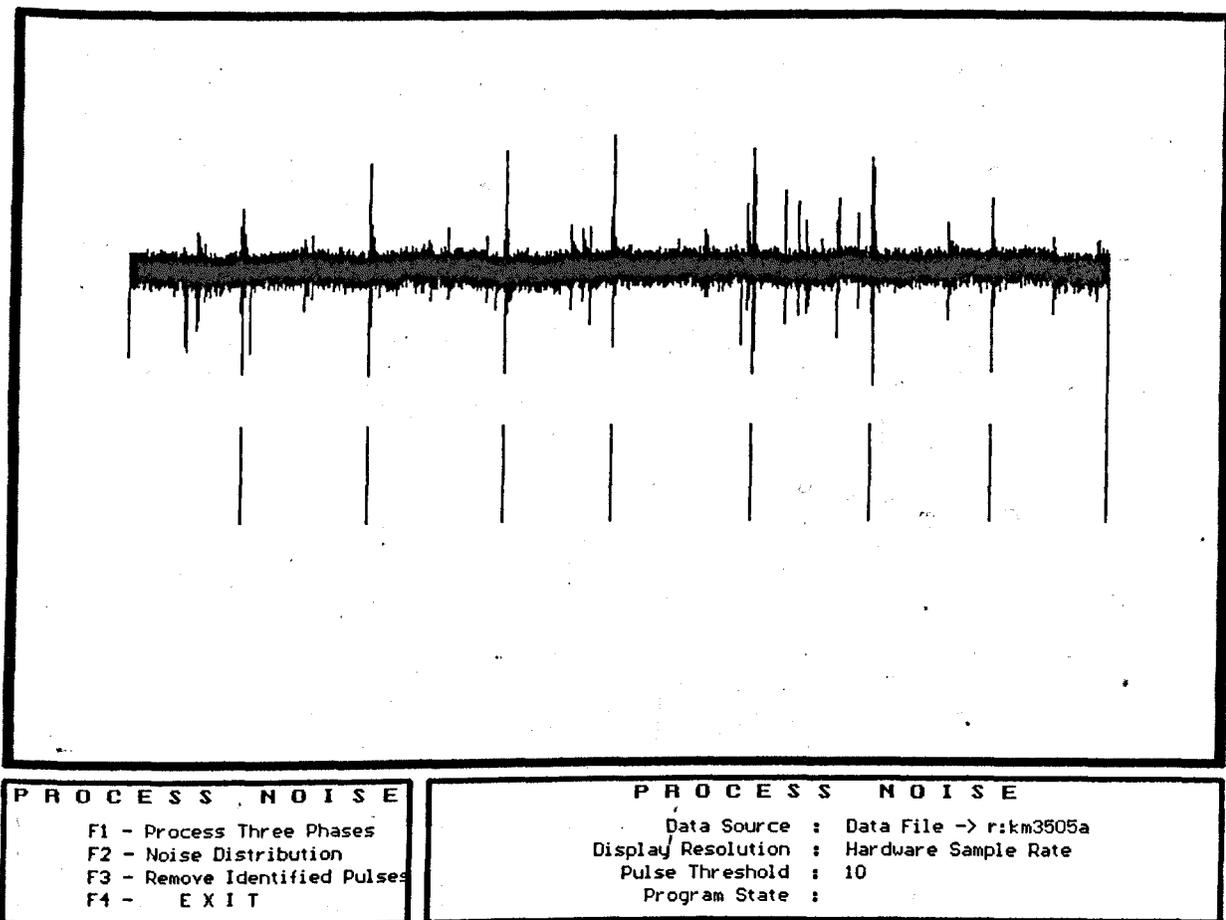


Figure 5.17 Noise Pulse Distribution

The diagram shows the summed record of the three phase data, with the prospective noise pulses highlighted by the series of lower lines. The prospective pulses are highlighted to allow the operator to view the pulses which will be removed so that he can override the decision if an error occurs.

In this case however, the highlighted pulses are indeed interference pulses and need to be removed. Once the highlighted pulses are accepted as noise pulses, the program calculates the necessary window size for each pulse and creates a copy of each phase data record, minus the interference pulses.

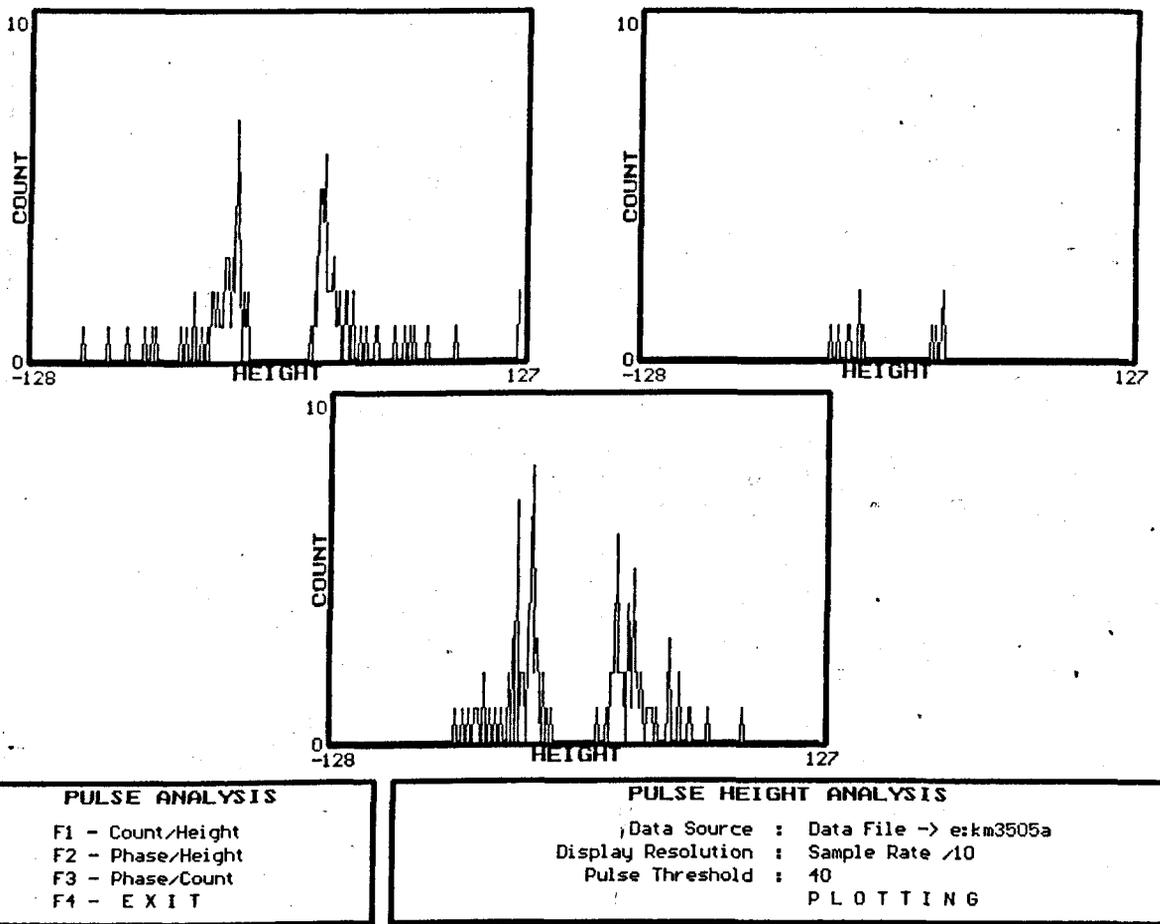


Figure 5.18 Pulse/Height Analysis Without Noise

Figure 5.18 depicts the pulse/height analysis of the three phase data record after the identified noise pulses have been removed. If this plot is compared with the previous pulse/height analysis plot, Figure 5.15, it is evident that a lot of the larger amplitude pulses have been removed, giving a much better view of the true level of discharge activity within the machine. (Note that the plot heights are auto-scaling making the comparisons difficult) The differences are particularly noticeable on the blue phase (top right) as it is the "quietest". The differences become even more obvious if a new phase/height analysis plot is taken.

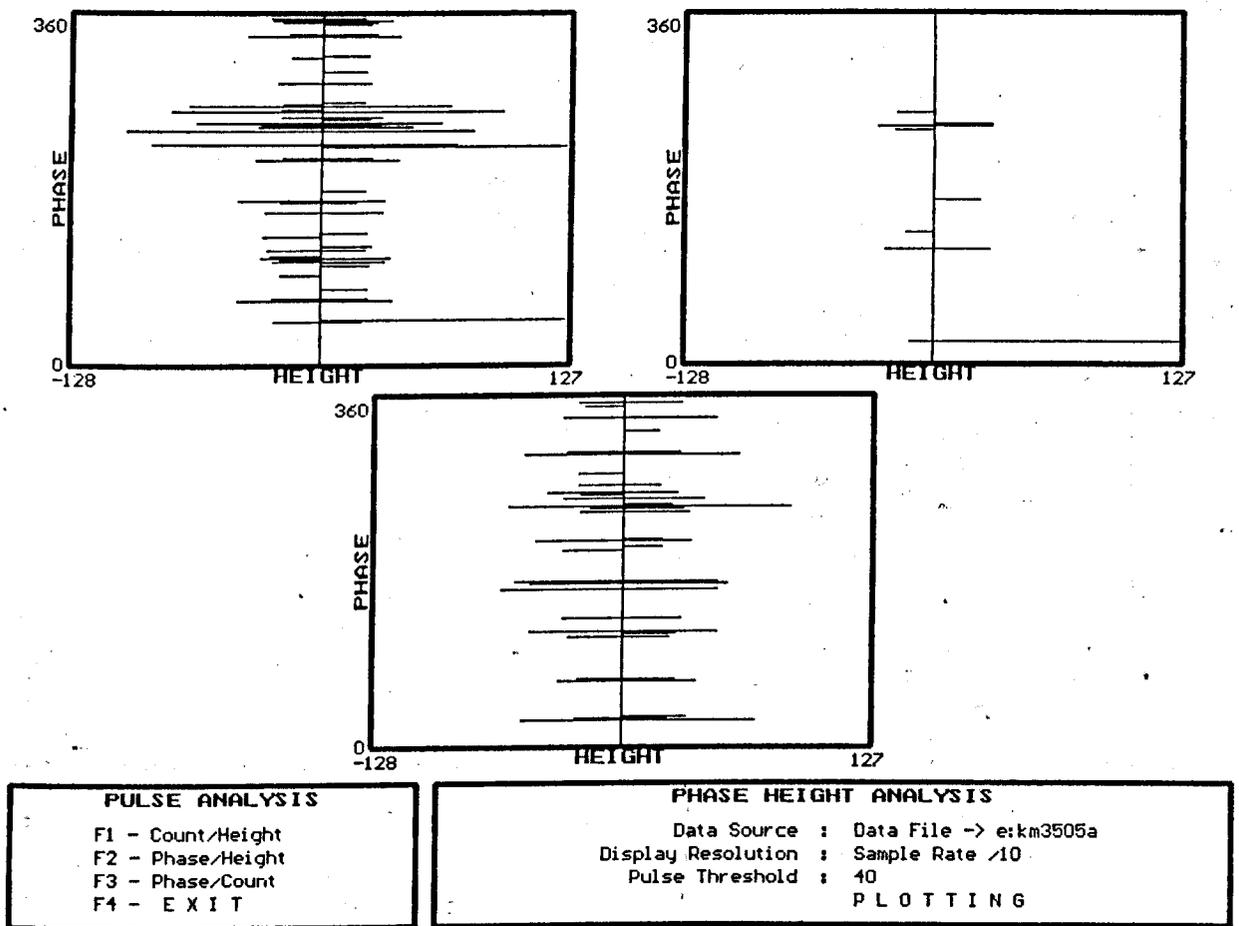
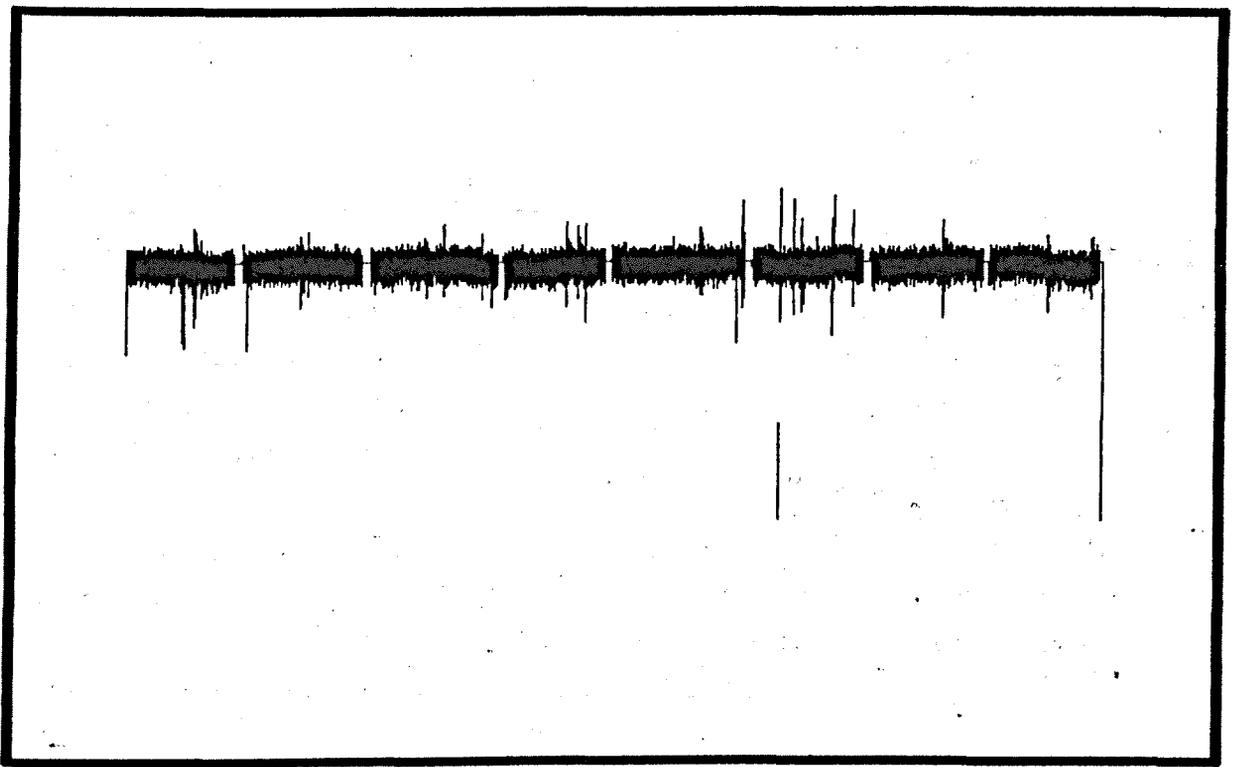


Figure 5.19 Phase/Height Analysis Without Noise

Figure 5.19 shows the phase/height analysis of the three-phase data once the thyristor interference pulses have been removed. If this plot is compared with that of Figure 5.16 it becomes apparent that the larger pulses that swamped the previous plot have been removed.

If the noise removal algorithm is invoked again, Figure 5.20 and compared with the previous result, Figure 5.17, the gaps where the interference pulses have been removed are visible.



P R O C E S S N O I S E
F1 - Process Three Phases
F2 - Noise Distribution
F3 - Remove Identified Pulses
F4 - E X I T

P R O C E S S N O I S E
Data Source : Data File -> e:km3505a
Display Resolution : Sample Rate /10
Pulse Threshold : 40
Program State :

Figure 5.20 Noise Pulse Distribution Showing Gaps

The above diagram shows the cumulative effect of windowing out the previously identified interference pulses. It is only included to graphically demonstrate the result of the process.

5.5 Frequency Analysis

The work, so far, has concentrated solely on time domain analyses of partial discharge activity. This work produced a number of analysis options, such as Pulse Height Analysis, where partial discharges are counted and categorised according to their amplitudes, with the resulting graph indicating the relative energies and repetition rates of the potentially damaging discharges.

The next logical step in the development process is to investigate frequency analysis of partial discharges, to see if further useful diagnostic information can be gained from them; this involves finding a method that would uniquely characterise the frequency content of individual partial discharge pulses.

Partial discharge pulses, as they propagate through a stator winding, are modified by the response of the winding. The response of a stator winding to such a high frequency disturbance is very complex, but essentially it behaves like a transmission line with frequency-dependant attenuation [11]. The frequency-dependant attenuation characteristic will obviously vary from machine to machine, but the fact still remains that the further a partial discharge pulse has to travel through a winding, the more its shape will be altered by the frequency-dependant attenuation characteristic of that winding. If a method is found which can accurately decompose a detected partial discharge pulse into its frequency content, then it may be possible to estimate where the inception of the discharge occurred within the winding.

Obviously such a facility would prove very useful as it could identify areas where localised damage is occurring, as well as saving time and effort during maintenance work by guiding the engineers to the areas of the winding directly affected by discharge activity.

However, before such schemes can be investigated, a method has to be found to characterise the frequency content of individual discharges. Difficulty arises due to the fact that each discharge only lasts for a very short time (typically $<10\mu\text{s}$), thus obtaining results with a good frequency resolution is difficult when using conventional frequency analysis algorithms; this is because the frequency resolution is indelibly linked to the length of sample being analysed. However, good frequency

resolution may not be necessary for most purposes, therefore it was decided to initially investigate Fast Fourier Transforms (FFT) applied to individual partial discharges.

5.5.1 Fast Fourier Transform

A physical process can be described either in the time domain, given some quantity as a function of time, e.g. $f(t)$, or else in the frequency domain, where the process is described by giving its amplitude as a function of frequency, e.g. $G(\omega)$. One can traverse between these two representations by means of the following equations:

$$G(\omega) = \int_{-\infty}^{\infty} h(t)e^{-j\omega t} dt$$

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} G(\omega)e^{j\omega t} d\omega$$

In most common situations, $f(t)$ is sampled at evenly spaced intervals in time and is of finite duration. Therefore the expression for the Fourier transform becomes,

$$X(k) \equiv X\left(\frac{2\pi k}{N}\right) = \sum_{n=0}^{N-1} x_n e^{-j\frac{2\pi kn}{N}} \quad k = 0, 1, 2, \dots, N-1$$

for a finite-duration sequence $x(n)$ of length N . This produces frequency samples $X(\omega)$ at equally spaced frequencies $\omega_k = 2\pi k/N$. This relation is called the Discrete Fourier Transform (DFT).

There is a very simple relationship between the resulting frequency resolution of the transform, the number of samples N and the sample rate Δt ,

$$\Delta f = \frac{1}{N\Delta t} \quad \text{or} \quad \frac{f_s}{N}$$

where Δf is the frequency resolution: i.e. the line spacing of the resulting spectrum, and f_s is the sampling frequency.

The main drawback of the DFT algorithm is that it takes N^2 operations to compute. The discrete Fourier transform can, in fact, be computed in $N \log_2 N$ operations with an algorithm called the Fast Fourier Transform, or FFT. Such is the efficiency of the FFT over the DFT that the former is invariably preferred to the latter when it comes to implementing a frequency analysis algorithm. The only drawback with the FFT is that the number of samples needs to be a power of 2; this restriction does not apply to the DFT.

Fast Fourier Transform algorithms are now fairly common-place in signal processing systems, so instead of developing one from first principles it was decided to adapt an existing algorithm, the capabilities of which had been proven. The algorithm chosen was one listed in the book "Numerical Recipes in C" [20].

The main problem with spectral analysis of partial discharges using a conventional FFT is the lack of frequency resolution. If we examine the relationship that defines the spectral line resolution,

$$\Delta f = \frac{1}{N\Delta t} \quad \text{or} \quad \frac{f_s}{N}$$

it is apparent that the resolution of the resulting spectrum depends on the number of samples applied to the FFT. Now, partial discharges have a very short duration, therefore they will have a correspondingly short sample length, between 100 and 200 samples when digitised at 10MHz. This means that short data records have to be applied to the FFT algorithm, resulting in poor spectral resolution.

Applying some numbers to the equation:

$$\text{let } N = 512$$

$$f_s = 10 \text{ MHz}$$

$$\Delta f = \frac{f_s}{N} = \frac{10 \times 10^6}{512} = 19.5 \text{ KHz}$$

Therefore, if a 512 sample Fast Fourier Transform is employed, the resulting spectrum would have a frequency resolution of approximately 20kHz. This means that the frequency spectrum of the analysed discharge is displayed as discrete lines 20kHz apart. If a 256 sample FFT were employed, then the resultant frequency spectrum would have discrete lines 40kHz apart, giving less information.

However, before any frequency analysis algorithm could be implemented some software routines had to be developed which would allow individual discharges to be selected for analysis. Figure 5.21 shows a screen representation of how the user selects the data for application to the FFT.

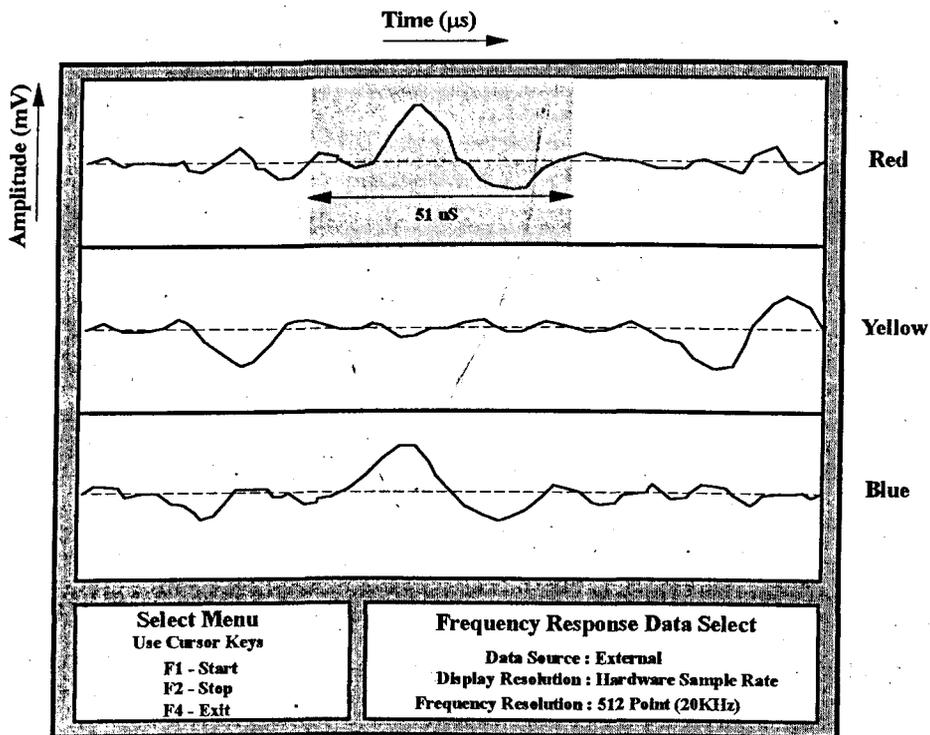


Figure 5.21 Data Selection

The user can scan through all of the captured partial discharge data and select any discharge from any phase for analysis. This data is then placed in an array, and made the required length by padding with zeros, before being passed to the FFT algorithm for analysis. An example of the resulting analysis can be seen in Figure 5.22.

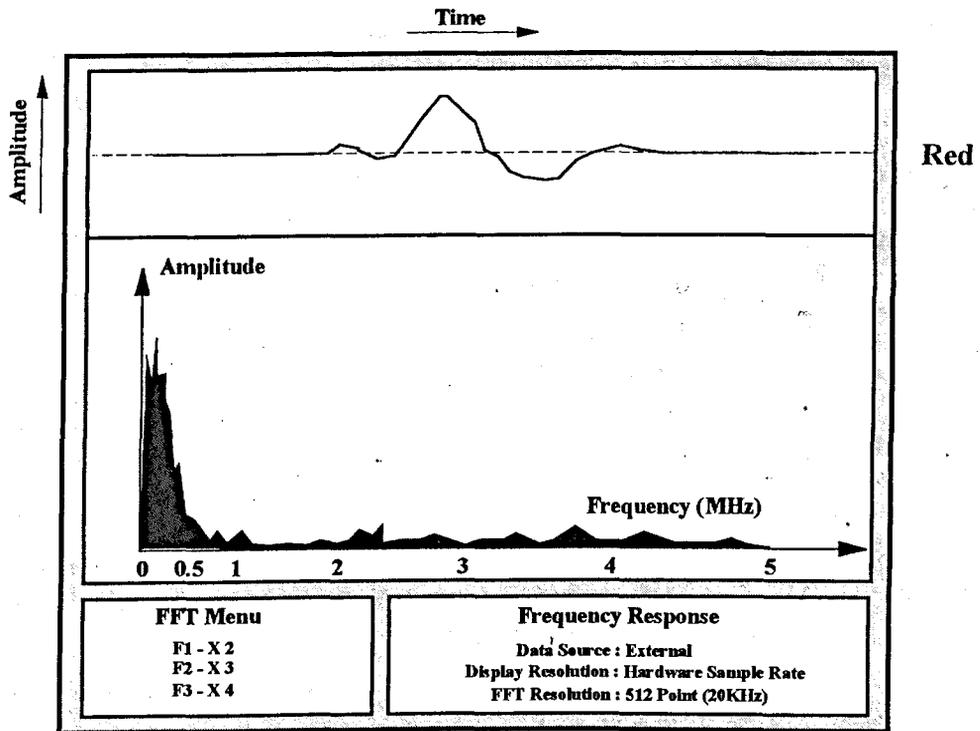


Figure 5.22 FFT Analysis

A potential solution to the frequency resolution problem would seem to come in the guise of zero-padding, whereby the length of the sample record passed to the FFT is increased by padding with zeros. Whilst this does indeed result in a spectrum with more frequency lines, it does not increase the resolution of the plot, Figure 5.23.

From Figure 5.23 it is apparent that zero-padding the sample record just causes the FFT algorithm to produce intermediate spectral lines that are interpolations between the original spectral lines, derived from the non-zero samples. This limitation is fundamental to the conventional FFT and means that other spectral analysis algorithms may be more suitable for analysis of partial discharges.

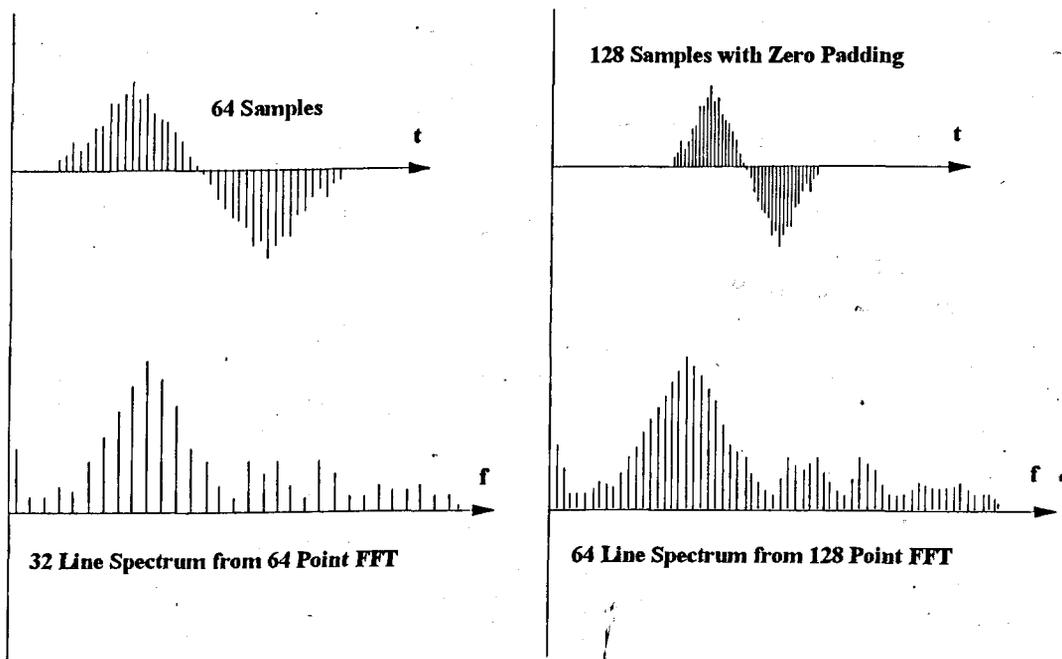


Figure 5.23 Effect of Zero-Padding the Data Record

Having perfected the frequency analysis and resulting display routines, experiments were carried out to see if the simple FFT algorithm could distinguish between discharges that originate from different locations within a stator winding. The experimental set-up is shown in Figure 5.24.

A large laboratory-based 6.6kV stator winding was used for this experiment, with the insulation stripped at regular intervals from the inter-coil connections so that the artificial discharge source, represented by the voltage generator E in series with the capacitor C, could be connected to the winding at alternative positions.

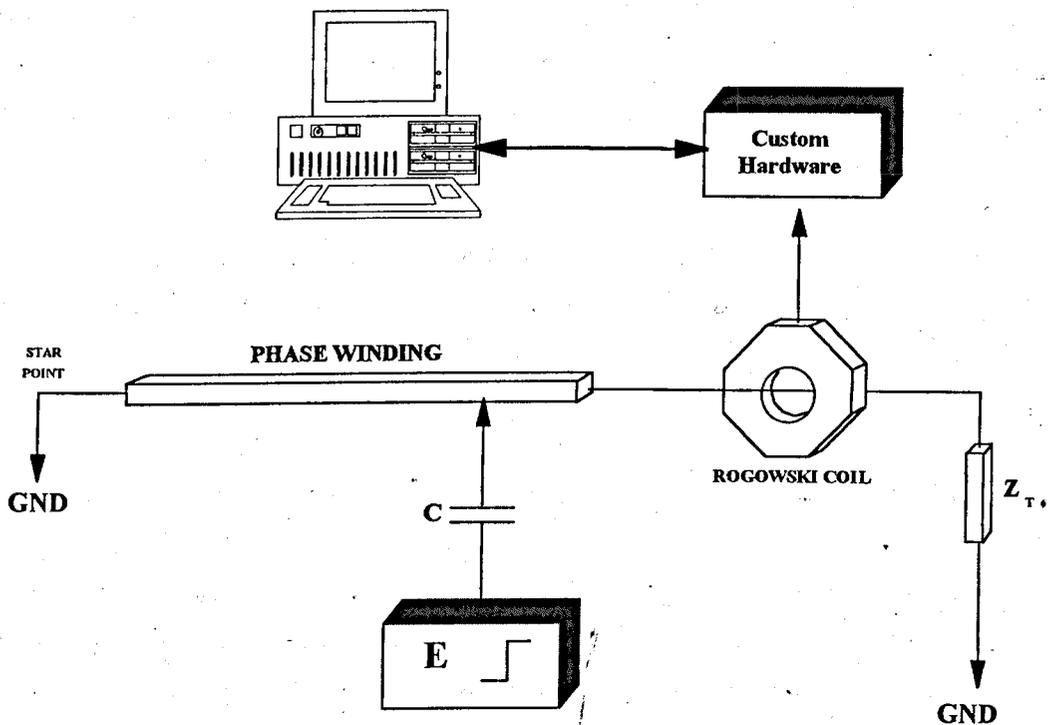


Figure 5.24 Experimental Set-up

Artificial discharges of known magnitude were injected at different locations in the winding, up to 10 coils from the line end, and the resulting signals saved on disk for later processing.

Once the data gathering had been completed, the resulting files were processed. The artificial discharge data was applied to a 256 point FFT algorithm, resulting in a spectral plot with discrete frequency lines approximately 40kHz apart. In order to get a comprehensive frequency profile of each data record, a number of pulses were analysed, the average profile found, as well as the standard deviation. These were then displayed graphically for comparison with the results from pulses injected at different locations.

Obviously, it would have been very time consuming, and prone to error, for the operator to manually select each pulse before calculating the average frequency profile. To overcome this, the software algorithms developed to detect pulses for the time analysis routines were copied and customised. These operated first by scanning

through the data store looking for a discharge peak, and once found, moving back a set number of samples in the store to before the inception curve of the pulse, then placing the next 220 data samples in the centre of a register that had been pre-loaded with zeros. This register was then applied to the FFT routine for analysis.

The number of data samples that could have been applied to the FFT was obviously 256, but the discharges were deemed to have ceased after approximately 200 samples. However, the full register width of 256 data samples was used for analysing pulses injected from deeper in the winding.

The following graphs are direct computer-screen reproductions and display the findings of the FFT analysis for artificial discharges injected at 1 coil intervals, up to 10 coils into the stator winding. In each case, the top graph shows one FFT analysis: the upper part displaying the pulse as a function of time. (i.e. the contents of the data register used in the analysis with 256 samples of data corresponding to a time record of 25.6 μ s.) The lower part displaying the corresponding spectral plot.

The bottom graph shows the mean frequency profile of 50 such pulses, along with the standard deviation for each averaged spectral line. The standard deviation plot indicates when the results of the averaging process become unreliable.

i.e. If the amplitude in the Mean plot is x and the standard deviation amplitude is y , then the actual value at that spectral point is $x \pm y$. Obviously, as y approaches x the true results of the mean become ambiguous.

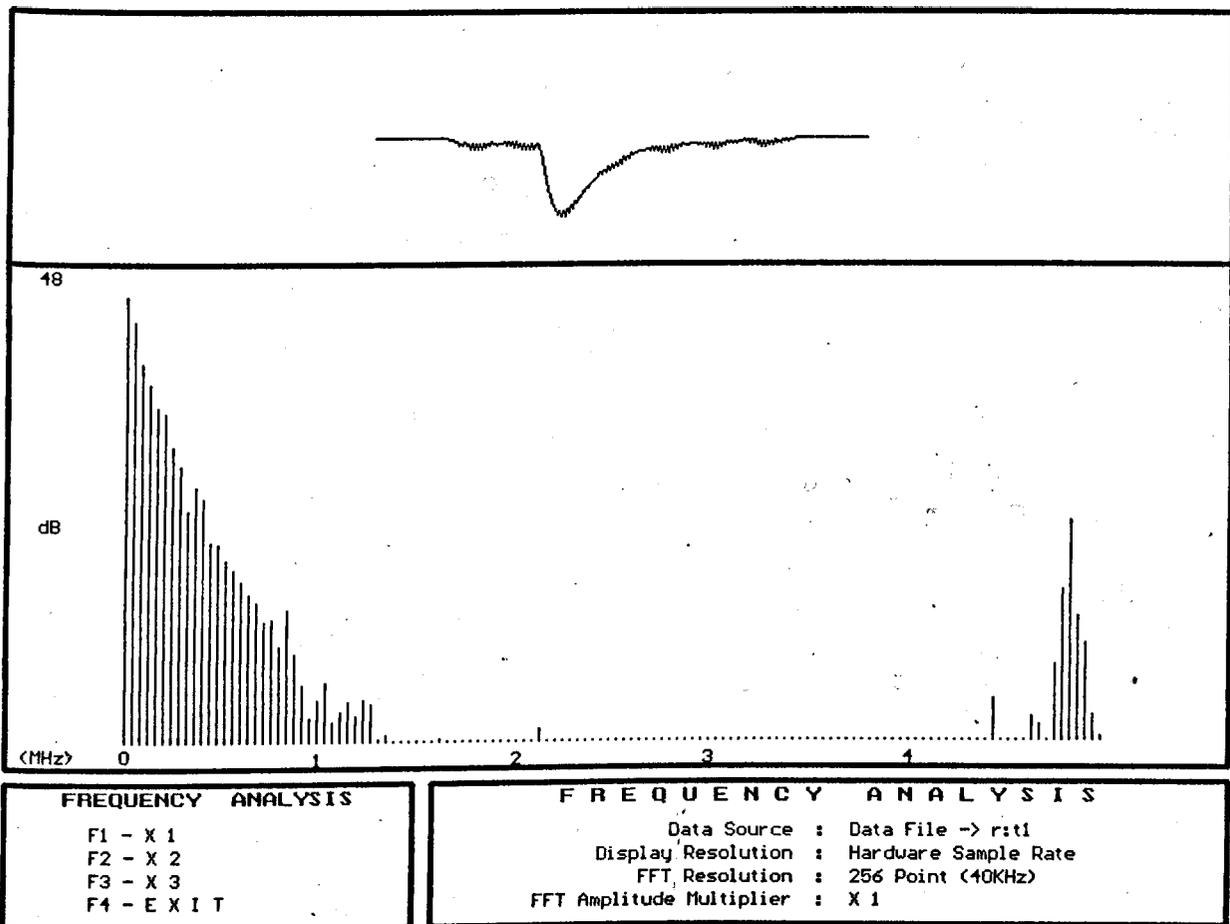
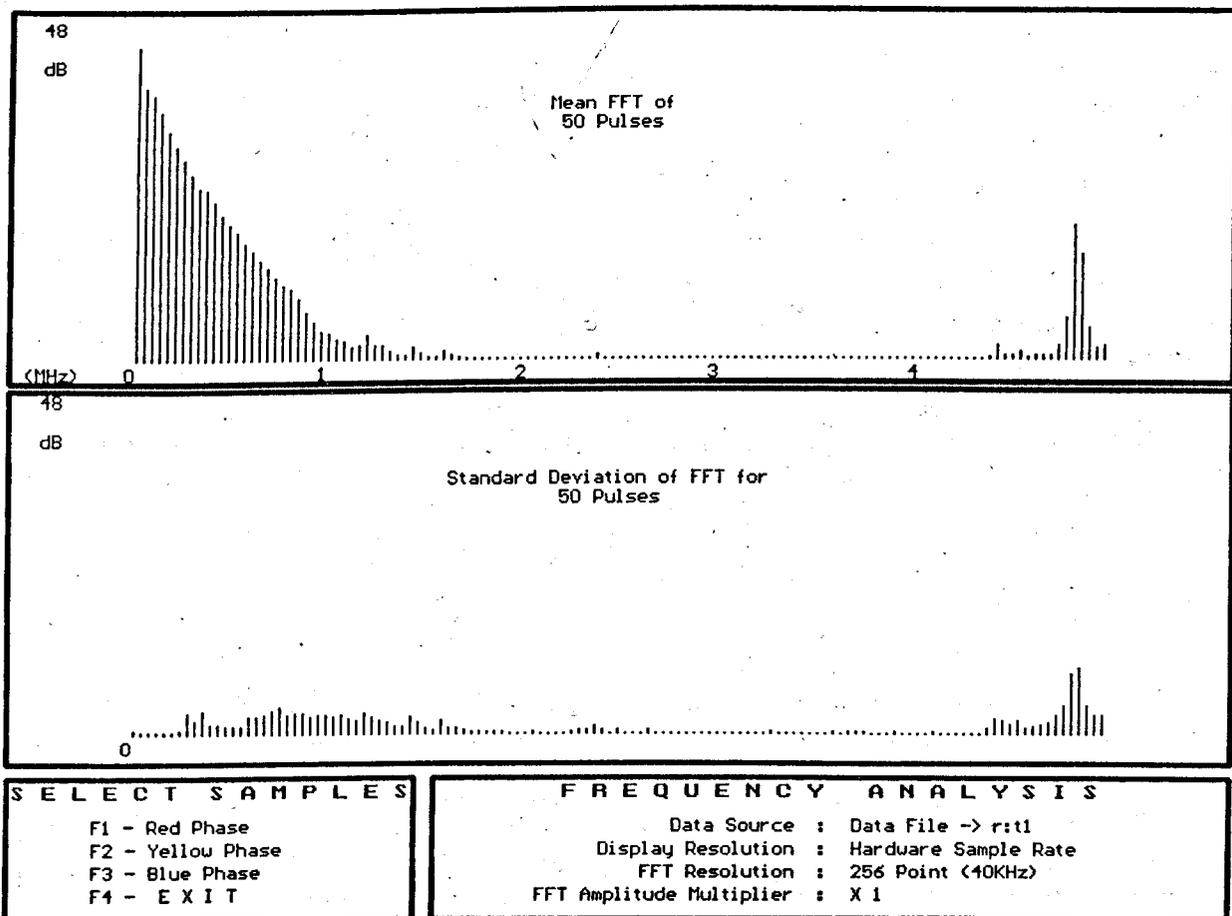
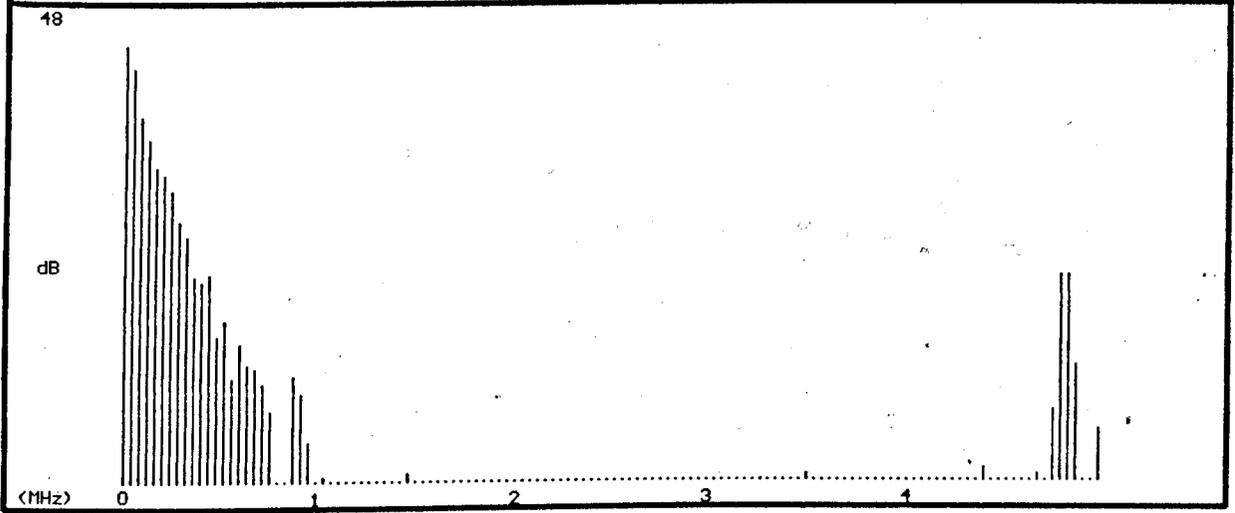
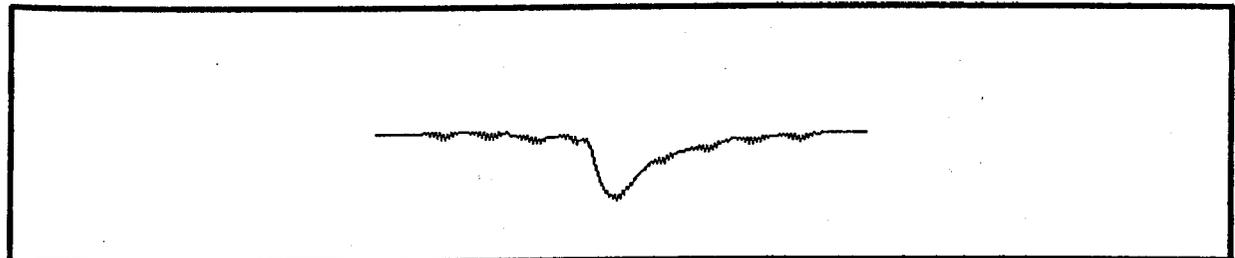


Figure 5.25 FFT Results for 1 Coil Inside Stator

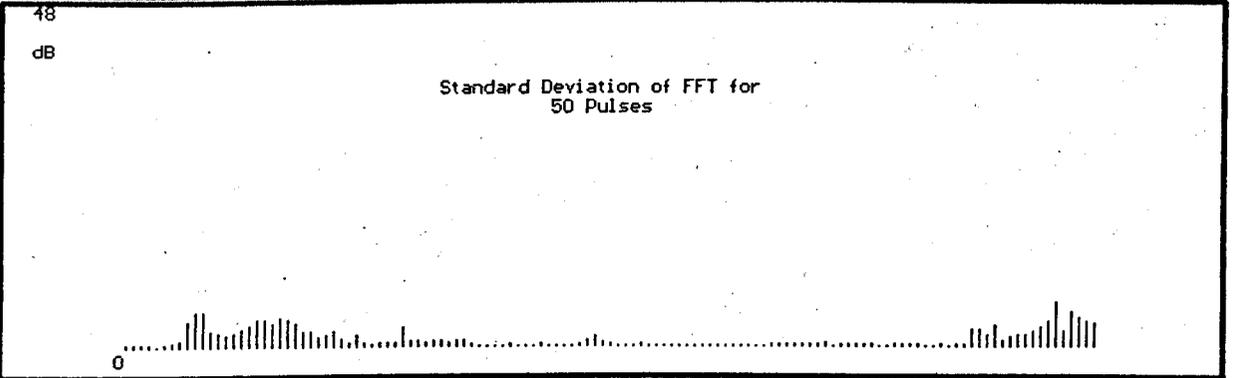
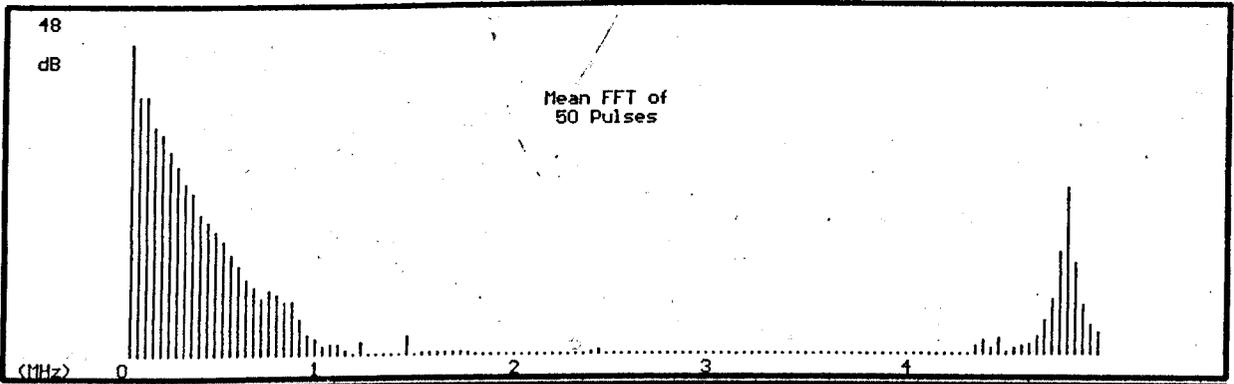




FREQUENCY ANALYSIS
 F1 - X 1
 F2 - X 2
 F3 - X 3
 F4 - E X I T

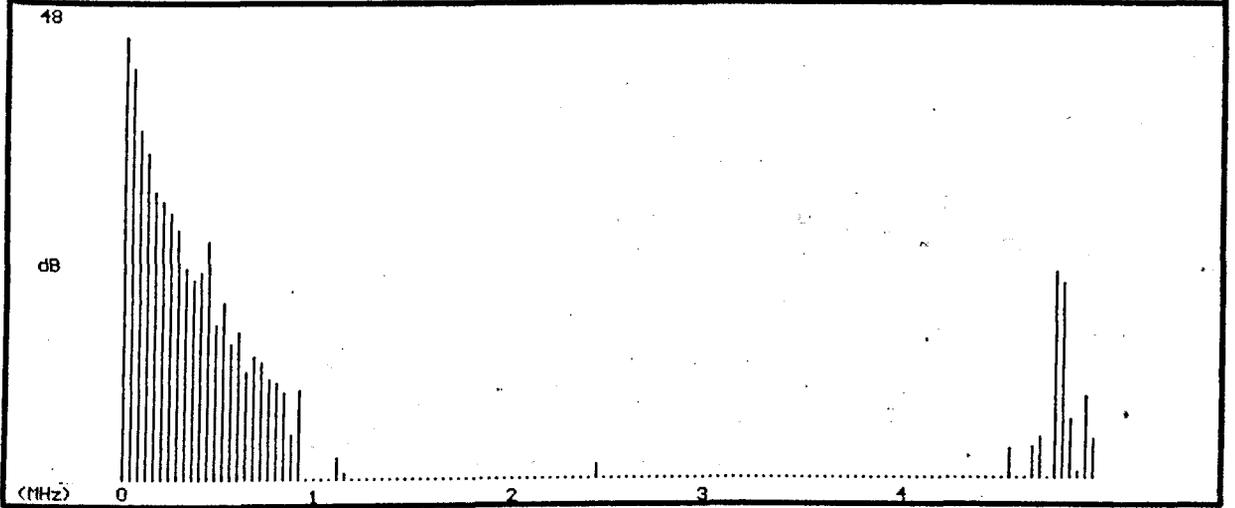
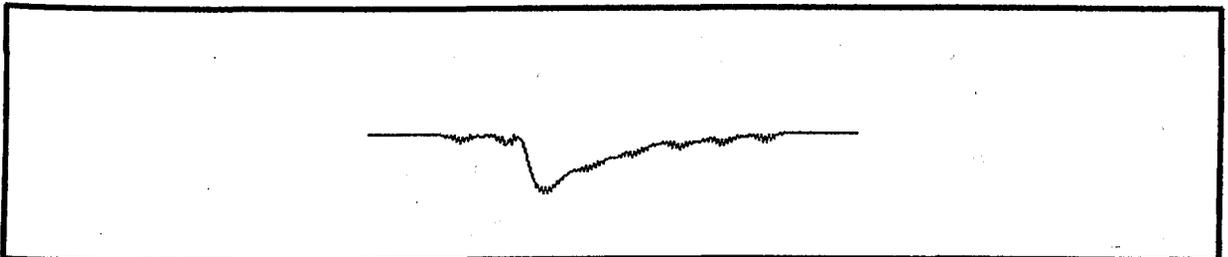
FREQUENCY ANALYSIS
 Data Source : Data File -> r:tl
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

Figure 5.26 FFT Results for 2 Coils Inside Stator



SELECT SAMPLES
 F1 - Red Phase
 F2 - Yellow Phase
 F3 - Blue Phase
 F4 - E X I T

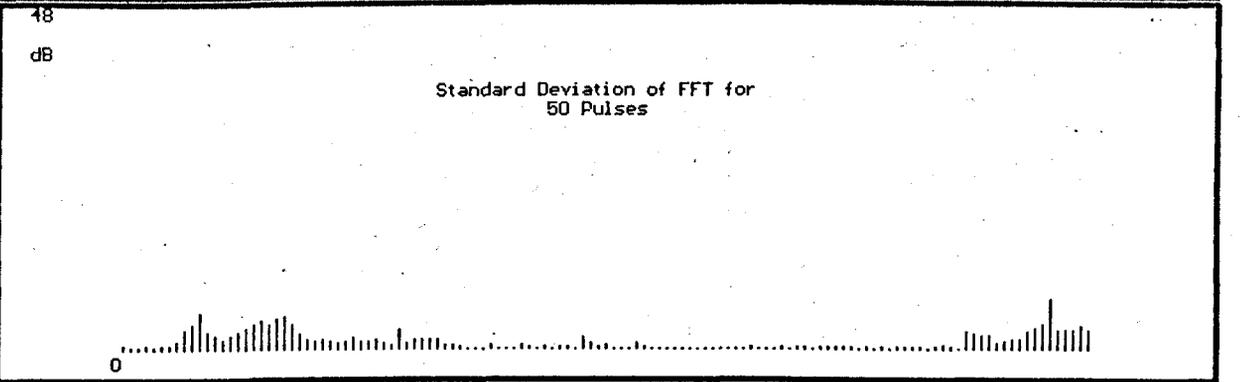
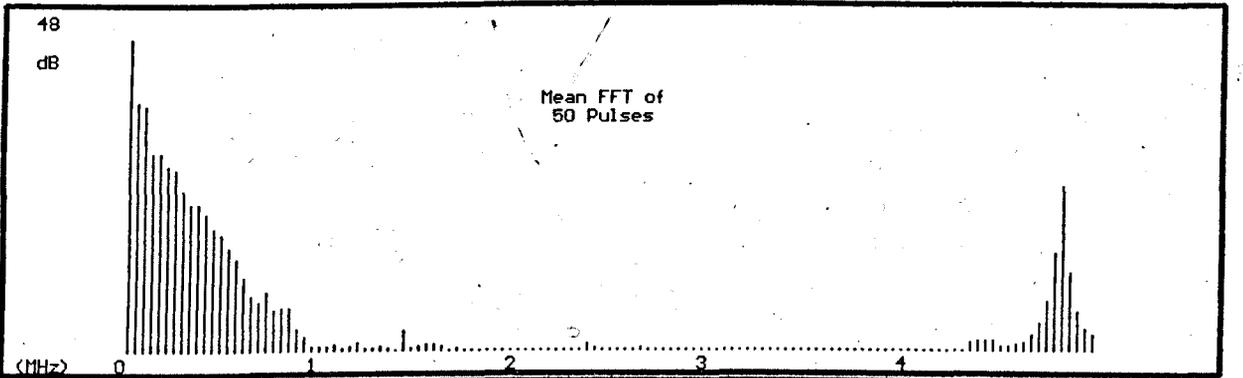
FREQUENCY ANALYSIS
 Unable to open file External
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1



FREQUENCY ANALYSIS
 F1 - X 1
 F2 - X 2
 F3 - X 3
 F4 - E X I T

FREQUENCY ANALYSIS
 Data Source : Data File -> r:t1
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

Figure 5.27 FFT Results for 3 Coils Inside Stator



S E L E C T S A M P L E S
 F1 - Red Phase
 F2 - Yellow Phase
 F3 - Blue Phase
 F4 - E X I T

FREQUENCY ANALYSIS
 Data Source : Data File -> r:t1
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

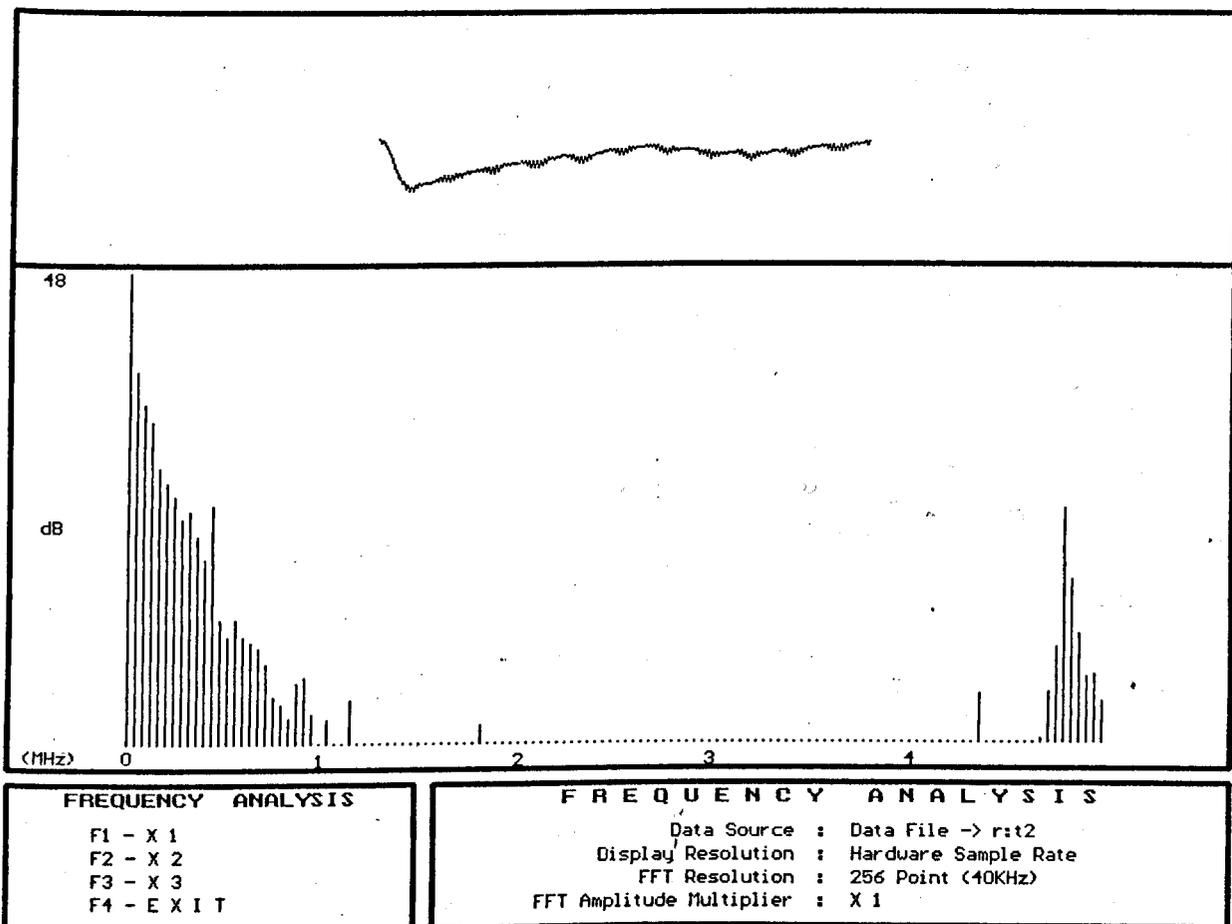
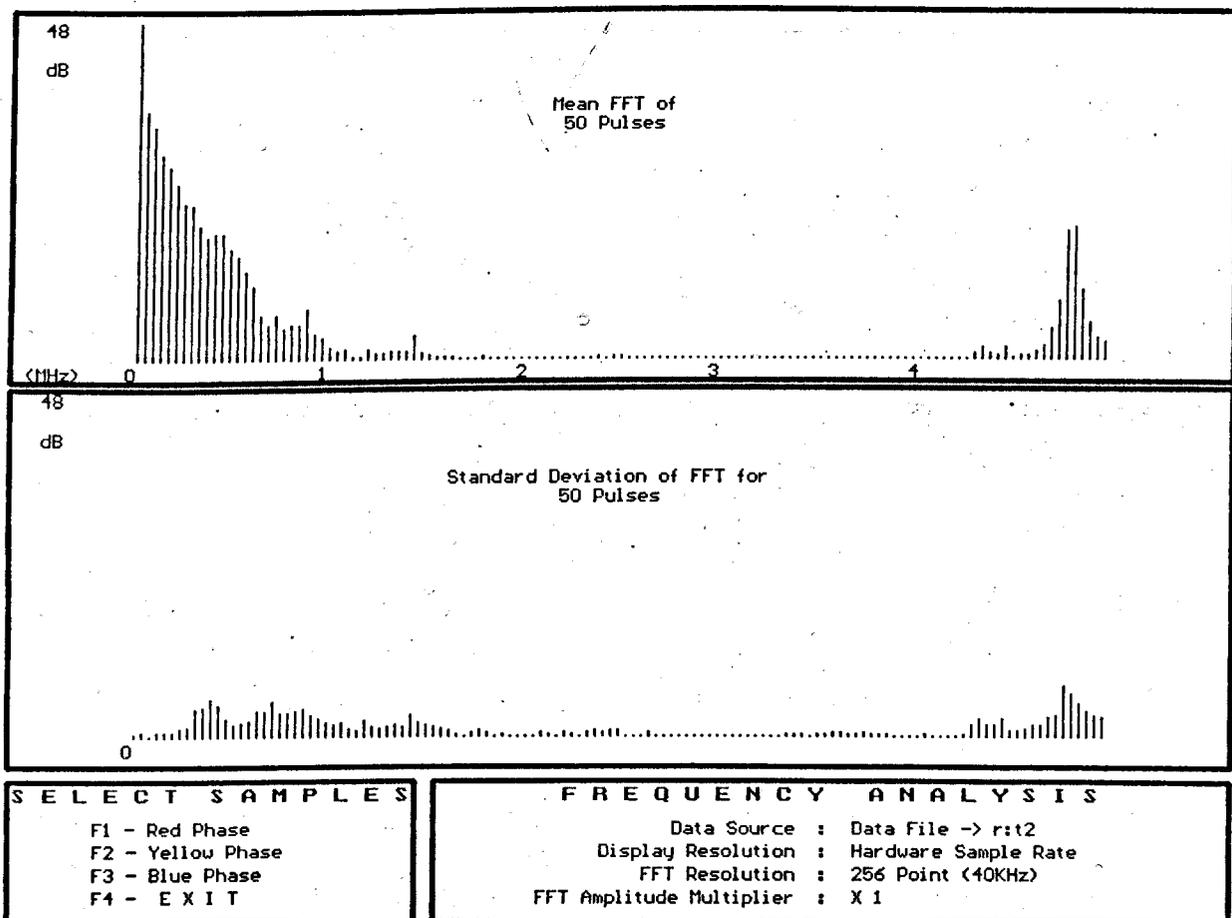
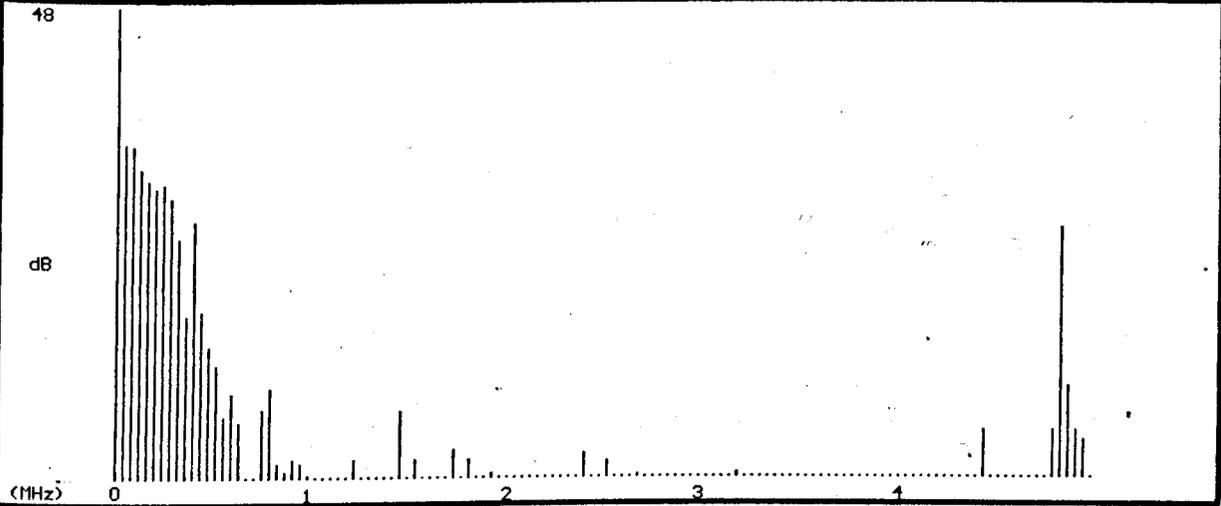
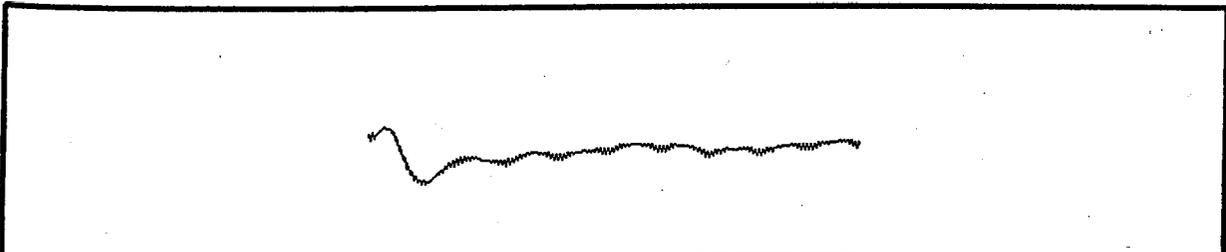


Figure 5.28 FFT Results for 4 Coils Inside Stator

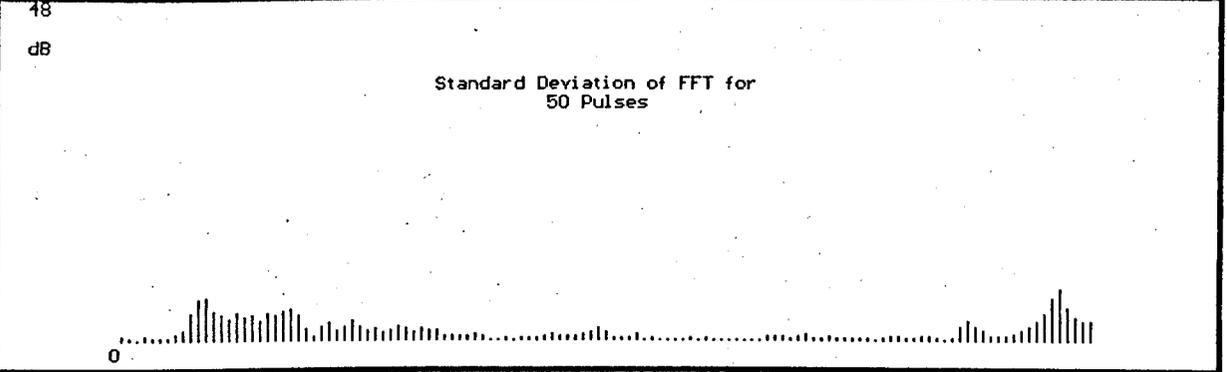
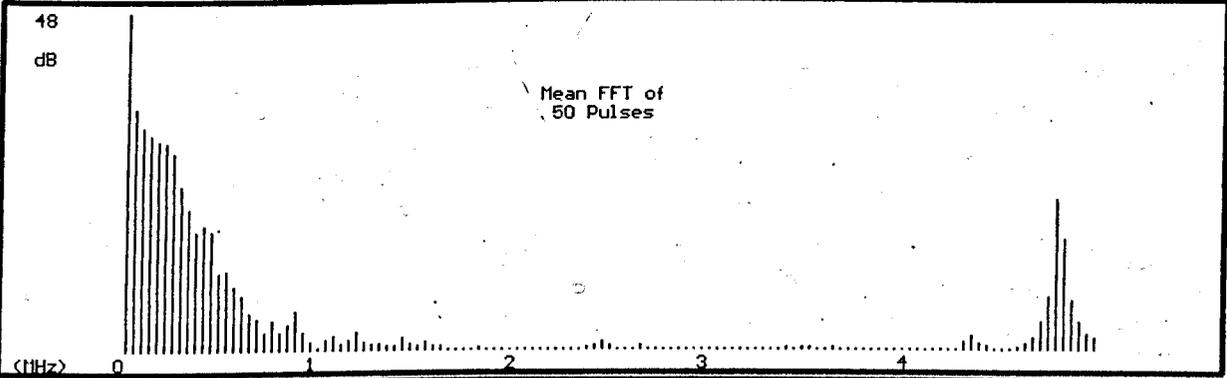




FREQUENCY ANALYSIS
 F1 - X 1
 F2 - X 2
 F3 - X 3
 F4 - E X I T

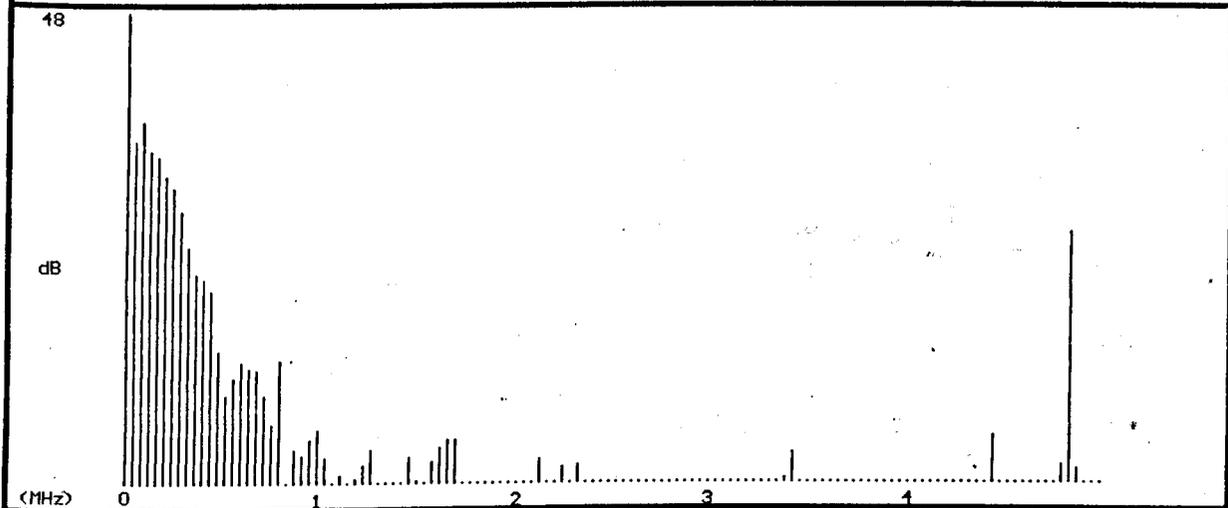
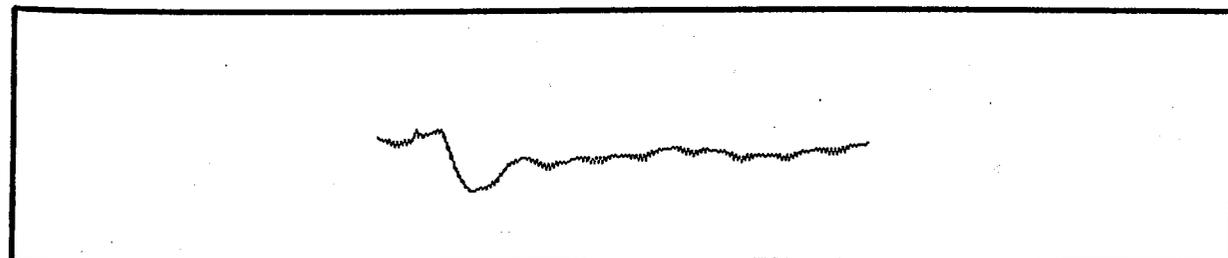
FREQUENCY ANALYSIS
 Data Source : Data File -> r:t2
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

Figure 5.29 FFT Results for 5 Coils Inside Stator



S E L E C T S A M P L E S
 F1 - Red Phase
 F2 - Yellow Phase
 F3 - Blue Phase
 F4 - E X I T

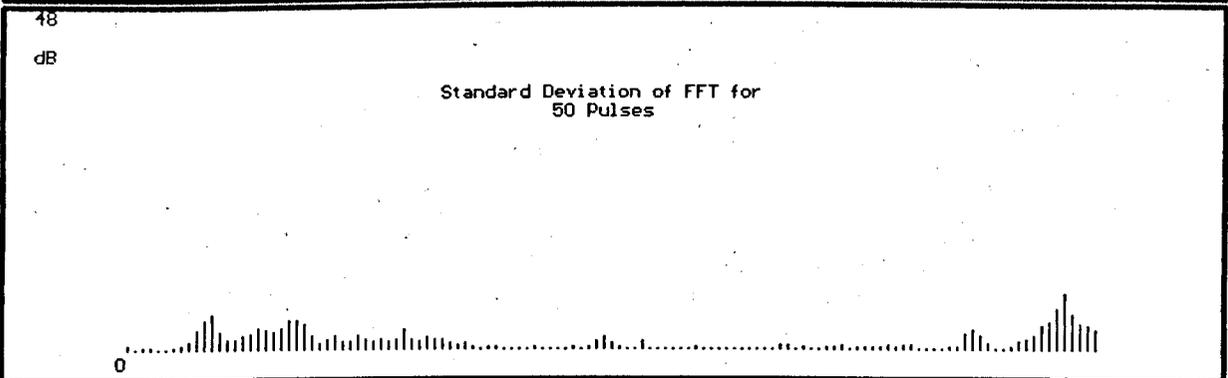
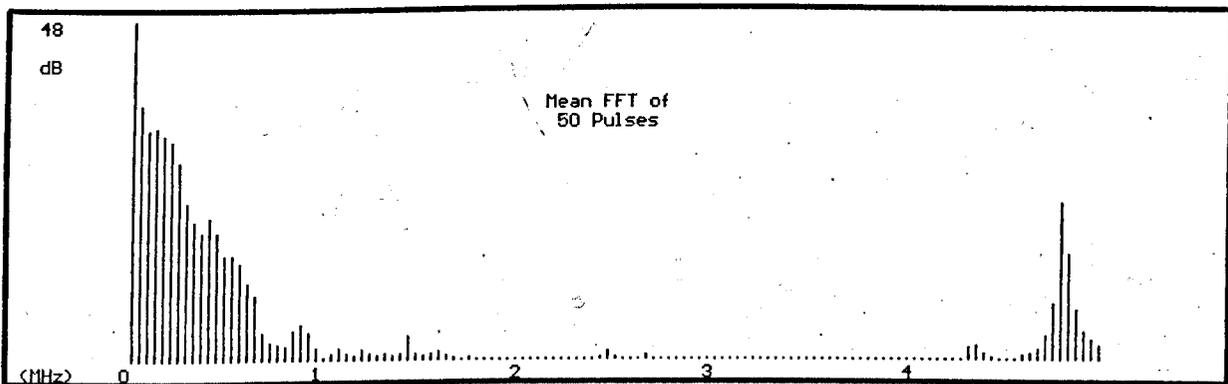
FREQUENCY ANALYSIS
 Data Source : Data File -> r:t2
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1



FREQUENCY ANALYSIS
 F1 - X 1
 F2 - X 2
 F3 - X 3
 F4 - E X I T

FREQUENCY ANALYSIS
 Data Source : Data File -> r:t2
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

Figure 5.30 FFT Results for 6 Coils Inside Stator



SELECT SAMPLES
 F1 - Red Phase
 F2 - Yellow Phase
 F3 - Blue Phase
 F4 - E X I T

FREQUENCY ANALYSIS
 Data Source : Data File -> r:t2
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

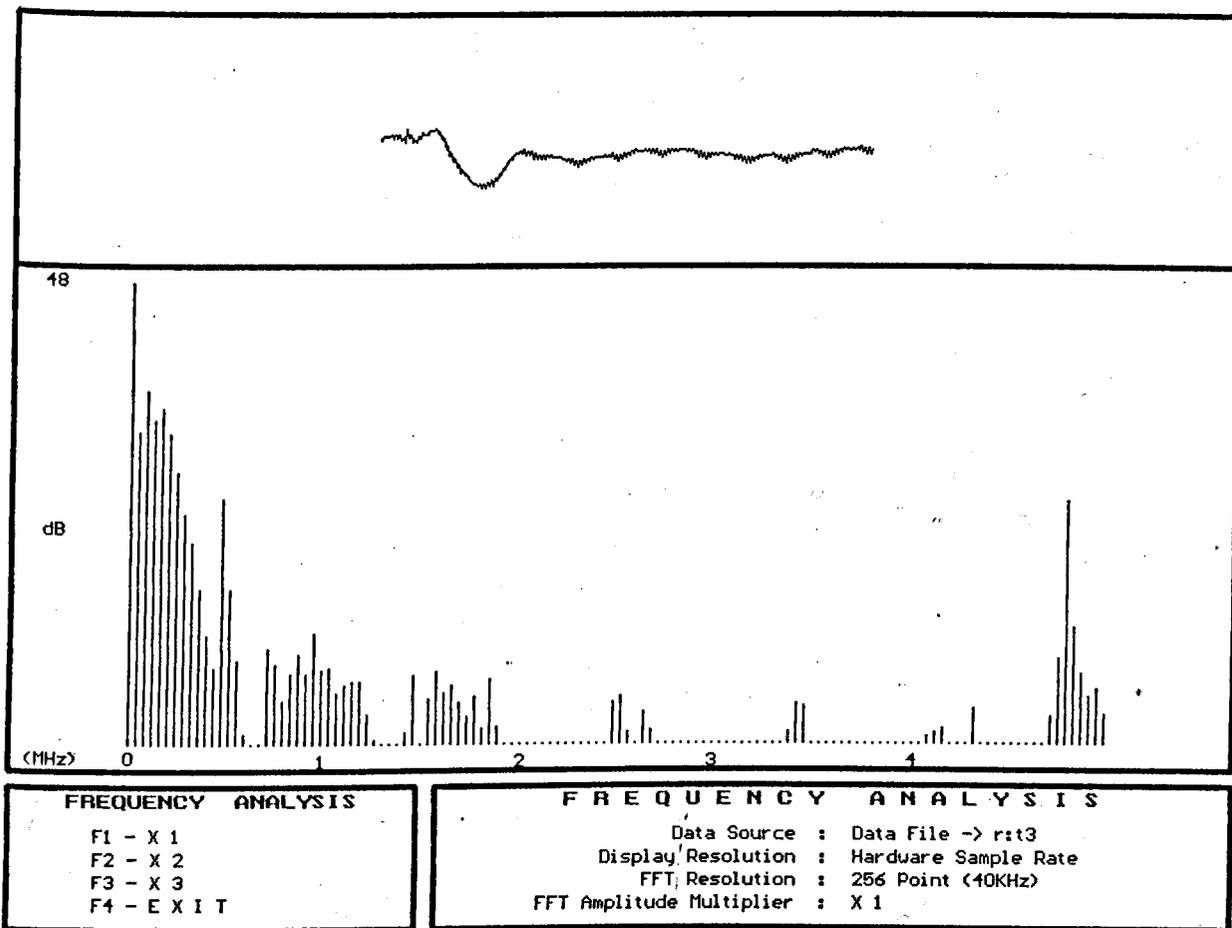
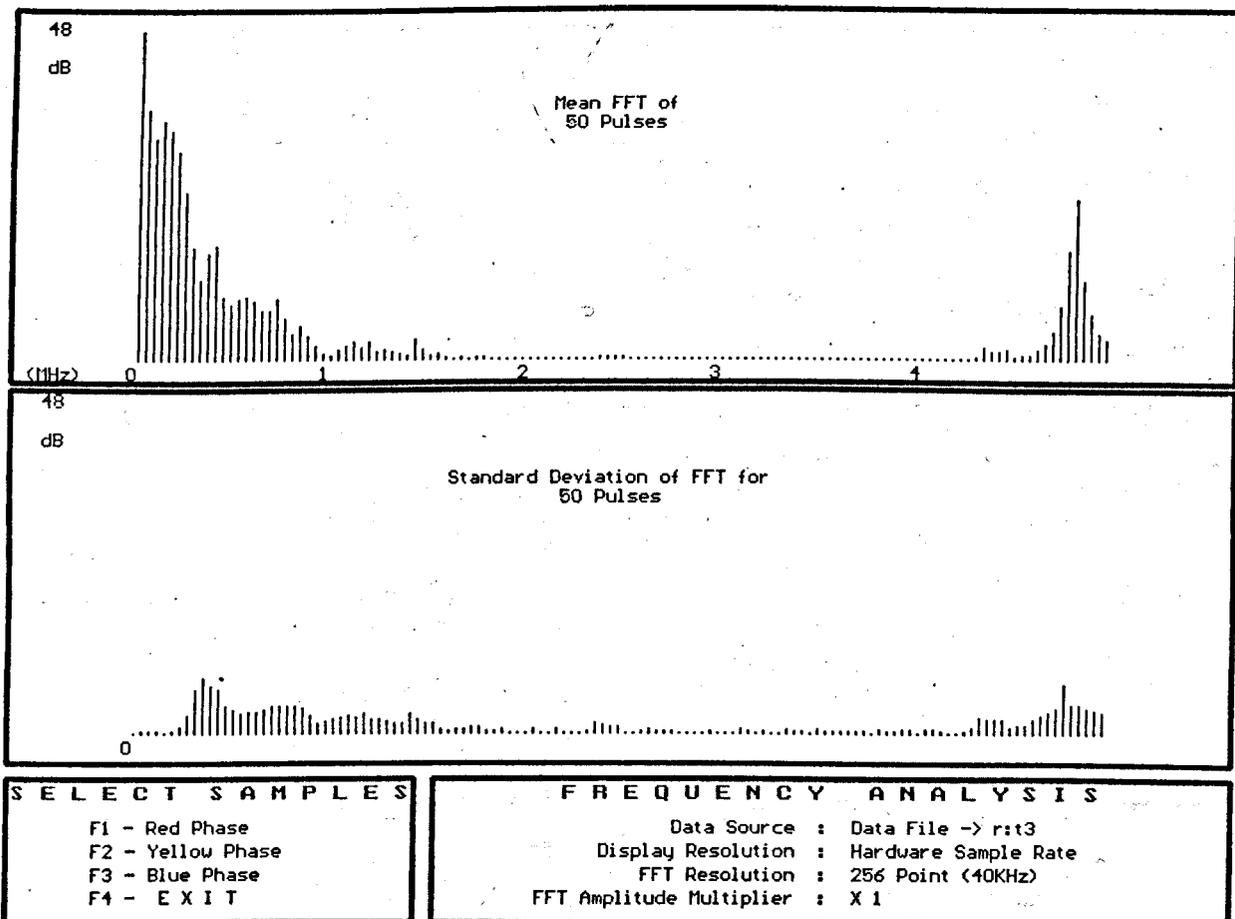
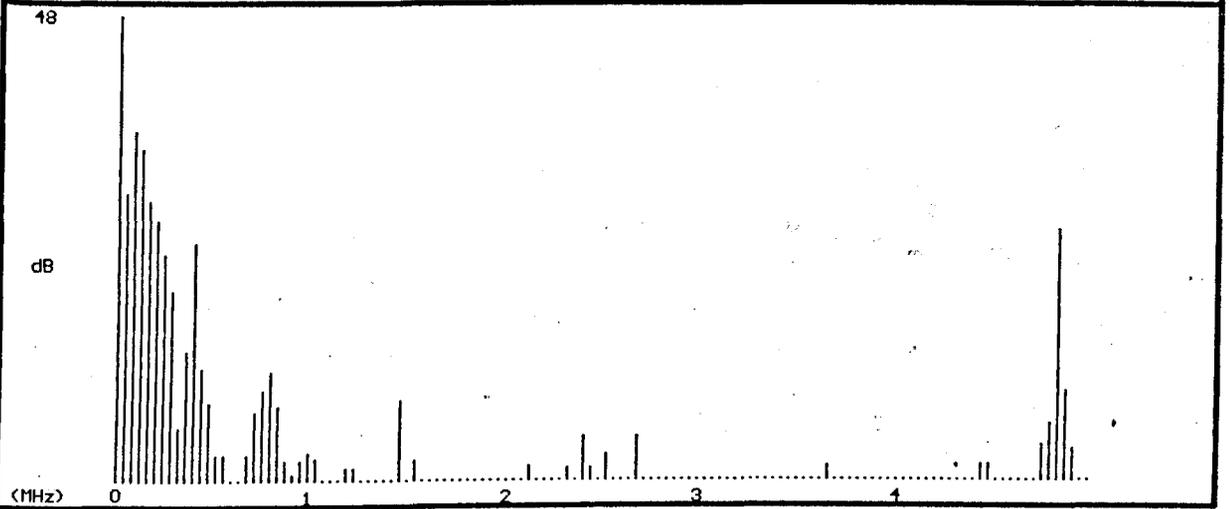
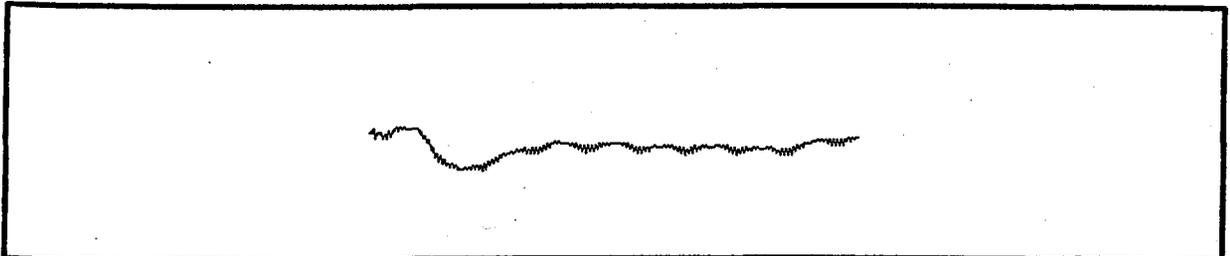


Figure 5.31 FFT Results for 7 Coils Inside Stator





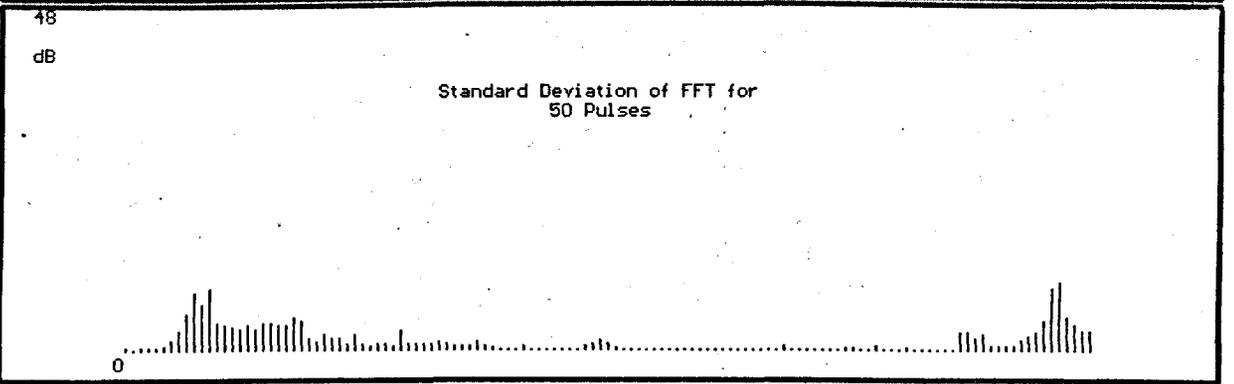
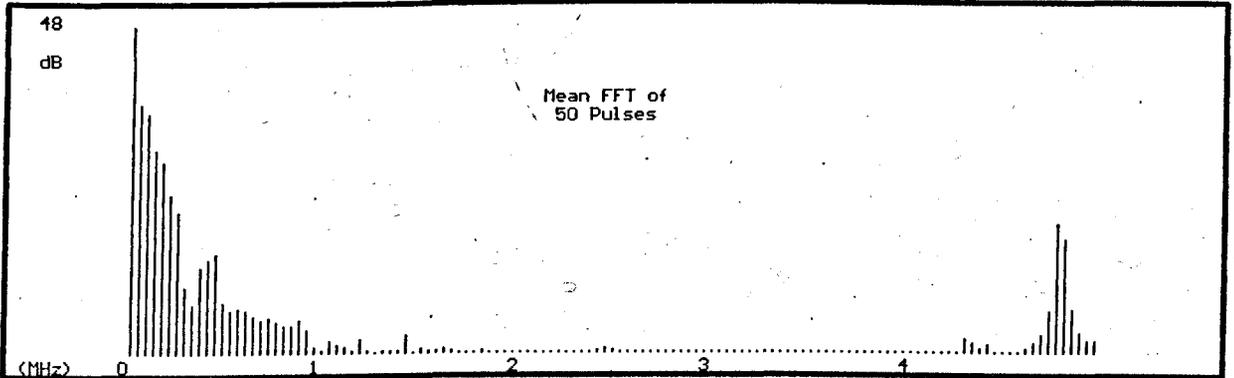
FREQUENCY ANALYSIS

F1 - X 1
 F2 - X 2
 F3 - X 3
 F4 - E X I T

FREQUENCY ANALYSIS

Data Source : Data File -> r:t3
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

Figure 5.32 FFT Results for 8 Coils Inside Stator



SELECT SAMPLES

F1 - Red Phase
 F2 - Yellow Phase
 F3 - Blue Phase
 F4 - E X I T

FREQUENCY ANALYSIS

Data Source : Data File -> r:t3
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

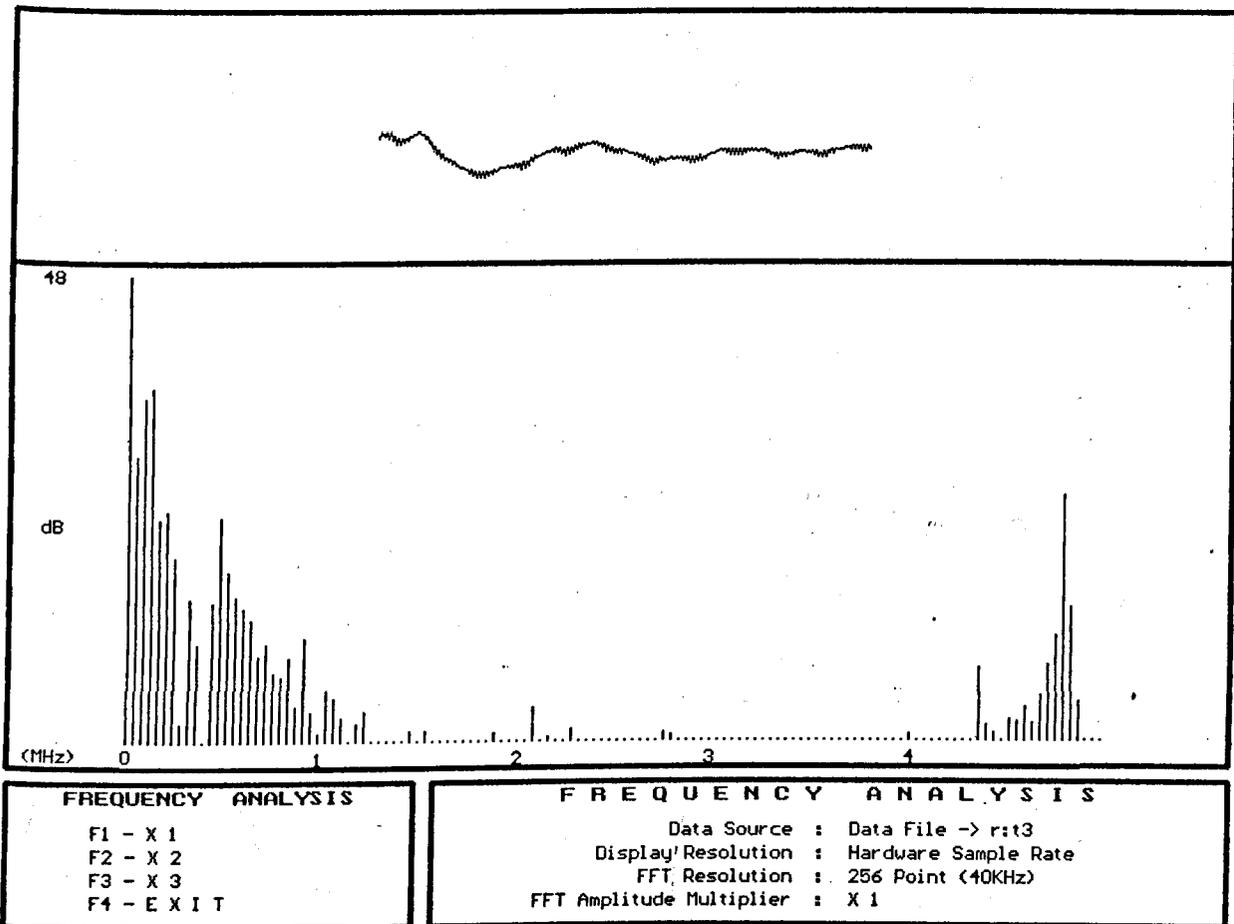
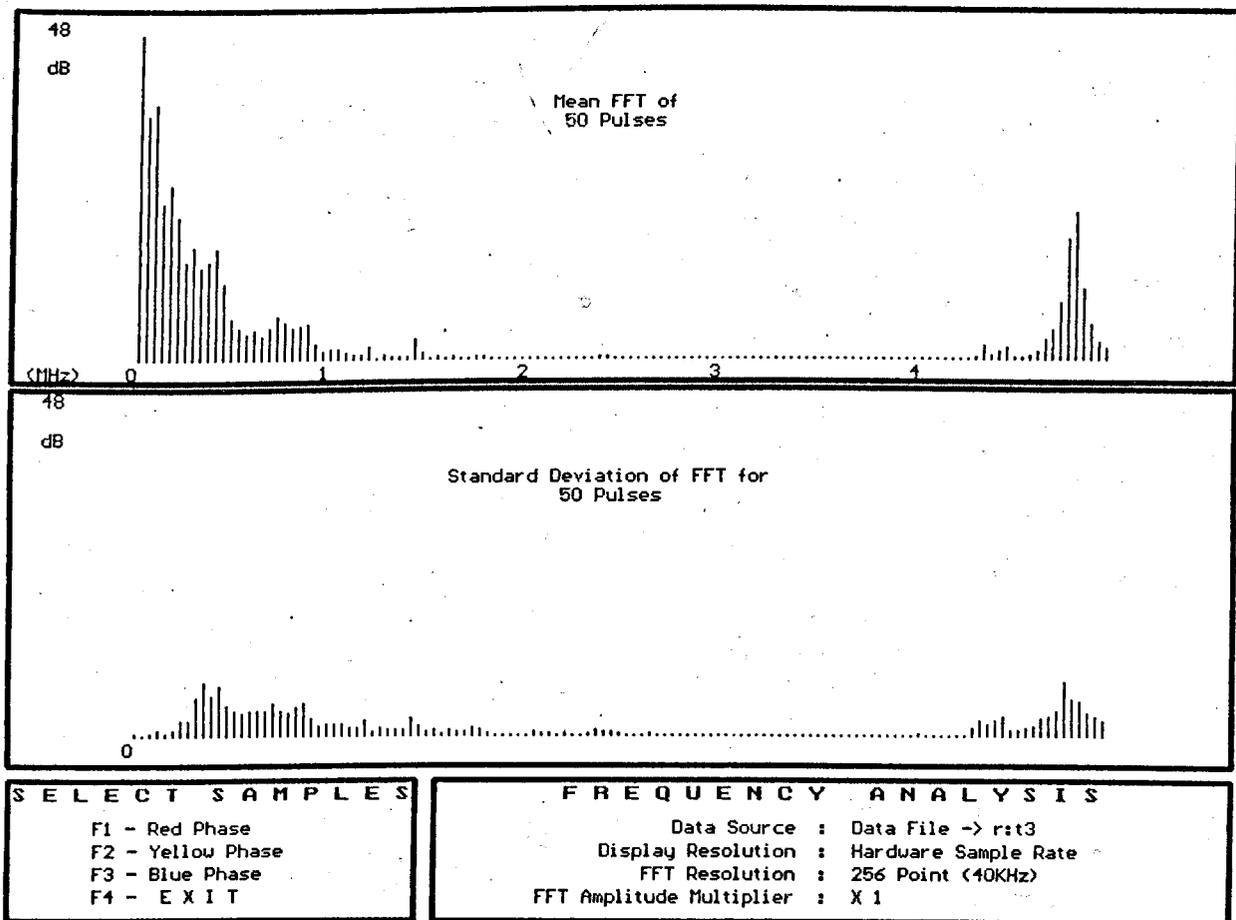
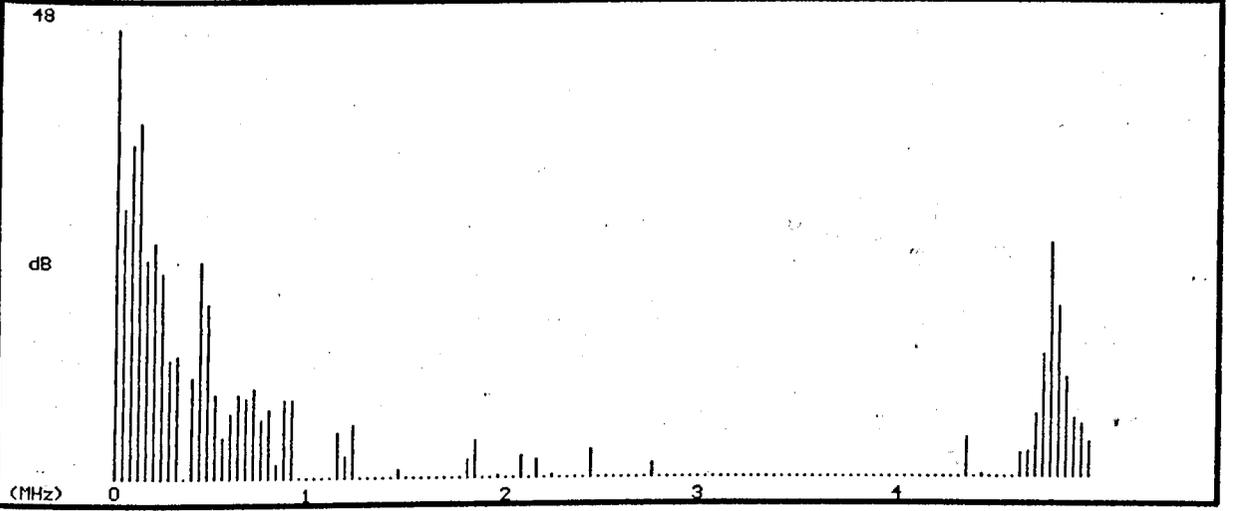
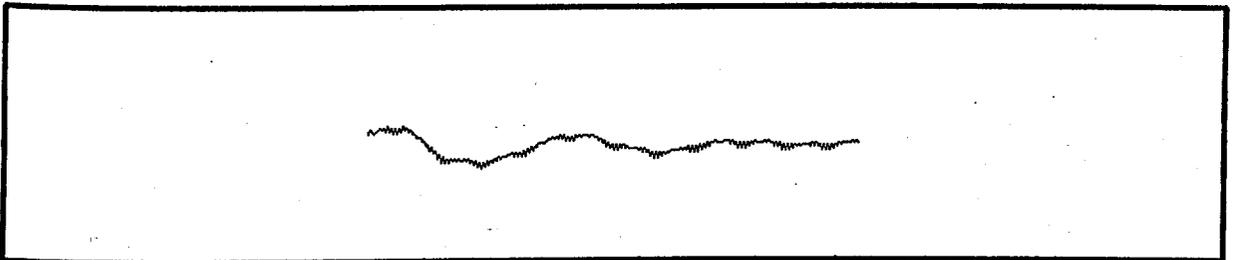


Figure 5.33 FFT Results for 9 Coils Inside Stator

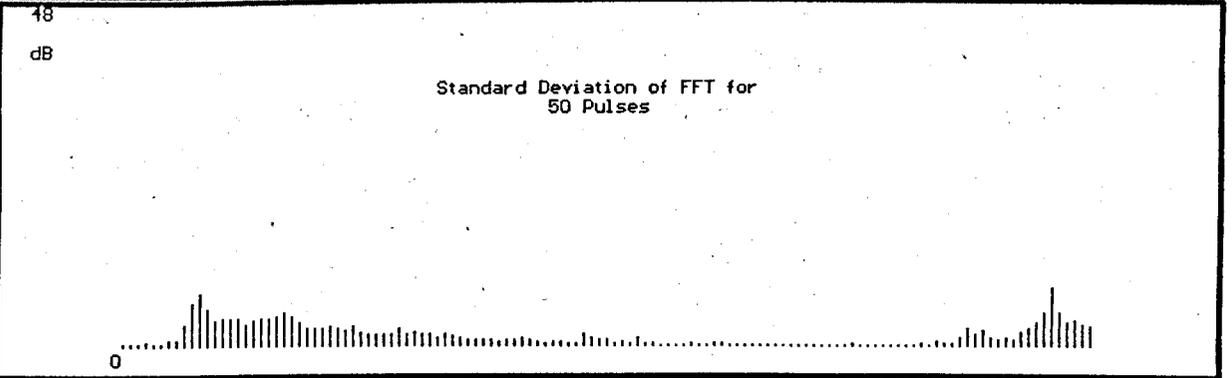
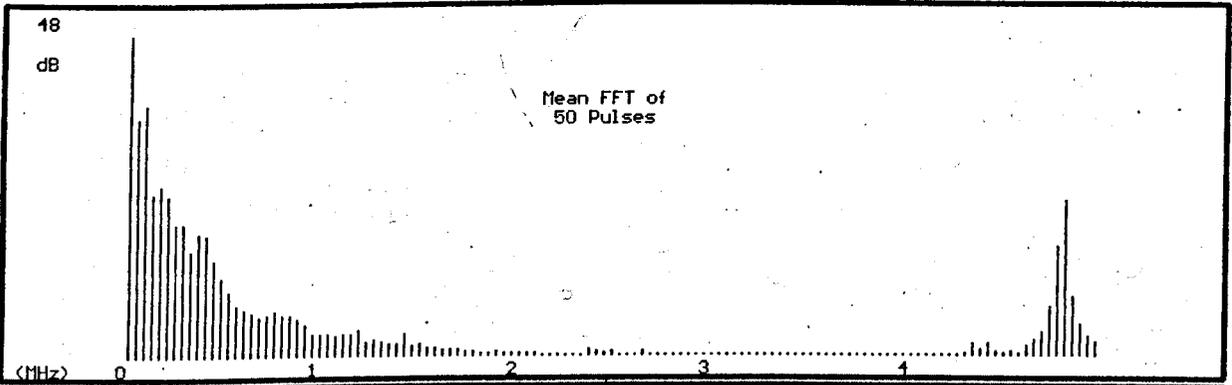




FREQUENCY ANALYSIS
 F1 - X 1
 F2 - X 2
 F3 - X 3
 F4 - E X I T

FREQUENCY ANALYSIS
 Data Source : Data File -> r:test10
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

Figure 5.34 FFT Results for 10 Coils Inside Stator



SELECT SAMPLES
 F1 - Red Phase
 F2 - Yellow Phase
 F3 - Blue Phase
 F4 - E X I T

FREQUENCY ANALYSIS
 Data Source : Data File -> r:test10
 Display Resolution : Hardware Sample Rate
 FFT Resolution : 256 Point (40KHz)
 FFT Amplitude Multiplier : X 1

Before the individual graphs are discussed, it is apparent from all of the plots that some high frequency interference has been superimposed onto the experimental data. This noise, from the ensuing spectral analyses, has components of sufficiently high frequencies to cause concern about the prospect of aliasing on the lower frequency bands.

The digitised interference can be plainly seen in the upper time plots manifesting itself as a number of small pulses, typically 2 or 3 pixels in height, re-appearing periodically. It is most likely caused by high frequency interference from the host computer finding its way into the power supplies of the digitising hardware. This means that it can bypass the anti-alias filters and impose itself directly on the analogue to digital converters themselves. However, the frequency plots plainly show that the interference does tend to be attenuated as the spectrum approaches 5MHz. This, coupled with the fact that the interference is common to all the data, led to the assumption that the interference problem was not too significant in this investigation.

Examining Figure 5.25, the FFT results for discharges injected one coil distant from the line-end of the stator, it can be seen that the plot of a single FFT analysis is somewhat different to the average analysis. Although the automatic averaging process takes a considerable amount of time, comparison of the two plots shows that it is necessary in order to obtain relatively consistent results, confirming that individual spectra are indeed "spectral estimates". Of course, if a larger population is used in the averaging process the results should become more representative. The exact number of FFT analyses used had to be a compromise between computation time and ensuing precision.

The results for Figure 5.25 show that pulses injected one coil distant from the termination have decreasing frequency components up to about 1.7MHz. However, comparing the mean and standard deviation of the components above 1MHz it can be seen that the standard deviations become larger than the mean components, indicating that individual pulses have considerably different spectral lines at these frequencies.

Close comparison of Figures 5.25 and 5.26 show that, as expected, when the injected pulses travel further their higher frequencies are attenuated. Figure 5.26 shows

spectral components only up to about 1.4MHz, with corresponding decreases in some of the lower components also.

Although Figure 5.27 shows a slight decrease in higher frequency components, it is not as marked as the decrease between coils one and two.

This trend is further noted in the graphs of Figure 5.28, four coils inside the stator. Comparing this to the results for the previous coil, it can be seen that the graphs of individual pulses show marked differences, whilst those of the averaged pulses show very little.

Close comparison of the results obtained for coils four and five, Figures 5.28 and 5.29, reveals small but significant differences from about 500kHz-1MHz

The initial results of the frequency analysis, whilst showing a difference between successive frequency profiles, does not seem to demonstrate a consistent change in specific frequency components that would allow a fixed algorithm to determine the pulses inception location. However, this conclusion relies purely on a visual interpretation of the results; chapter 6 will deal specifically with detailed analysis of the results.

5.5.2 Maximum Entropy Method

The FFT is not the only way to estimate the power spectrum of a process, nor is it necessarily the best way for all purposes. The method that has found most widespread use as an alternative to the FFT is that of Autoregressive (AR) spectral analysis. Apart from giving improved resolution and reduced leakage, which make the AR method suitable for time-varying signals, it also offers large data reduction facilities, making it particularly suitable in speech synthesis work. In addition, AR spectral analysis has inherent smoothing of the spectra, contrasting with the large spectral variance associated with the FFT.

Like other methods of spectral estimation, the AR technique has quite a long history. Yule [23] and Walker [24] used AR models for forecasting trends in economic time series and determining periodicity's in sunspot data. Burg[25] introduced the

Maximum Entropy Method (MEM) for the analysis of geophysical data, while Parzen [26] formally proposed AR spectral estimation. Since then AR spectral analysis has been applied in radar, sonar, imaging, radio astronomy, biomedicine, oceanography, ecological systems etc.

The MEM estimation, originated by Burg [25], is a function of continuously varying frequency. There is no special significance to specific, equally spaced frequencies as there was in the FFT case. In fact, since the MEM may have sharp spectral features, it can be evaluated on a very fine mesh near specific features and more coarsely farther away.

What this means in practice, is that for a small data sample length, say 256 samples, the resolution of the resulting spectral lines are not limited to 40kHz steps, as was the case for the FFT, but rather can be as small or as large as desired. This is because the power spectrum is estimated in two parts: firstly the Maximum Entropy Method spectral estimation is calculated using the data sample and by specifying the desired number of coefficients to be produced by the estimation. Once the coefficients have been computed, a different function can evaluate the estimated power spectrum as a function of $f\Delta$ (the frequency times the sampling interval). This means that the power spectrum can be estimated for arbitrarily small steps in frequency, merely by supplying small changes in $f\Delta$ to the power spectrum evaluation algorithm.

The number of coefficients produced by MEM spectral estimation has a direct bearing on the resulting power spectrum, and, of course, the time taken to compute the estimate. The estimation produces poles, corresponding to infinite power spectral density, on the unit z-circle. Such poles can provide an accurate representation for the underlying power spectra being estimated. The number of poles produced corresponds to the desired order, or number of coefficients, of the MEM estimation and, correspondingly, more poles mean more "peaky" power spectra, sometimes resulting in spurious spectral features. If the desired order of the estimate is too small, the result is a very smooth power spectrum, but crucial details may be lost. Care must therefore be taken, when applying the Maximum Entropy Method, to specify a suitable number of poles such that spectral details are preserved, but not at the expense of spurious spectral peaks.

The Maximum Entropy Method algorithms used here were again adapted from those produced in the book "Numerical Recipes in C" [20]. To evaluate the MEM algorithms and produce a meaningful comparison of performance, the same data that was used by the FFTs was applied. Figure 5.35 shows the MEM estimate of a pulse originating from coil 1, with the MEM order 80 and 400 spectral lines, each spaced 12.5kHz apart.

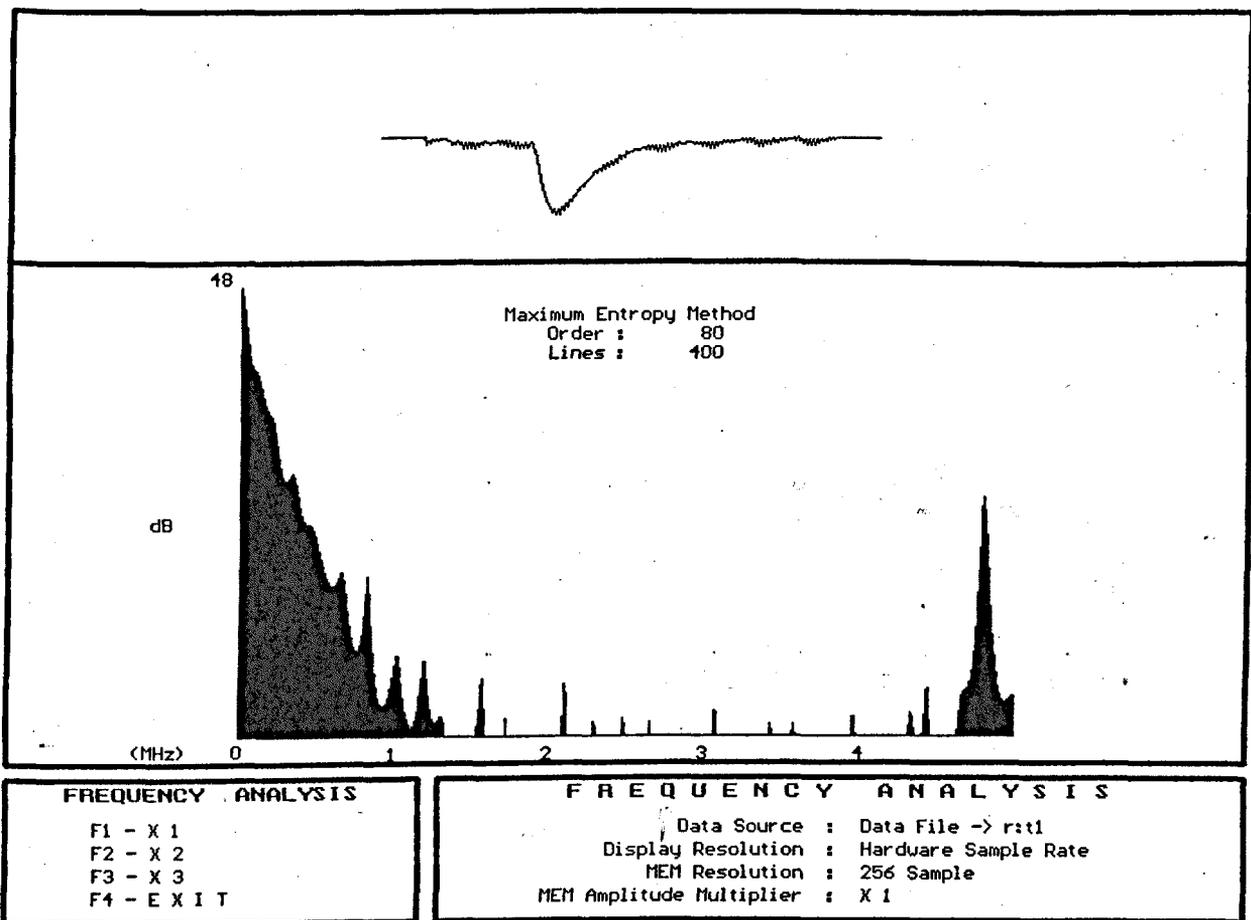


Figure 5.35 MEM Estimate of Pulse From Coil 1: 80 Poles
 (top part shows time-domain pulse, the bottom part shows MEM Estimate)

The above figure shows the effect of too large an order in the spectral estimation. If this plot is compared with that produced by the FFT, Figure 5.25, the presence of spurious peaks can be noted.

If the number of poles estimated is reduced, say to 40, then the resulting analysis, shown in Figure 5.36, shows a marked decrease in spurious peaks.

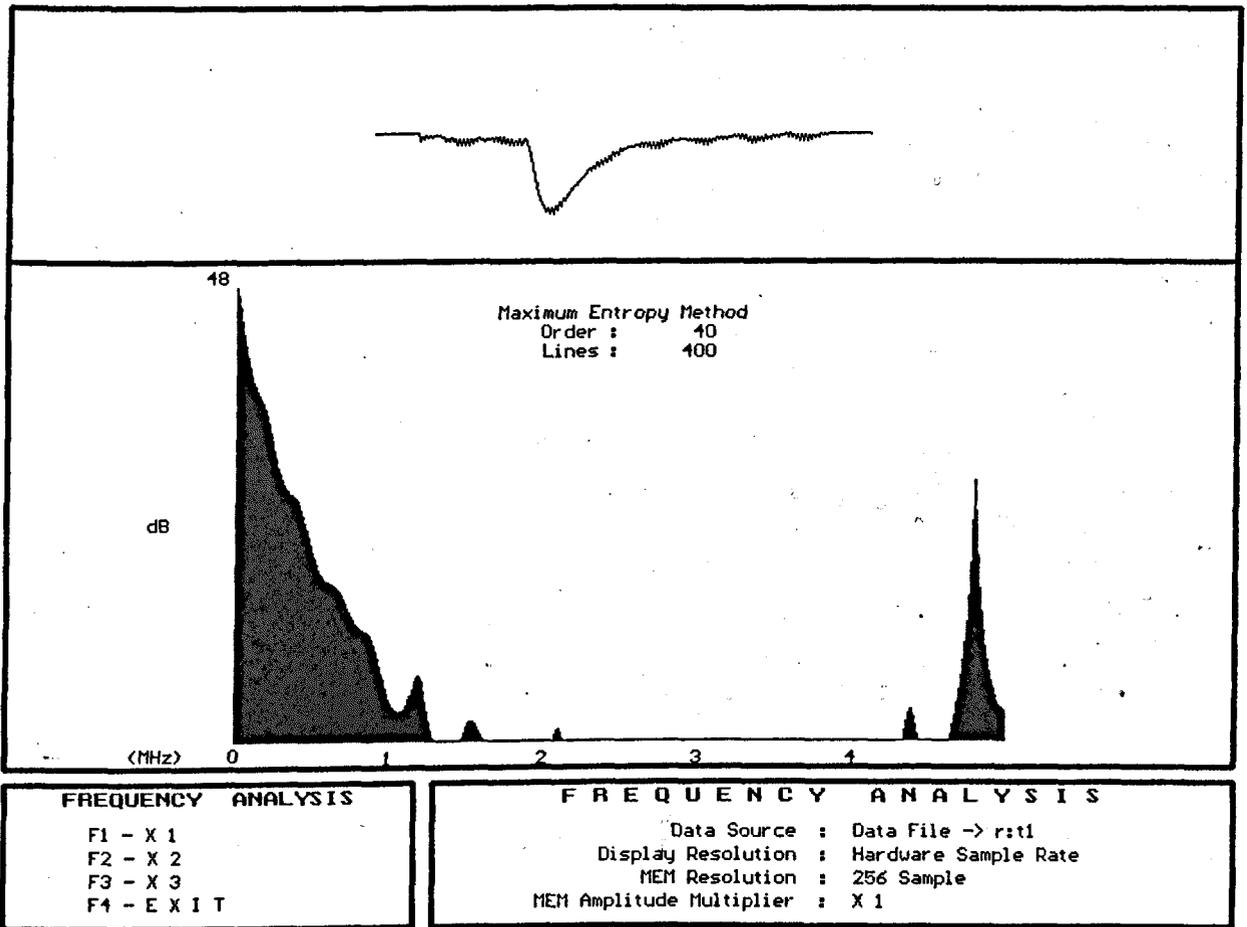


Figure 5.36 MEM Estimate of Pulse From Coil 1: 40 Poles

Although most of the spurious peaks have gone, comparison with the FFT result, Figure 5.25, shows that the peaks that remain are more pronounced, indicating that the order is still too high. Figure 5.37 shows the resulting estimate if the order is reduced further to 20.

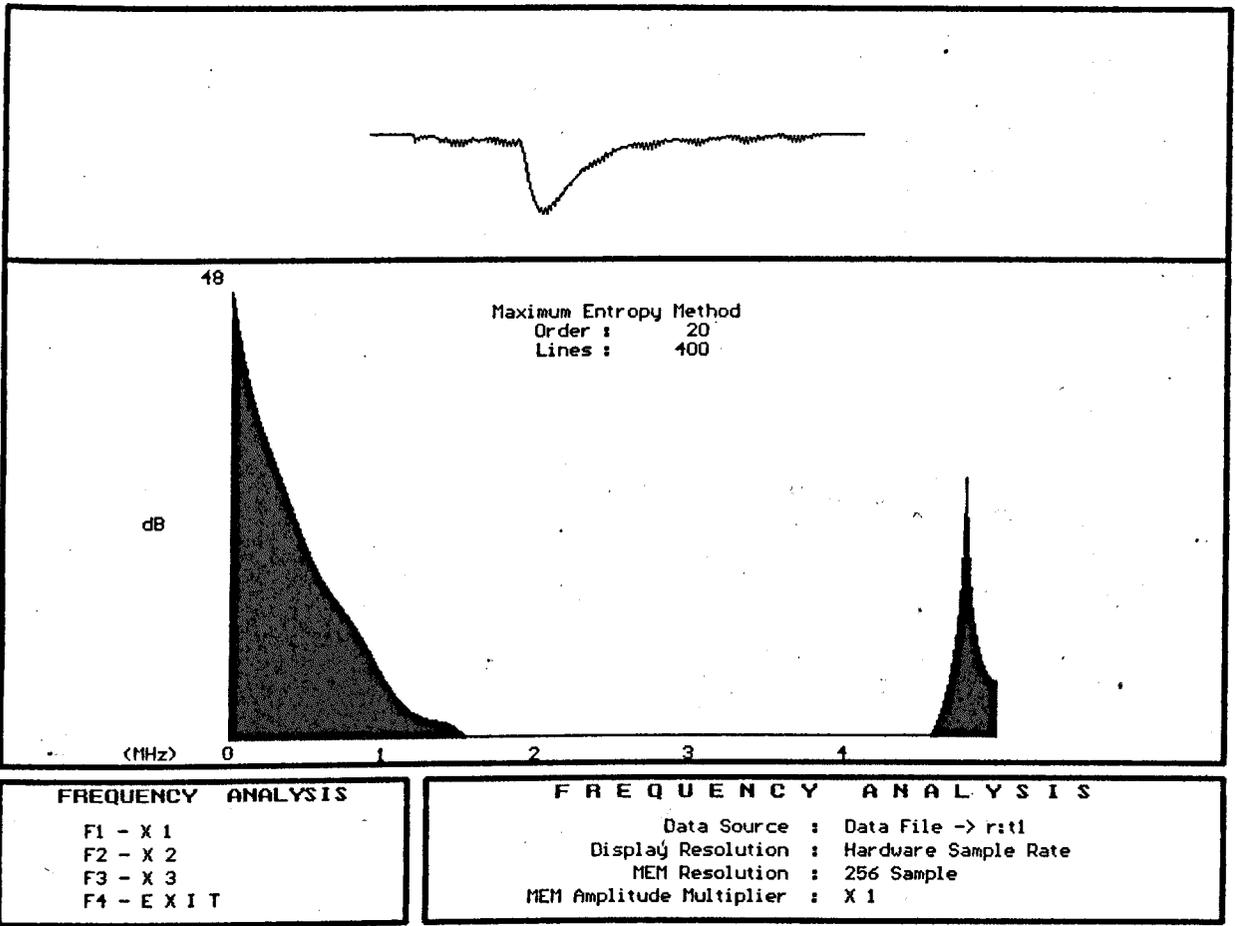


Figure 5.37 MEM Estimate of Pulse From Coil 1: 20 Poles

The above plot is quite similar to that of the FFT analysis, Figure 5.25, but, just for completeness, an estimate was produced using just 10 poles to see the effect, Figure 5.38.

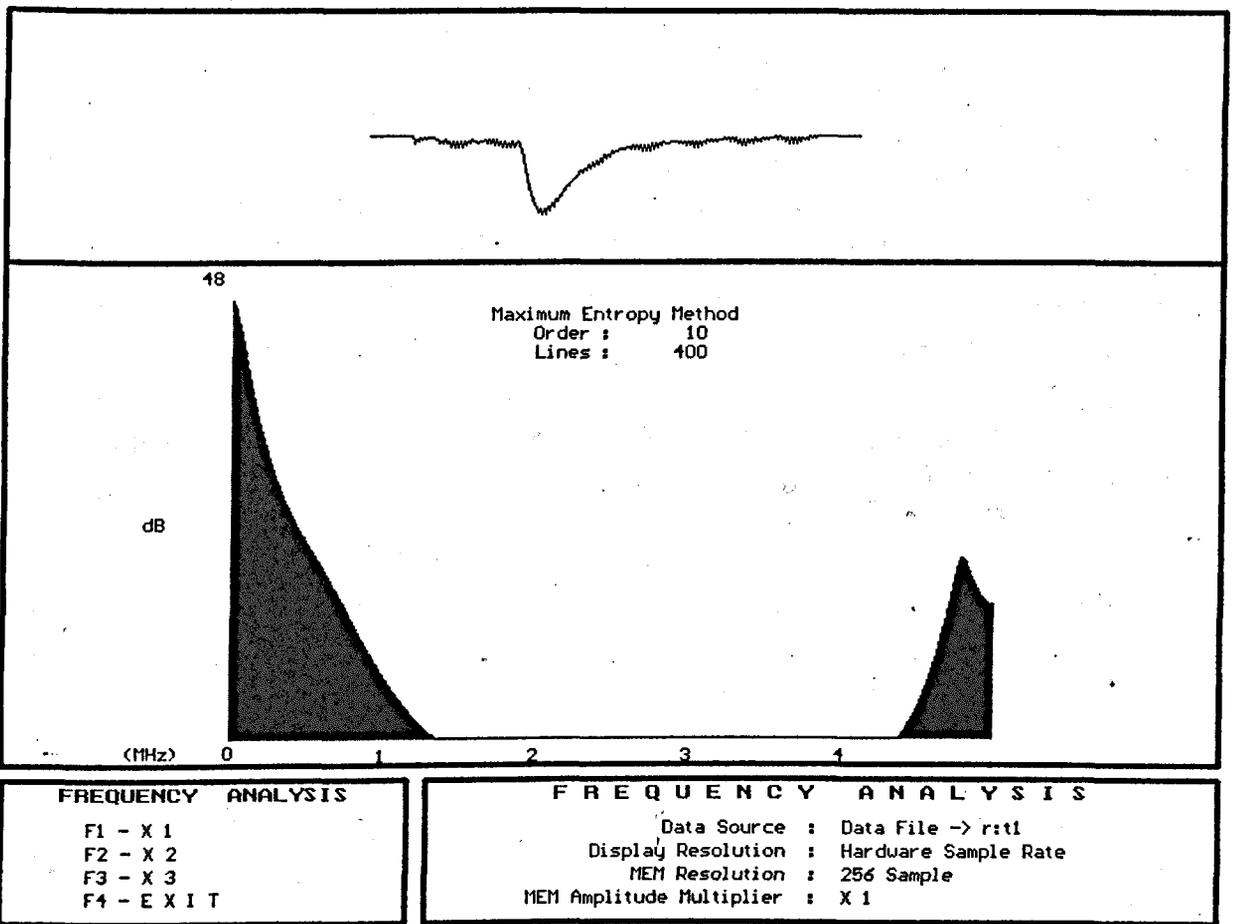


Figure 5.38 MEM Estimate of Pulse From Coil 1: 10 Poles

Here, the result of reducing the number of poles too much can be observed: the power spectrum is very smooth with little detail.

Having examined the different spectral estimates, it was decided to use a MEM order of 20 to evaluate the power spectrums of the experimental data. For these evaluations, the spectral line resolution was expanded to produce 400 lines, from 0-2.5MHz, each spaced 6.25kHz apart. In practice, a small alteration was made to speed up the processing: Most of the frequencies of interest lie below 2MHz, so a saving on computational time would be achieved by plotting only 320 lines, still spaced by 6.25kHz, in the frequency range 0-2MHz

The following graphs are direct computer-screen reproductions and display the findings of the MEM analysis for artificial discharges injected at 1 coil intervals, up to 10 coils into the stator winding. These graphs compare directly to those of Figures 5.25-5.34 for FFT analysis. As before, in each case the top graph shows one pulse as a function of time, and its corresponding spectral plot. The bottom graph shows the mean frequency profile of 50 such pulses, along with the standard deviation for each averaged spectral line.

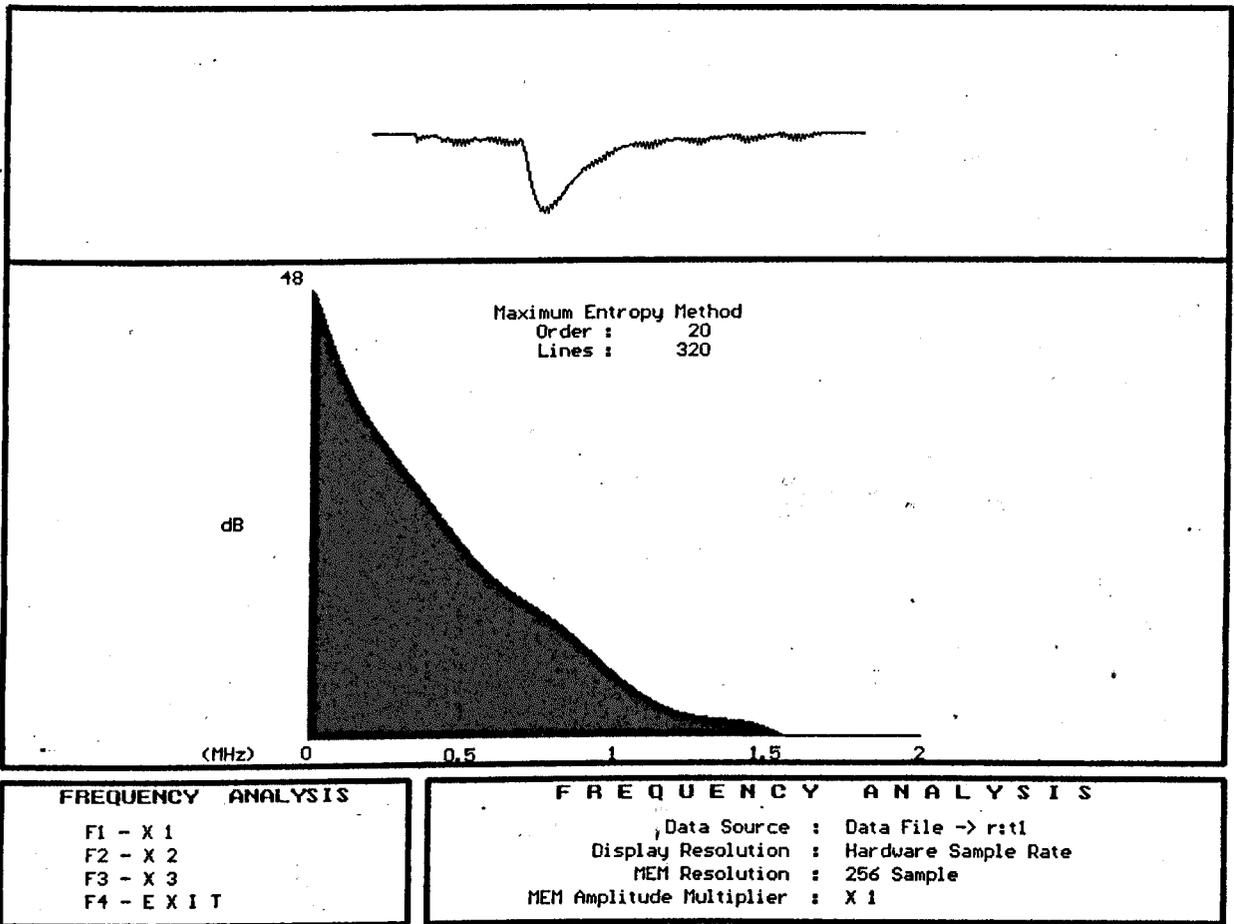
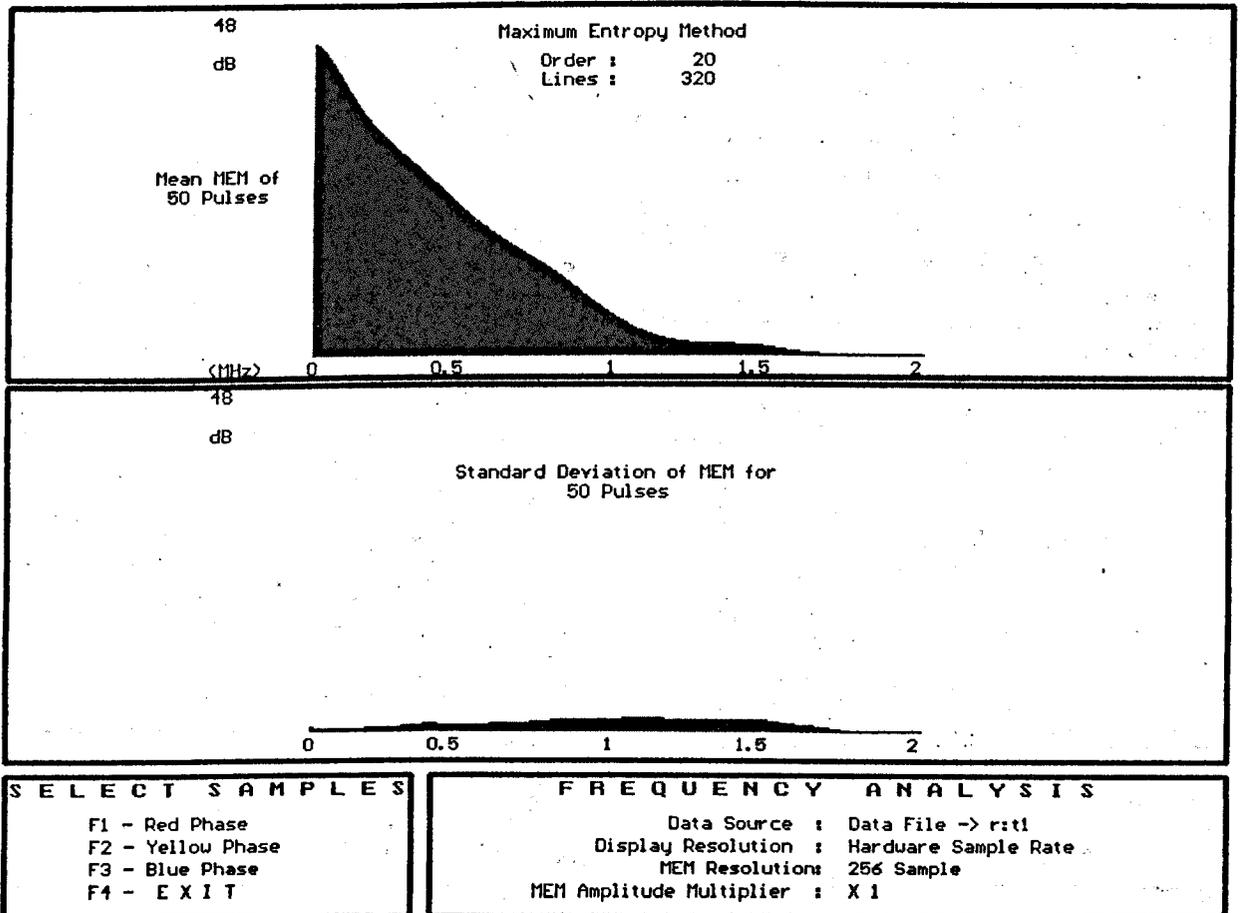
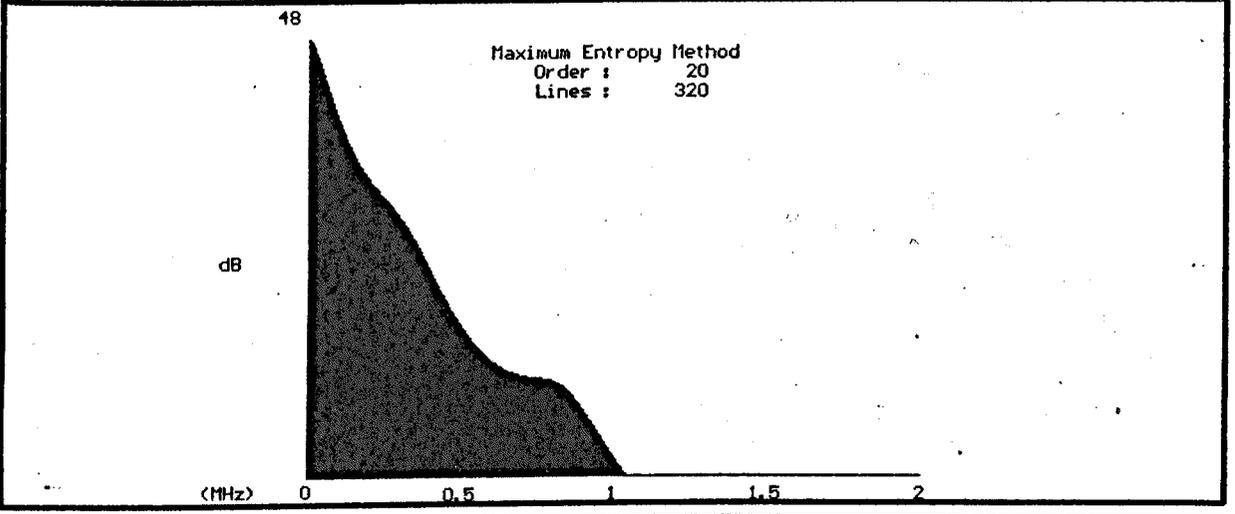
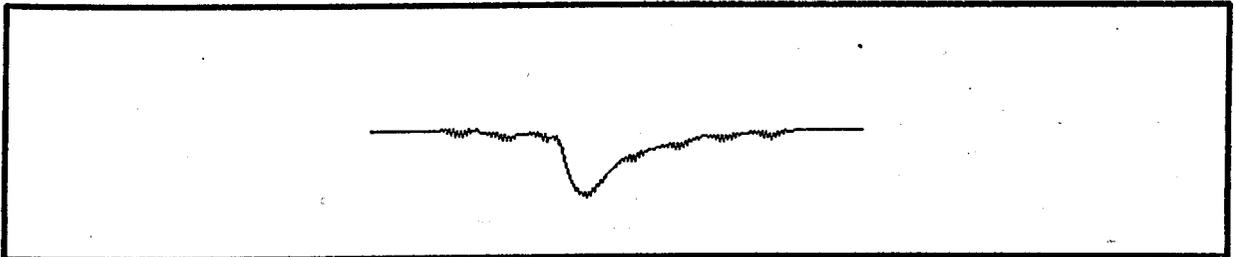


Figure 5.39 MEM Results for 1 Coil Inside Stator





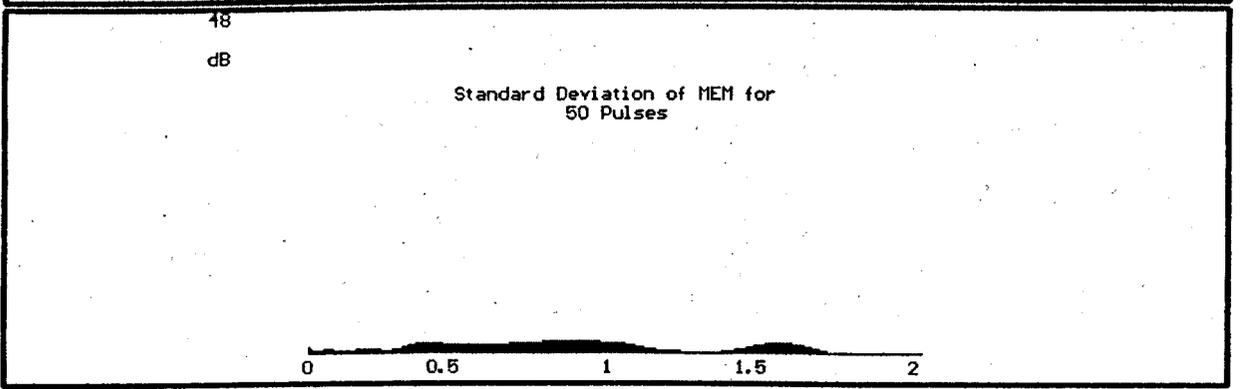
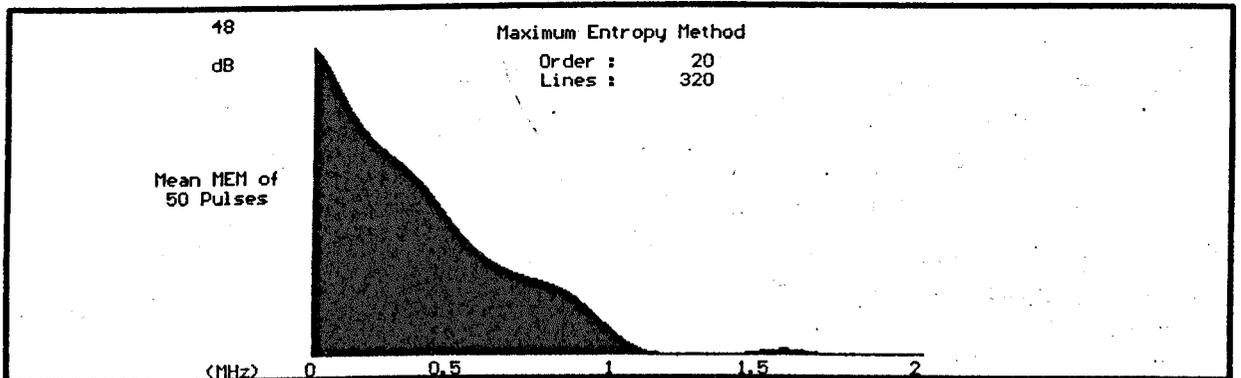
FREQUENCY ANALYSIS

F1 - X 1
F2 - X 2
F3 - X 3
F4 - E X I T

FREQUENCY ANALYSIS

Data Source : Data File -> r:tl
Display Resolution : Hardware Sample Rate
MEM Resolution : 256 Sample
MEM Amplitude Multiplier : X 1

Figure 5.40 MEM Results for 2 Coils Inside Stator

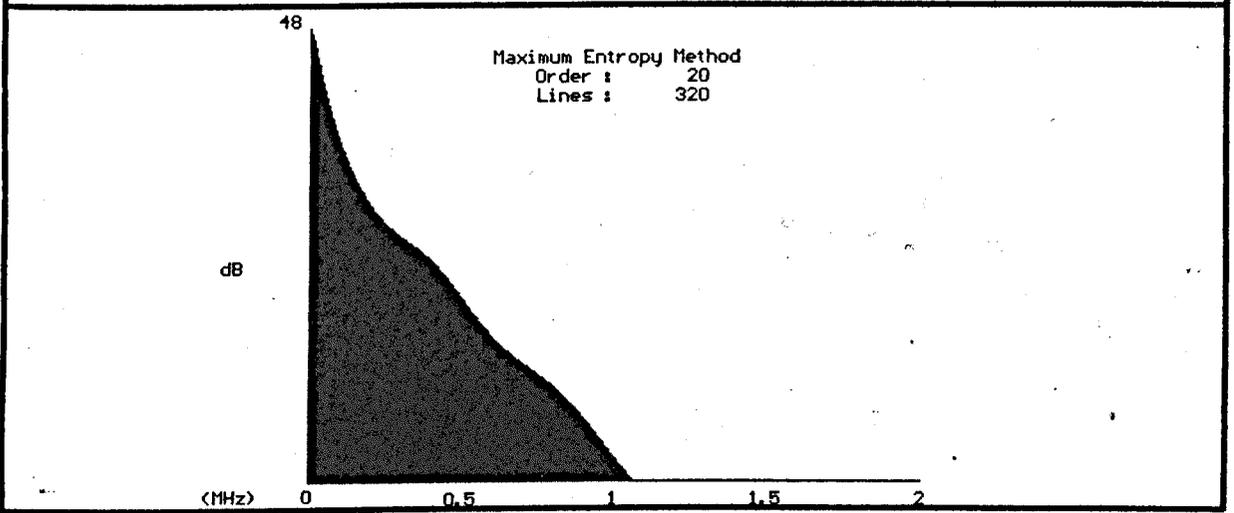
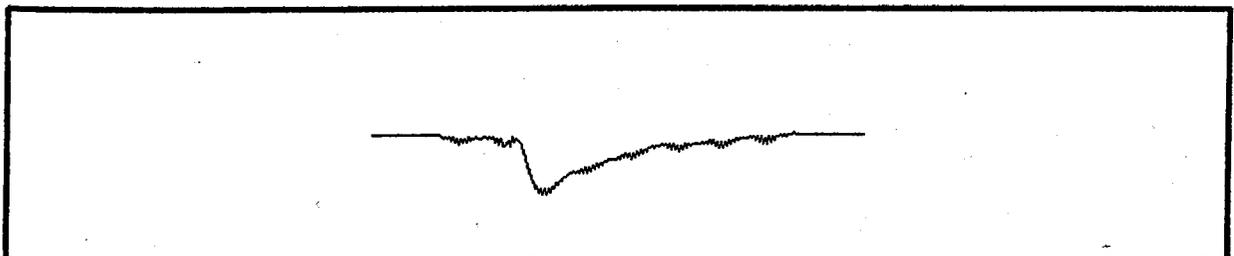


SELECT SAMPLES

F1 - Red Phase
F2 - Yellow Phase
F3 - Blue Phase
F4 - E X I T

FREQUENCY ANALYSIS

Data Source : Data File -> r:tl
Display Resolution : Hardware Sample Rate
MEM Resolutions : 256 Sample
MEM Amplitude Multiplier : X 1



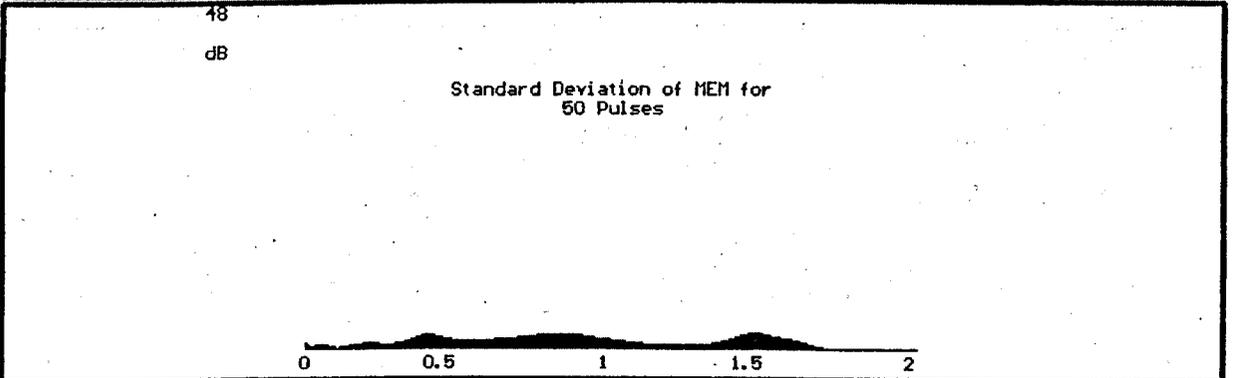
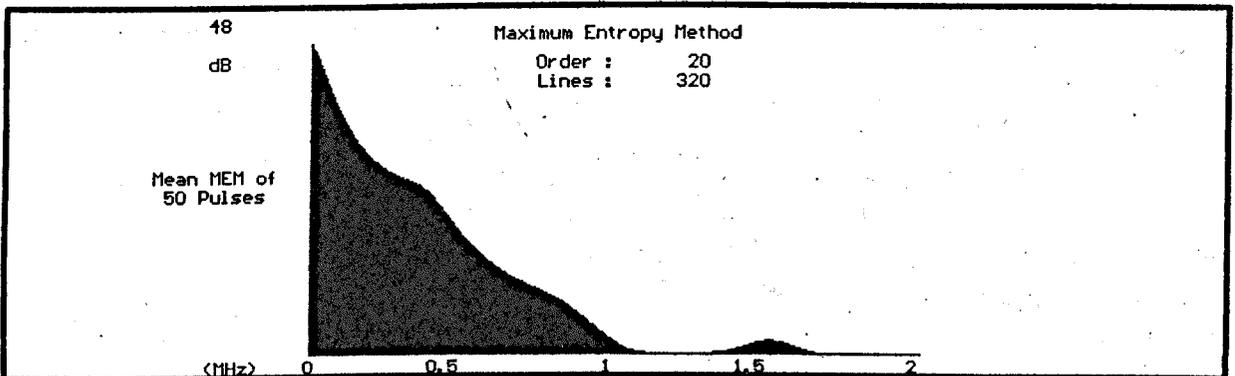
FREQUENCY ANALYSIS

F1 - X 1
F2 - X 2
F3 - X 3
F4 - E X I T

FREQUENCY ANALYSIS

Data Source : Data File -> r:tl
Display Resolution : Hardware Sample Rate
MEM Resolution : 256 Sample
MEM Amplitude Multiplier : X 1

Figure 5.41 MEM Results for 3 Coils Inside Stator

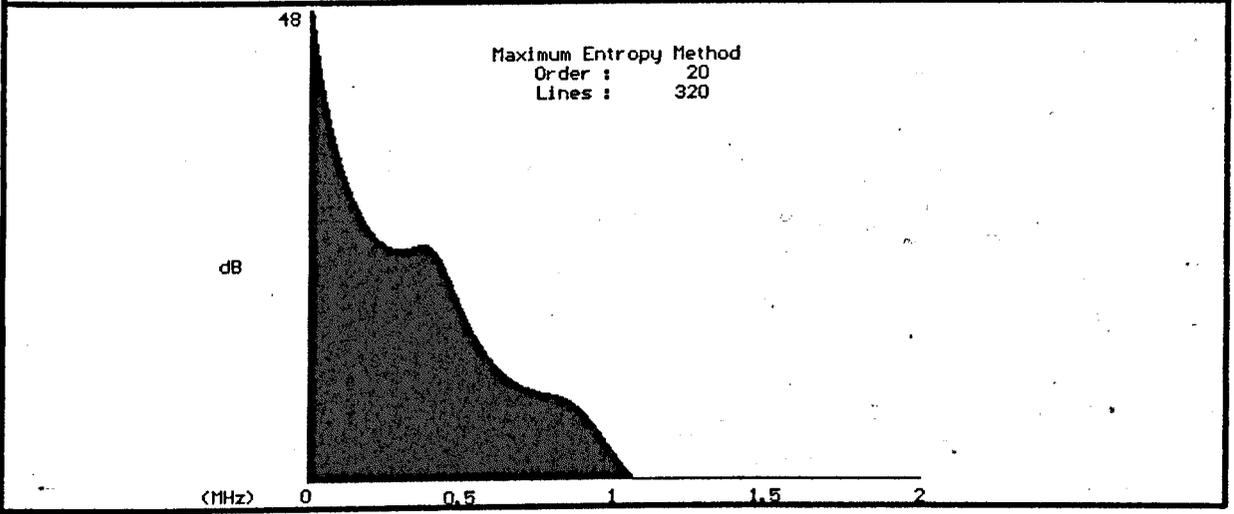
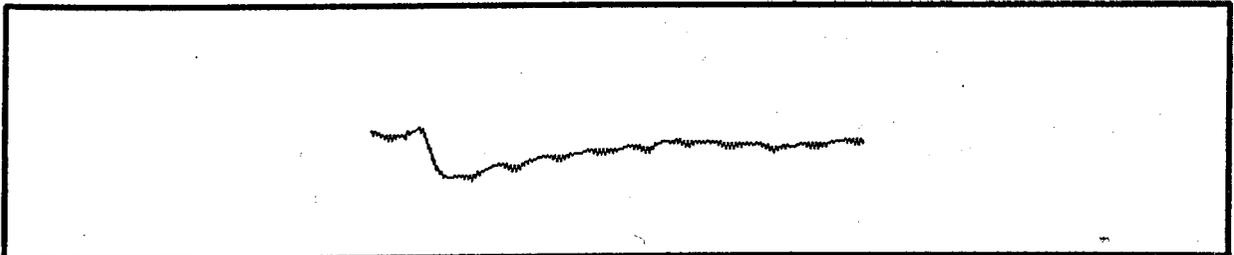


SELECT SAMPLES

F1 - Red Phase
F2 - Yellow Phase
F3 - Blue Phase
F4 - E X I T

FREQUENCY ANALYSIS

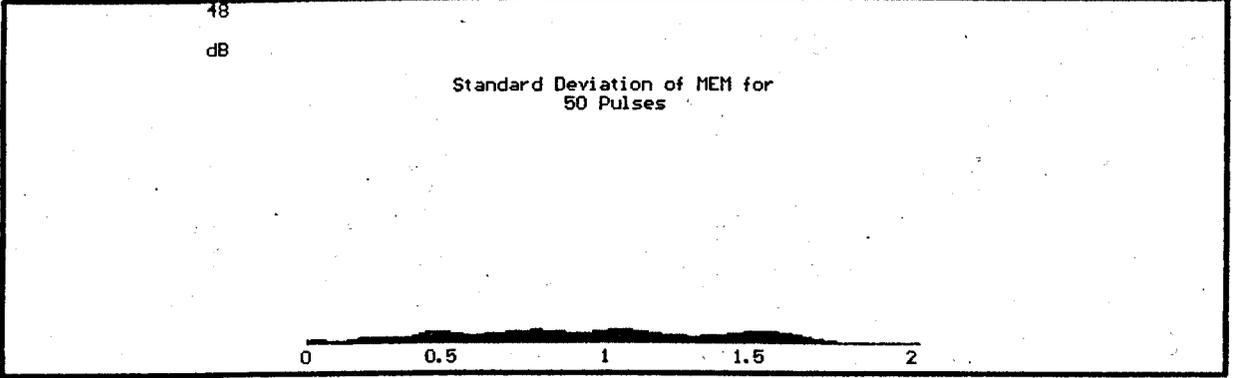
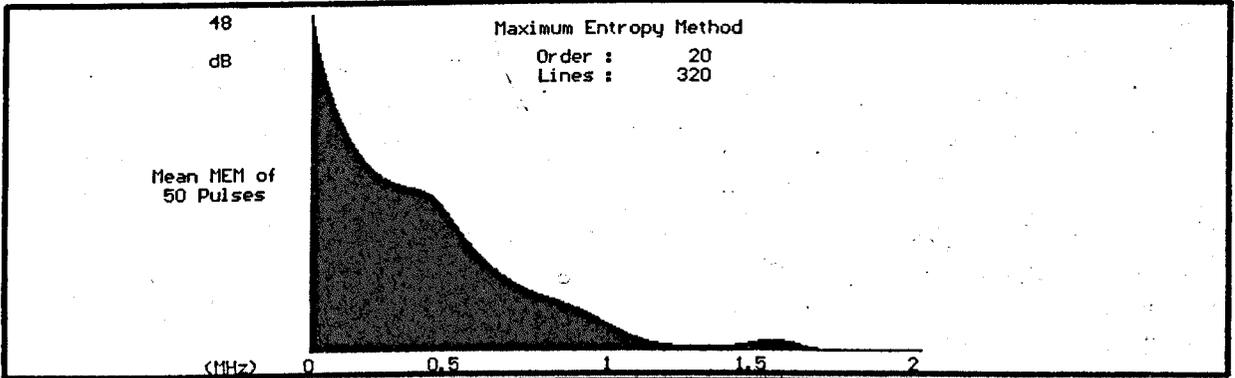
Data Source : Data File -> r:tl
Display Resolution : Hardware Sample Rate
MEM Resolution : 256 Sample
MEM Amplitude Multiplier : X 1



FREQUENCY ANALYSIS
 F1 - X 1
 F2 - X 2
 F3 - X 3
 F4 - E X I T

FREQUENCY ANALYSIS
 Data Source : Data File -> r:t2
 Display Resolution : Hardware Sample Rate
 MEM Resolution : 256 Sample
 MEM Amplitude Multiplier : X 1

Figure 5.42 MEM Results for 4 Coils Inside Stator



SELECT SAMPLES
 F1 - Red Phase
 F2 - Yellow Phase
 F3 - Blue Phase
 F4 - E X I T

FREQUENCY ANALYSIS
 Data Source : Data File -> r:t1
 Display Resolution : Hardware Sample Rate
 MEM Resolution : 256 Sample
 MEM Amplitude Multiplier : X 1

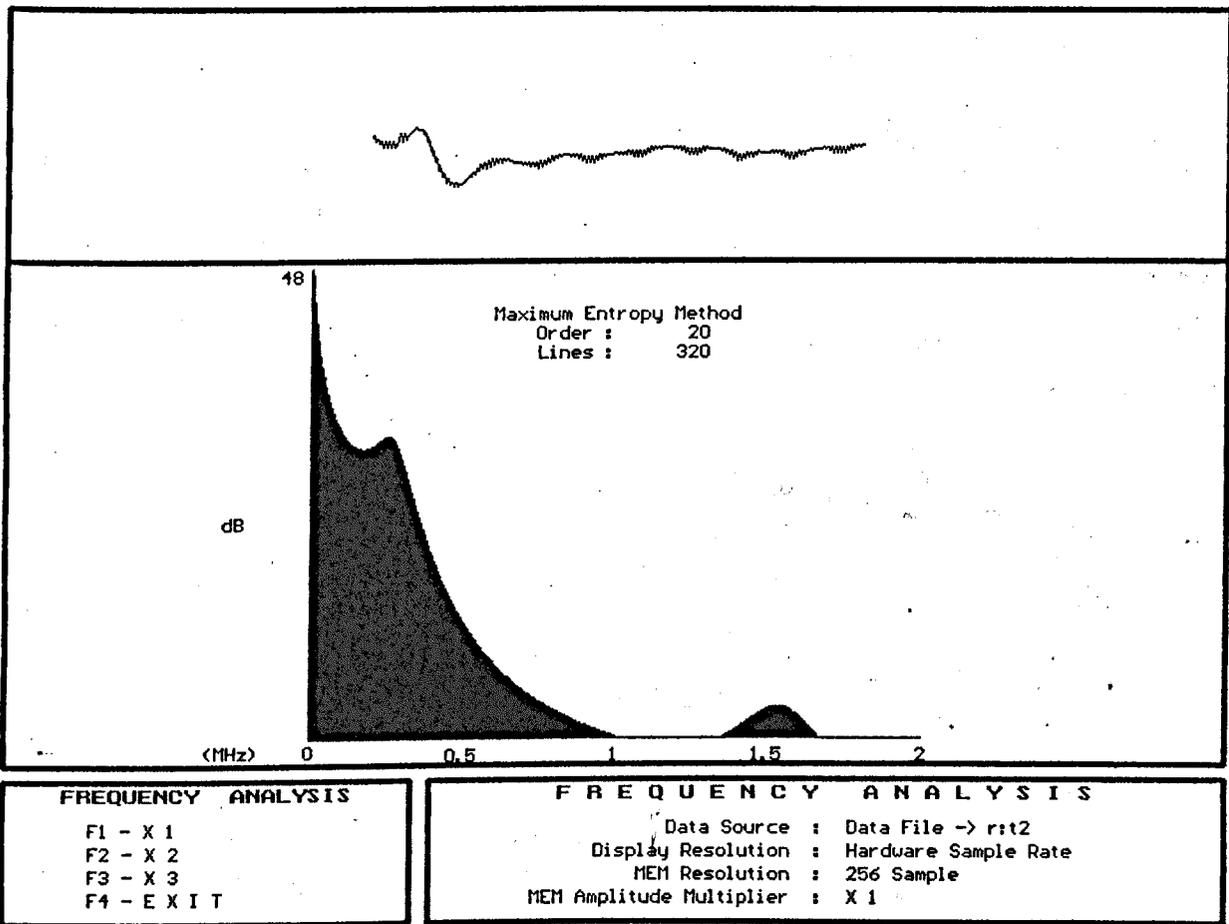
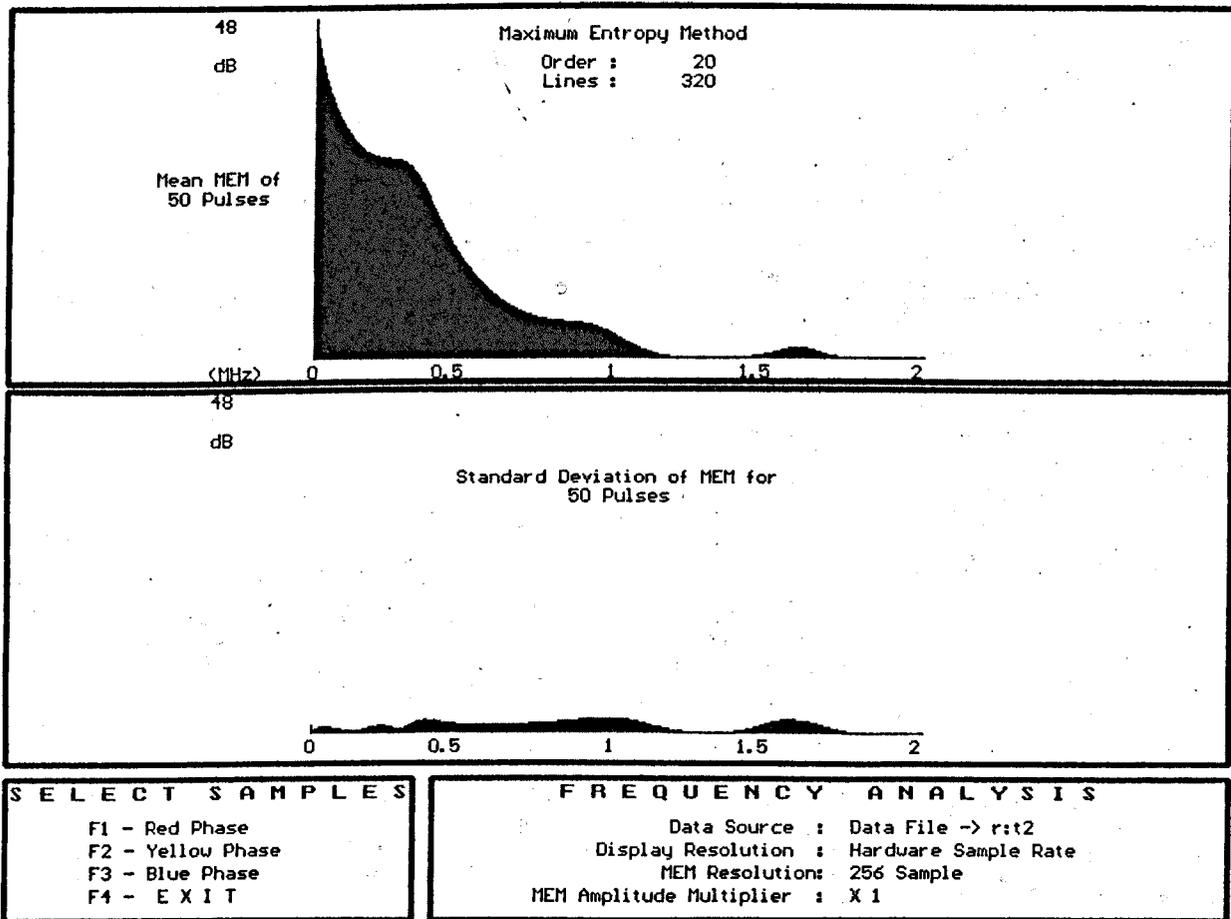
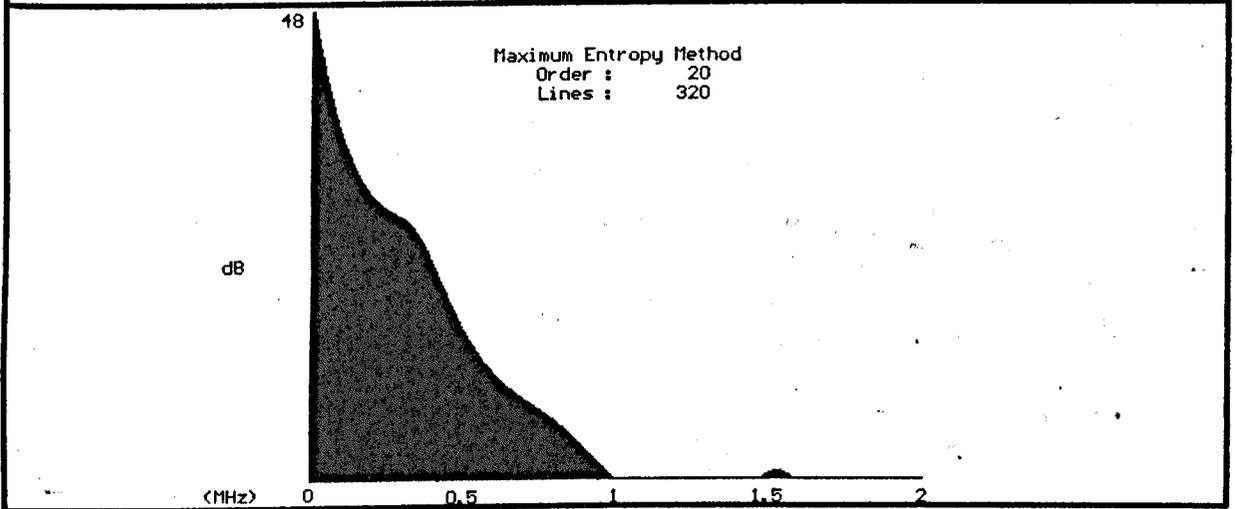
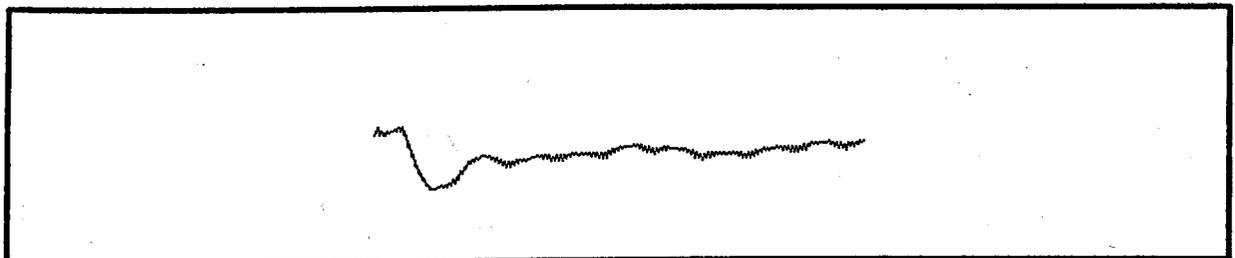


Figure 5.43 MEM Results for 5 Coils Inside Stator





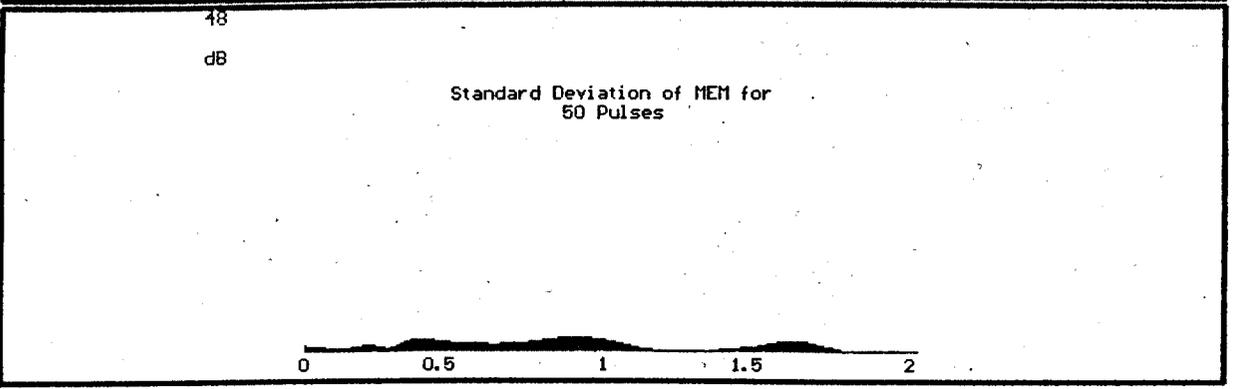
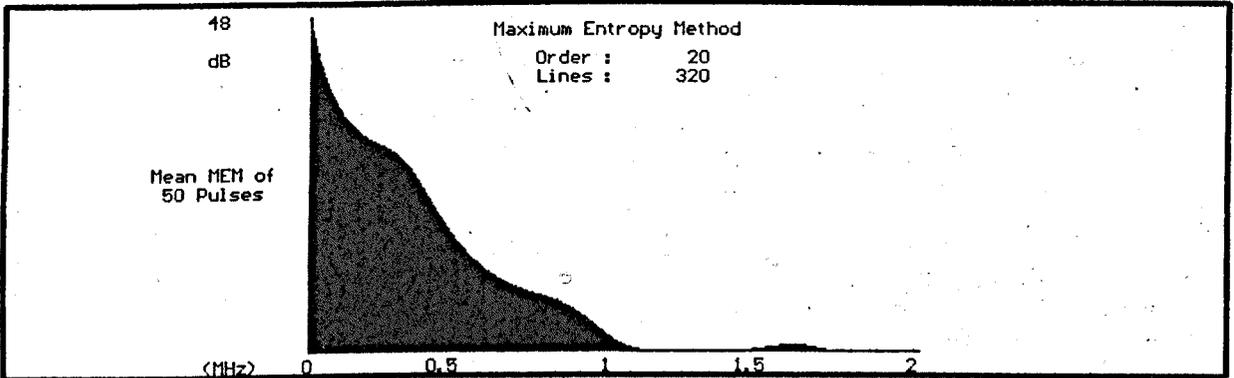
FREQUENCY ANALYSIS

F1 - X 1
F2 - X 2
F3 - X 3
F4 - E X I T

FREQUENCY ANALYSIS

Data Source : Data File -> rrt2
Display Resolution : Hardware Sample Rate
MEM Resolution : 256 Sample
MEM Amplitude Multiplier : X 1

Figure 5.44 MEM Results for 6 Coils Inside Stator



SELECT SAMPLES

F1 - Red Phase
F2 - Yellow Phase
F3 - Blue Phase
F4 - E X I T

FREQUENCY ANALYSIS

Data Source : Data File -> rrt2
Display Resolution : Hardware Sample Rate
MEM Resolution : 256 Sample
MEM Amplitude Multiplier : X 1

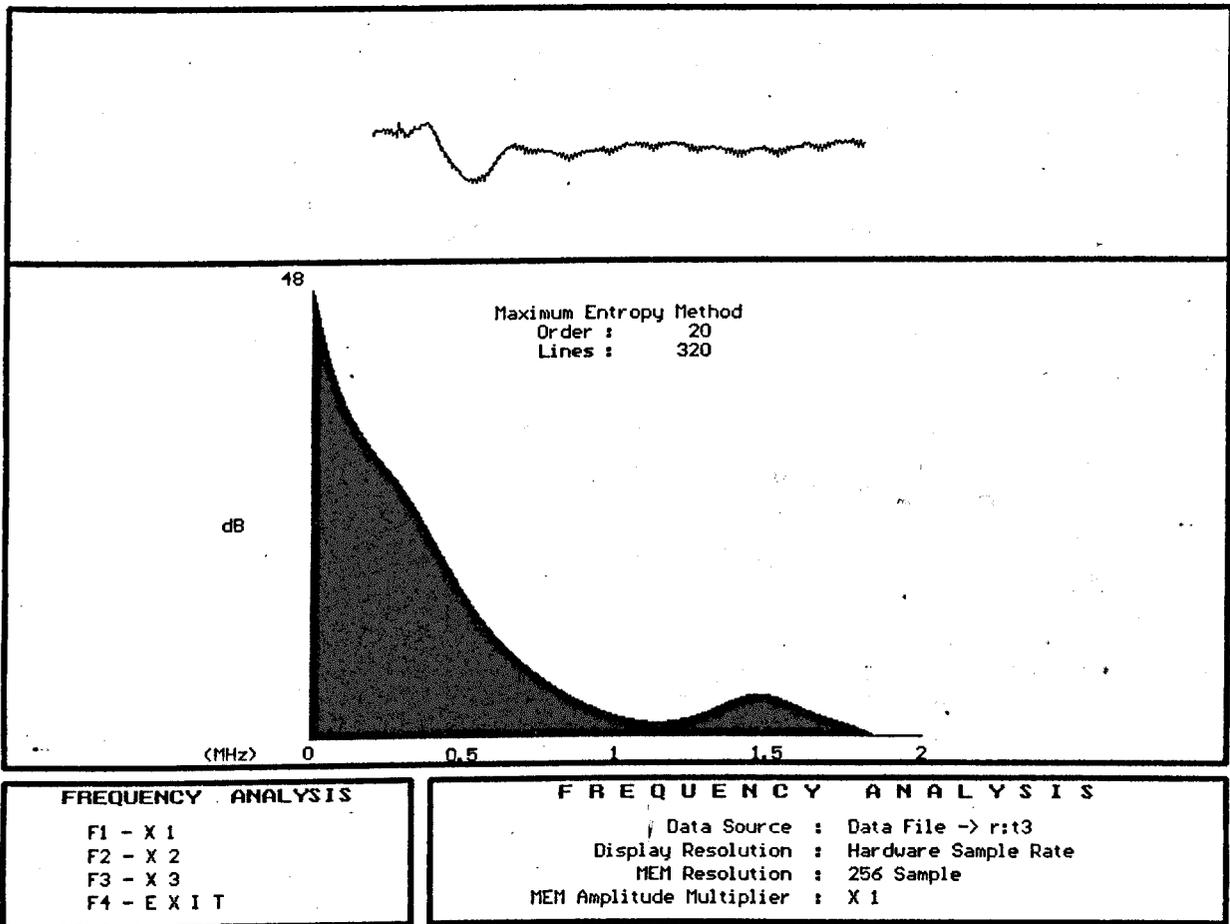
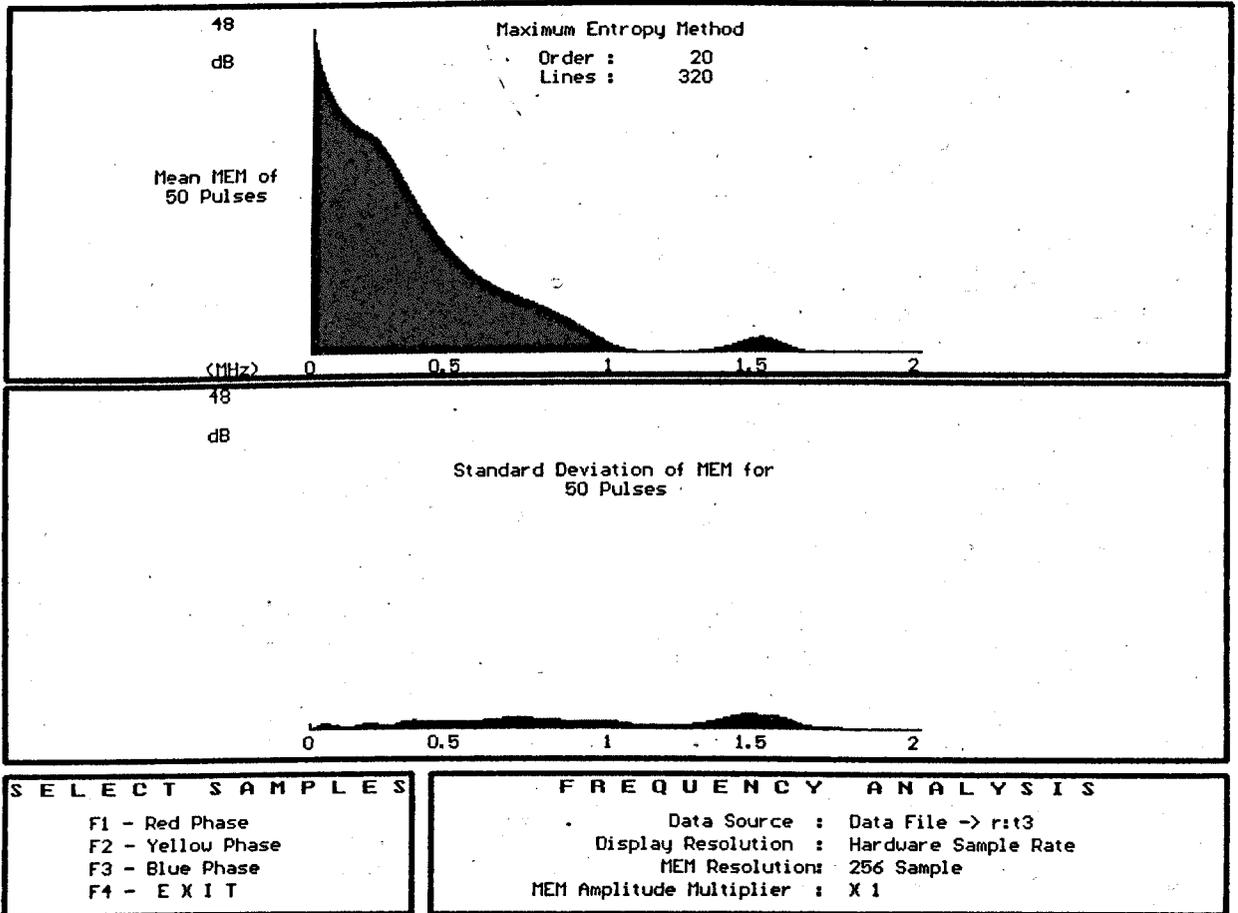


Figure 5.45 MEM Results for 7 Coils Inside Stator



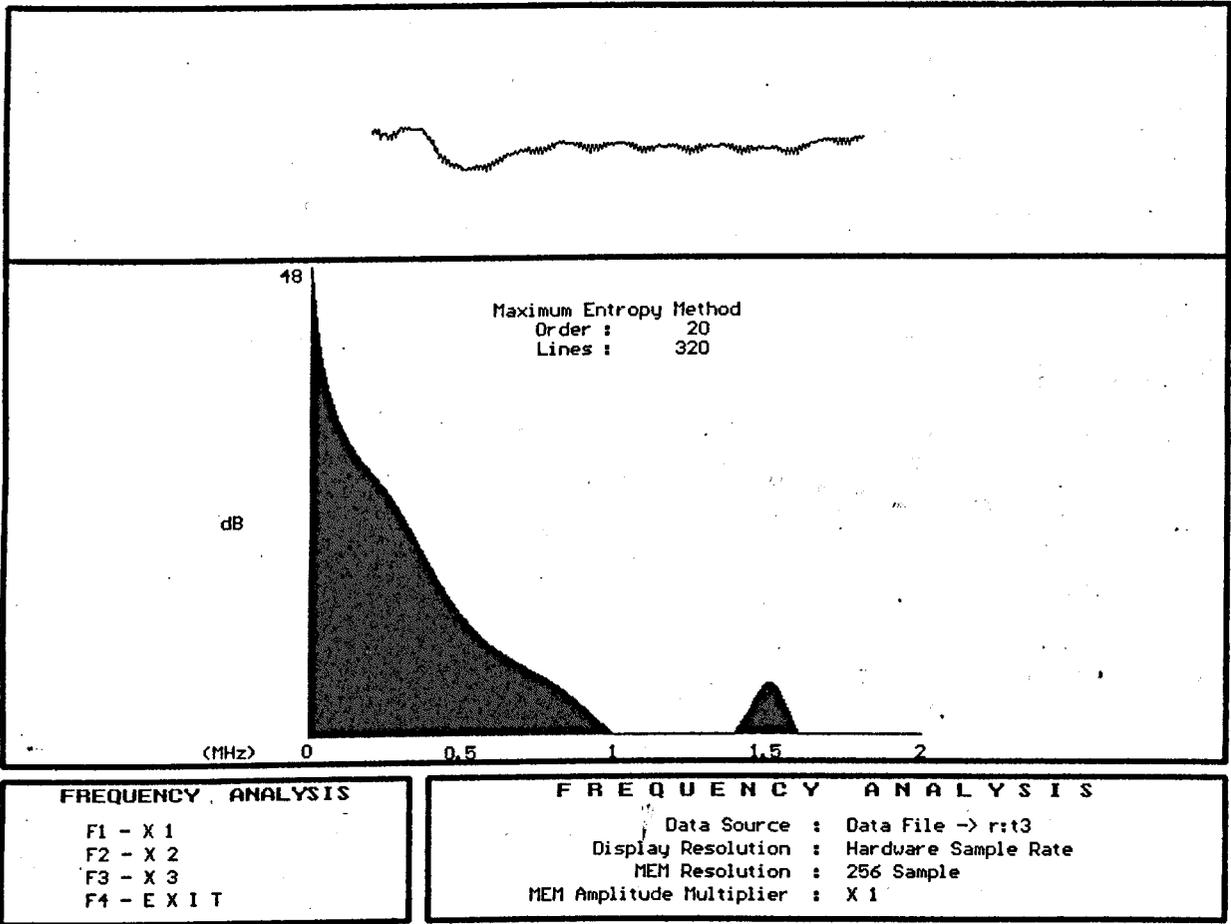
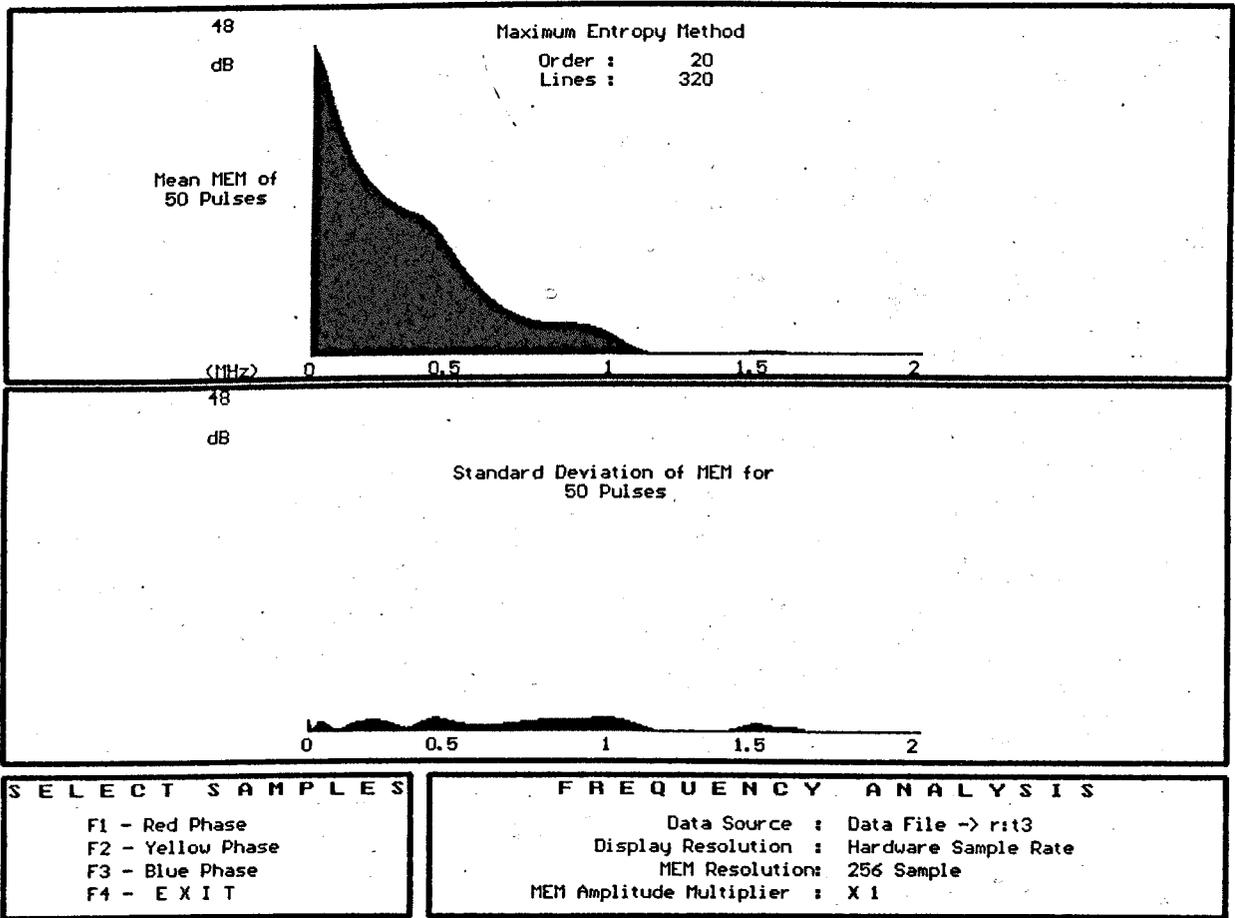


Figure 5.46 MEM Results for 8 Coils Inside Stator



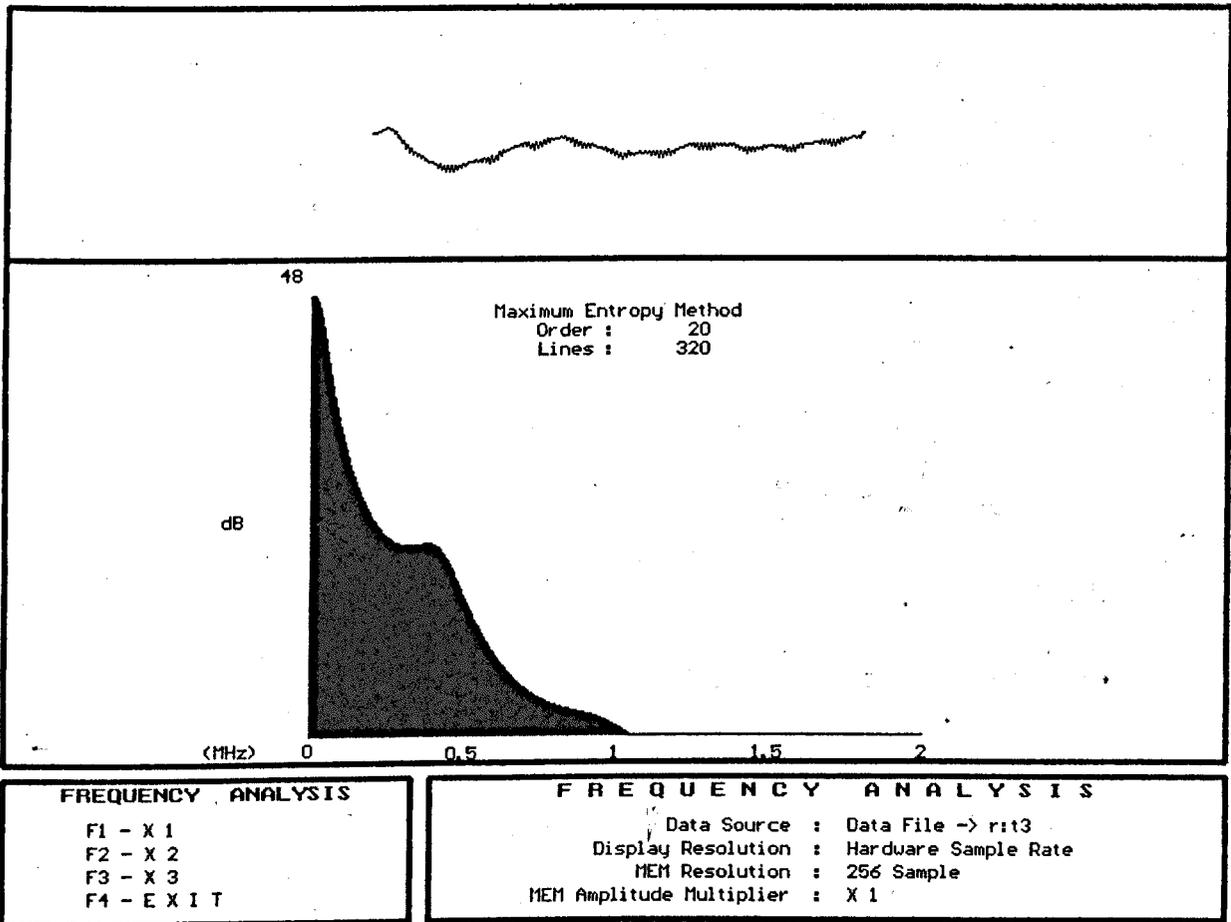
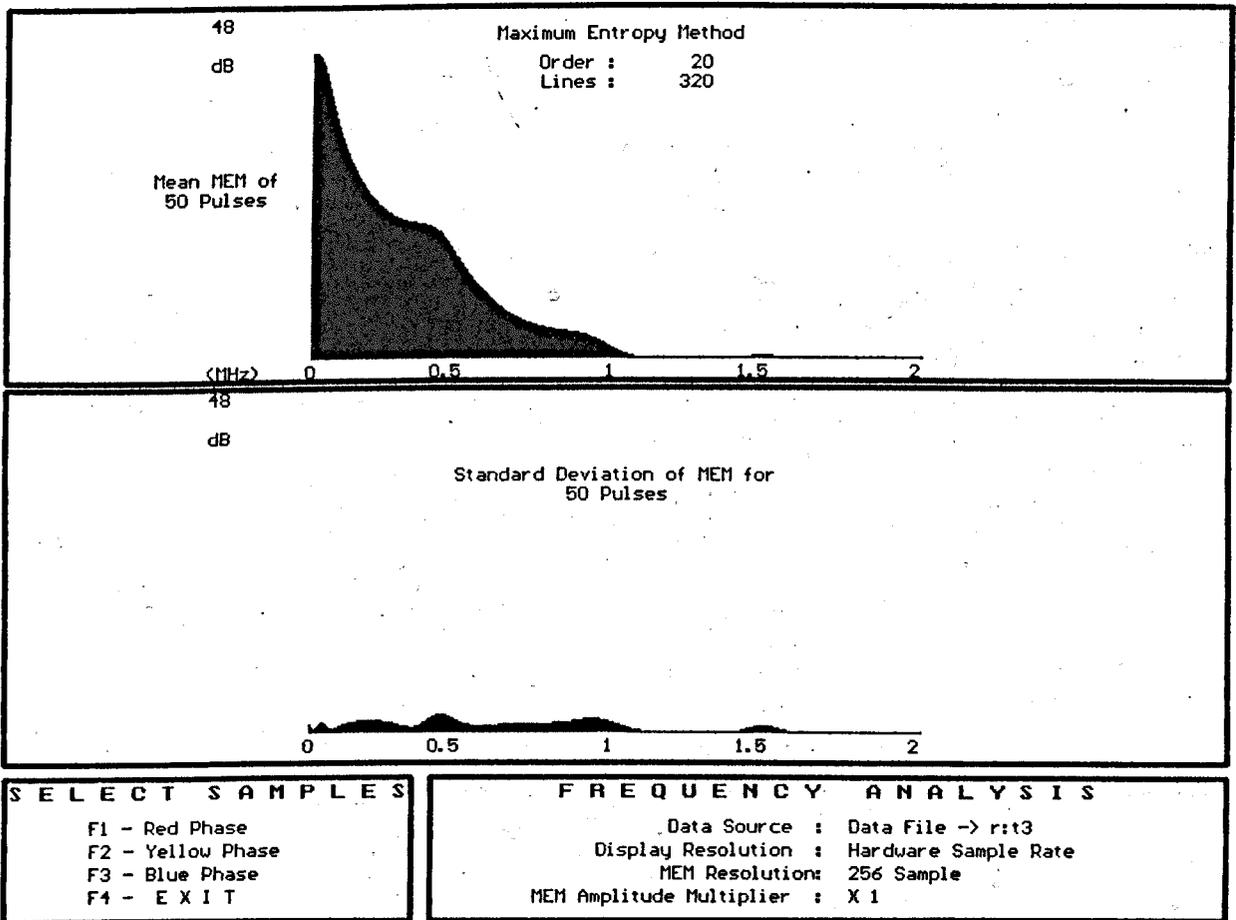
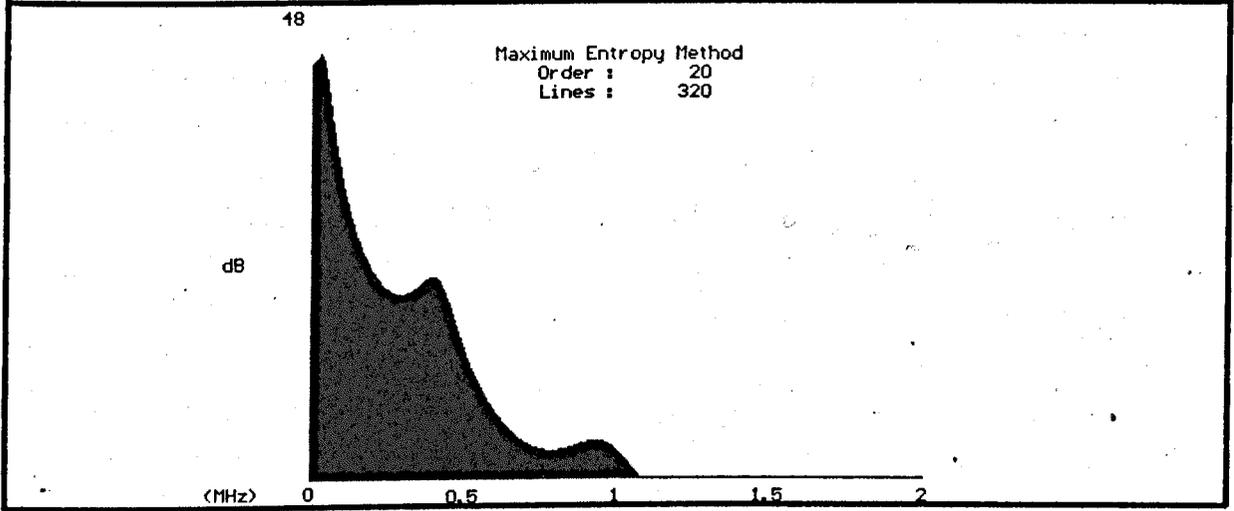
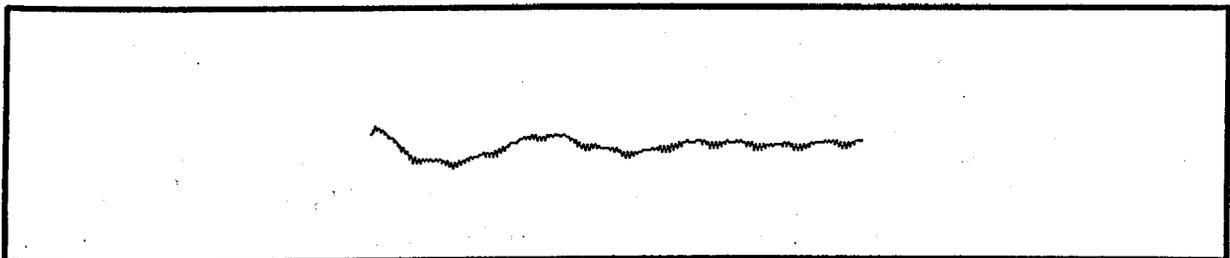


Figure 5.47 MEM Results for 9 Coils Inside Stator





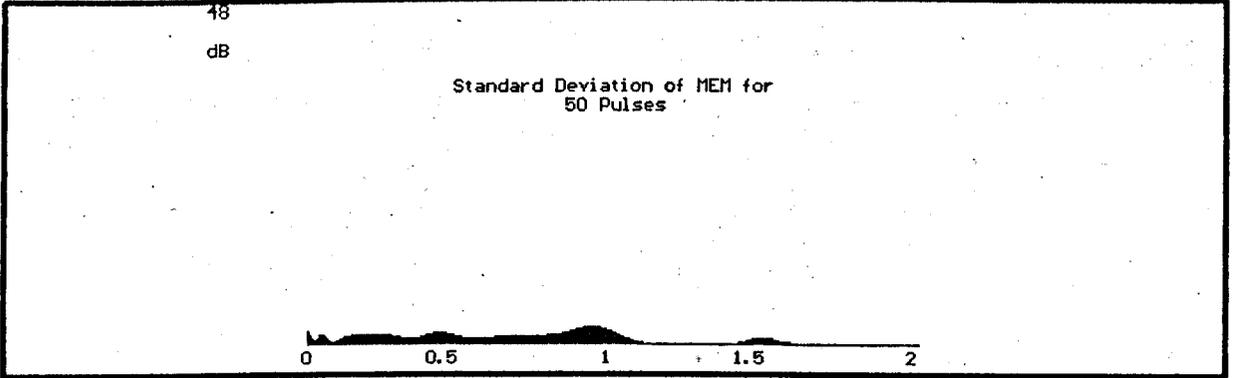
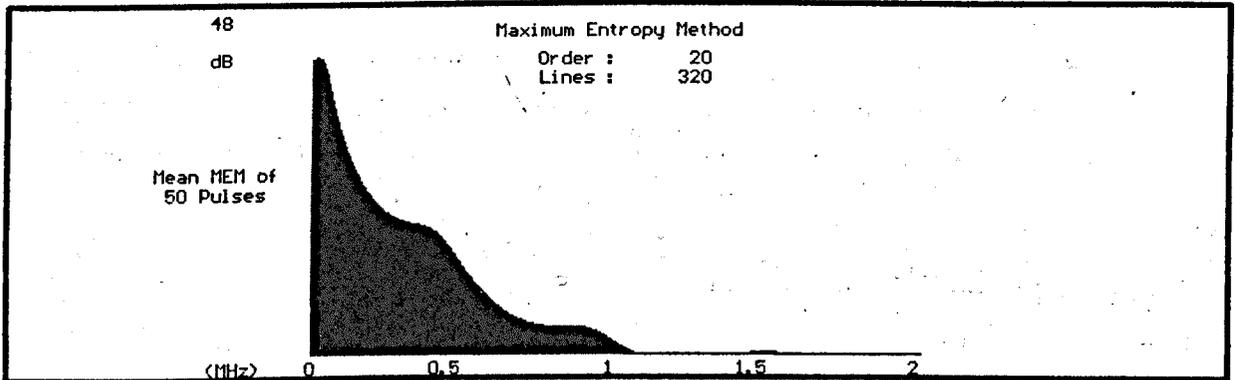
FREQUENCY ANALYSIS

F1 - X 1
 F2 - X 2
 F3 - X 3
 F4 - E X I T

FREQUENCY ANALYSIS

Data Source : Data File -> r:test10
 Display Resolution : Hardware Sample Rate
 MEM Resolution : 256 Sample
 MEM Amplitude Multiplier : X 1

Figure 5.48 MEM Results for 10 Coils Inside Stator



SELECT SAMPLES

F1 - Red Phase
 F2 - Yellow Phase
 F3 - Blue Phase
 F4 - E X I T

FREQUENCY ANALYSIS

Data Source : Data File -> r:test10
 Display Resolution : Hardware Sample Rate
 MEM Resolution : 256 Sample
 MEM Amplitude Multiplier : X 1

The MEMs were evaluated over the frequency range 0 - 2MHz with a spectral line spacing of 6.25kHz. This means that the high frequency noise components are not visible at all on these plots. However, this does not mean that they were no longer there, but merely disregarded for the purposes of this investigation, for the same reasons as stated earlier.

Examining Figure 5.39, the MEM results for discharges injected one coil distant from the line-end of the stator, it can be seen that the plot for the single pulse does not differ significantly from that of the mean, bearing in mind the differences in scale. It is not until the standard deviation plot is examined that it becomes apparent that there must be significant differences from pulse to pulse. From figure 5.39, it appears that the MEM produces spectral components up to 1.25MHz, before the amplitude of the standard deviation makes further measurement assumptions meaningless.

Figures 5.40 and 5.41 and 5.42 show little differences from each other, all with significant spectral components up to around 1MHz, but when compared with the plot for one coil distant (Figure 5.39), they are seen to have attenuation across the entire spectral range.

Figure 5.43, the plot for discharges injected 5 coils distant from the line-end shows a spectral peak at around 400kHz which is not as prominent in the previous plots. Comparing this with the FFT analysis (Figure 5.29), a 400kHz component can also be seen, indicating that the measured peak here is unlikely to be directly attributable to a spurious by-product of the MEM analysis.

Figure 5.44 shows a less prominent peak at 400kHz, but otherwise is very similar to that of the previous plot.

Figures 5.45, 5.46, 5.47 and 5.48 are very similar in appearance, showing gradual spectral attenuation as pulses have to travel further through the winding.

Both spectral analyses, Fast Fourier Transform and the Maximum Entropy Method, can only operate on one pulse at a time. Therefore, in order to get an average frequency profile of the injected pulses, additional software routines had to be developed that could automatically identify these pulses, pass them to the analysis routines, one by one, and finally find the mean frequency profile, as well as the standard deviation of each spectral line.

The standard deviation of each spectral line is a direct measure of the stability, or certainty, of that particular line being at a given level after the analysis of any pulse originating from the same point: i.e. if the mean level of the 100kHz component of a pulse from coil 2 is 40dB, given that 50 pulses were analysed, and the standard deviation of the 100kHz component is 2dB, then we can say with a large degree of certainty that any pulse originating from coil 2 will have a 100kHz component of 40 ± 2 dB. This means that the standard deviation of the spectral lines will play an important part in identifying which, if any, lines can be used to estimate where the analysed pulses originated from.

The mean and standard deviations of each spectral plot were stored in text files for later, more detailed, analysis.

5.6 Summary

The three-phase monitoring software, whilst retaining the "look and feel" of the prototype system, was considerably enhanced to incorporate the three-phase display and pulse analysis routines. Additionally, a statistical noise removal scheme was developed which can automatically identify and remove the unwanted thyristor switching pulses that can interfere with correct diagnoses of partial discharge activity.

Finally, two different frequency analysis schemes were implemented to investigate the frequency content of partial discharges at different inception points within stator windings.

The next chapter deals with a more detailed analysis of the spectral data obtained earlier.

CHAPTER 6

ANALYSIS OF SPECTRAL DATA

6.1 Introduction

Visual inspection, alone, cannot realistically discriminate between pulses which originate one coil apart. However, two display options were investigated to try and highlight the differences that do exist between the frequency profiles.

6.1.1 Three Dimensional Display Options

The first display option that was investigated involves directly displaying all 10 mean coil profiles, isometrically, beside one another. In addition to this, the degree of uncertainty of each spectral line is also shown by shading each spectral line to \pm the standard deviation of each line. The result of this can be seen in Figure 6.1.

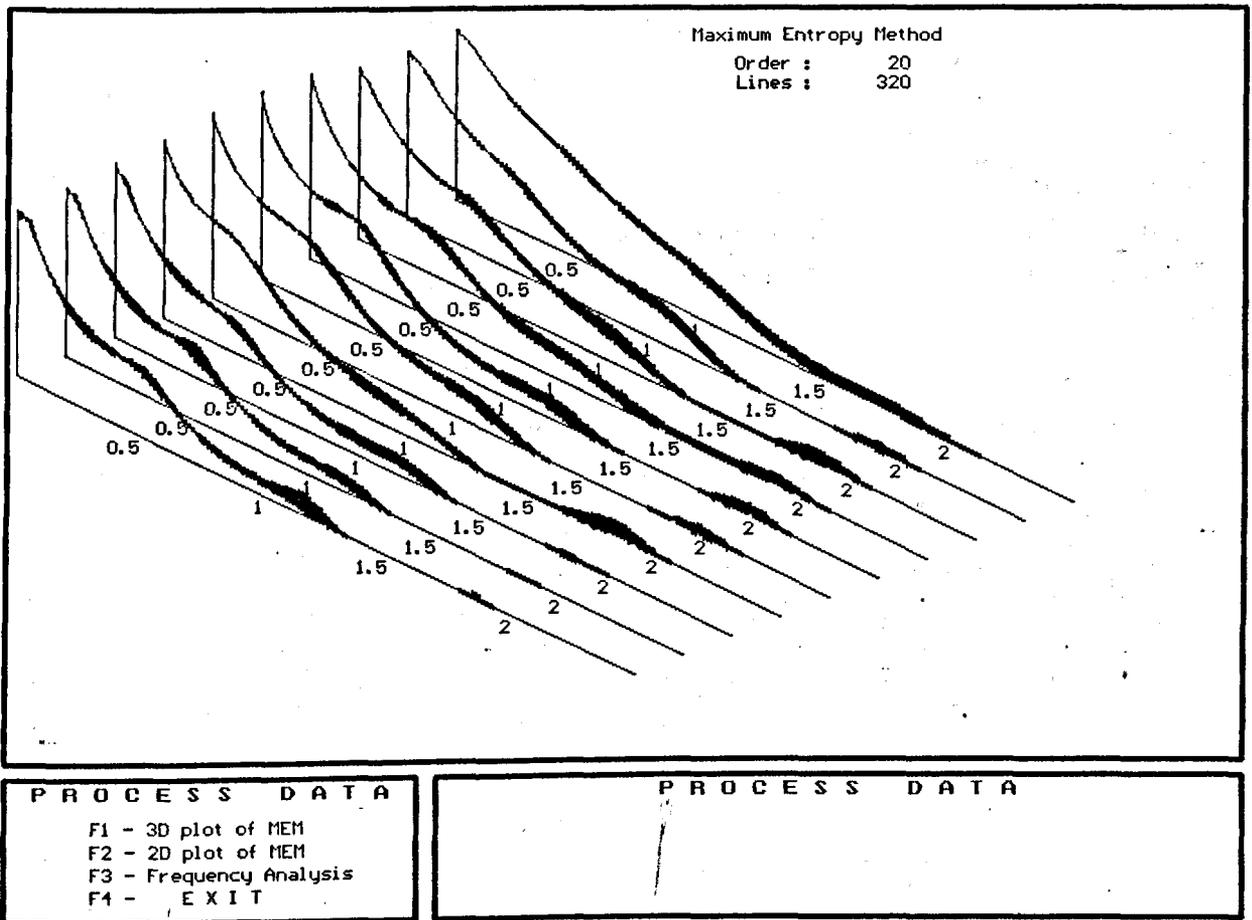


Figure 6.1 Comparison of Frequency profiles - 1

The above plot shows the mean frequency profiles, using the MEM order 20, for artificial partial discharge pulses injected at ten different distances inside a large machine stator winding. Each profile shows the mean of 50 analyses, with the standard deviation of each spectral line superimposed as a shaded area.

The profile appearing at the rear of the plot is for pulses injected 1 coil from the line-end of the stator, whilst the profile at the front is for pulses injected 10 coils from the line-end.

It can be seen from the diagram that pulses that have been injected near the line-end of the stator have a higher overall frequency content than those that have had to travel further through the winding before being detected. This adheres to the general

frequency-dependant attenuation characteristic expected from the stator winding. This information can be displayed in another way, depicting a better view of the attenuation of certain frequencies as the pulses travel further through the winding, Figure 6.2.

Here we see a sort of 'relief map' of the frequency profiles: all spectral lines of the same frequency are joined together.

If one particular line is chosen, as we travel left to right along the graph, we can see how that particular frequency changes as discharge pulses have to travel further through the winding.

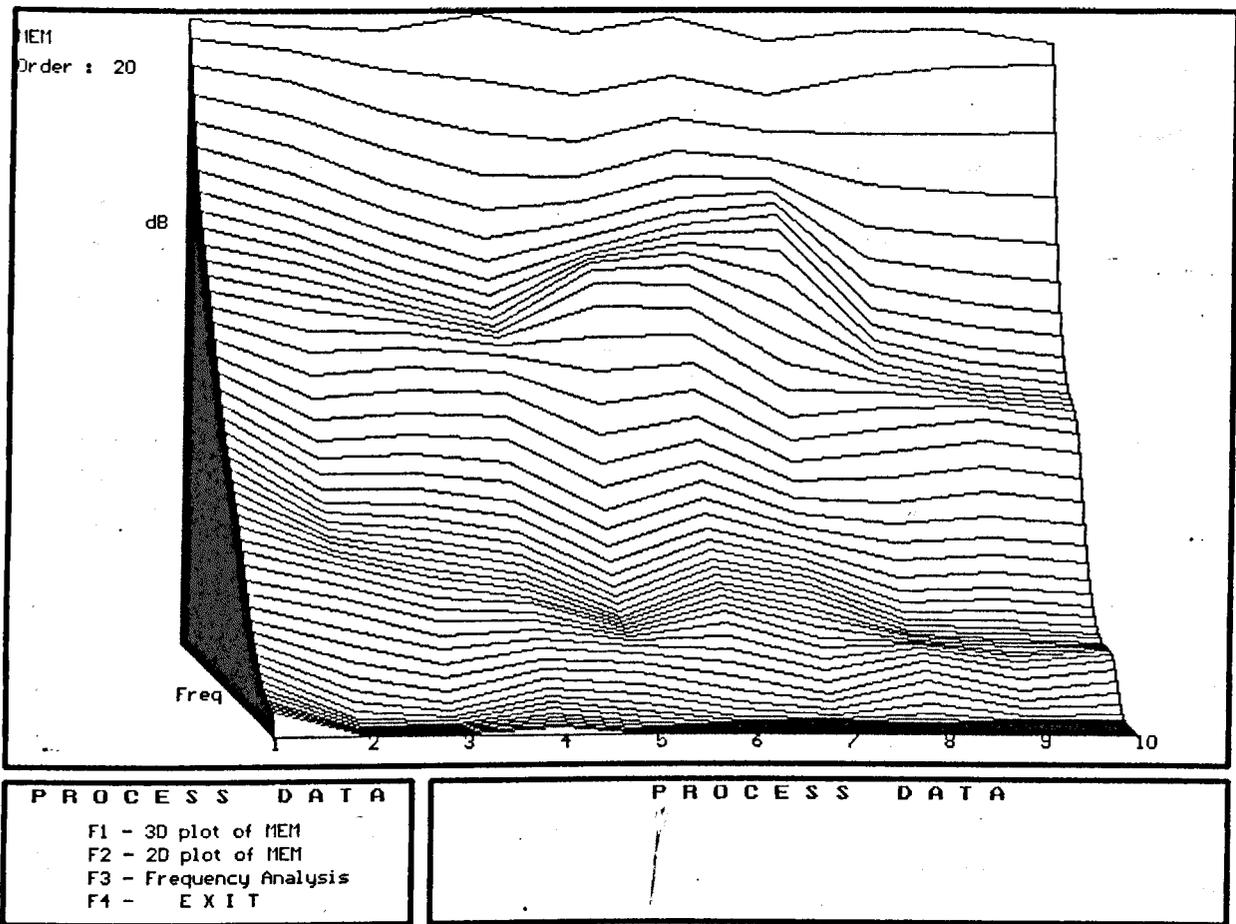


Figure 6.2 Comparison of Frequency Profiles - 2

Irrespective of how convoluted the display of frequency data is, there are no alternatives to processing the 'raw' numbers in order to accurately ascertain if partial discharge location is at all feasible.

Considering the MEM analysis, each winding location in the experiment produced 320 mean values and 320 standard deviations. Given that experiments were carried out at 10 different winding locations, 6400 items of data needed to be processed. Furthermore, this figure would increase dramatically if individual spectral estimations were used instead of the average values.

Bespoke software routines could have been developed to analyse the data, incorporating both the processing and display features. However, until the appropriate analyses were identified, this would be extremely time consuming.

Software products exist, most notably spreadsheets, that have been specifically designed to manipulate numbers. Such products could be utilised if the data was in a format that the spreadsheet could recognise. Not only this, but a spreadsheet has a number of built-in graphical display routines that could display the results without further software development.

6.2 Microsoft Excel as a Data Analysis Tool

Microsoft Excel is an integrated spreadsheet package that runs in the Microsoft Windows environment. It allows multiple spreadsheet files to be open simultaneously, and can produce graphs of many different styles from the data contained in those spreadsheets.

6.2.1 Spreadsheet Introduction

A spreadsheet is basically a large grid, initially empty, each cell of which can be uniquely addressed with a letter and number combination. Individual columns are given letter identifiers, whilst every row is assigned a number, thus single cells within the grid are addressed using the column letter and the row number, e.g. A12, BX234 etc.

Individual cells can contain numbers, characters, or formulae, and the size of the cells can be changed to accommodate display of their contents. An example of a simple spreadsheet can be seen in Table 6.1.

The spreadsheet in Table 6.1 shows an example of yearly accounts for a fictitious firm. It serves merely to illustrate the versatility and strengths of spreadsheets, and to provide a good introduction to their application in analysis of large amounts of data.

	A	B	C	D	E	F
1	Sales	1989	1990	1991	1992	
2						
3	Widgets	£54,568	£158,567	£478,578	£457,667	
4	Sprockets	£578	£1,758	£6,567	£64,746	
5	Cogs	£2,452	£26,767	£34,234	£4,676	
6	Total Sales	£57,598	£187,092	£519,379	£527,089	
7						
8	Expenses					
9	R & D	£4,766	£78,678	£80,698	£34,256	
10	Marketing	£3,214	£57,867	£168,989	£143,356	
11	Admin	£46,757	£34,654	£165,767	£175,767	
12	Total Expenses	£54,737	£171,199	£415,454	£353,379	
13						
14	Operating Income	£2,861	£15,893	£103,925	£173,710	
15	Profit Margin	4.97%	8.49%	20.01%	32.96%	
16						
17						
18						

Table 6.1 Example Spreadsheet

Individual cells within a spreadsheet can hold text (A1), numbers (A3), or formulae (B6). The cells that hold formulae, do not display the actual formulae themselves, but rather the result of their calculation. So, for example, cell B6 shows the Total Sales for 1989, but what it actually contains is a formula for calculating the total sales for 1989: i.e. $B6 = B3 + B4 + B5$. In practice, this means that whenever any data is changed, in cells B3-B5, their total is automatically re-calculated. Not only this, but other calculations using this result are automatically updated, so, for example, the Operating Income, cell B14, subtracts the Total Expenses ($B12 = B9 + B10 + B11$) from the Total Sales, and if any cell used to calculate the Total Income is changed, then every result that references this cell, directly or indirectly, is automatically re-calculated.

Formulae are not restricted to operating on individual columns: they can reference any cell within the spreadsheet, or even cells in another spreadsheet. Additionally, formulae can be copied so that their cell references remain relative to their positions: e.g. The formula for Operating Income ($B14 = B6 - B12$) can be copied with a few keystrokes into the adjoining cell (C14) such that the cell references automatically adjust to column C, i.e. $C14 = C6 - C12$. This saves repetitively re-typing formulae, some of which can be very long.

The results produced by the spreadsheet can be left as numerical values, displayed in tables, as shown in Table 6.1, or can be presented in any number of graphical formats: bar, pie, line, or scatter charts, or whatever combination suits the data.

6.2.2 Data Analysis Using Microsoft Excel

Each spectral analysis produced, in the case of the MEM, 320 mean values, and 320 standard deviation values, each corresponding to individual frequencies 6.25kHz apart in the range 0-2MHz. This meant that the frequency analysis for each experiment (10 in all) produced 640 pieces of data, and when all of this data was collected together, some 6400 numbers had to be processed.

Before any spreadsheet could be set up, the analysed data, produced from the FFT and MEM routines, had to be converted to a format that the spreadsheet could recognise. In the end it proved a simple matter to get the data into a suitable format, merely by saving the mean and standard deviation data in text files. These text files could then be directly read into the spreadsheet.

The spreadsheet was set up to hold all of the mean frequency profiles and the standard deviation results for all 10 experiments. Next, taking the experimental data for pulses injected 1 coil into the winding as a reference, the differences between the mean frequency profiles of the other coils and the reference profile were calculated: i.e. how much each frequency dropped with respect to a pulse injected at coil 1, was found for those injected at coils 2,3,4 etc. This could then be plotted in a graph to visually display the difference. Figures 6.3 to 6.11 show the results of the calculations for each coil, using the data produced from the Maximum Entropy Method frequency analysis routines.

The graphs show the original mean frequency profiles of pulses from coils 2 through 10, and the mean frequency profile of pulses injected at coil 1, as well as the difference between them and the standard deviation result for the deeper coil.

The standard deviation graph is important if we consider that to identify a pulse as coming from, say, coil 3 then, the difference in frequency profile between 1 & 3 must

exceed the standard deviation of coil 3, as, statistically speaking, the average profile of a pulse from coil 3 can rise or fall by up to one standard deviation.

Each graph has two vertical axes, one (the solid line) contains the scaling information for the frequency profiles, whilst the other (dotted line) scales the mean differences and standard deviation plots.

Figure 6.3 MEM Analysis for Coil 2 vs Coil 1

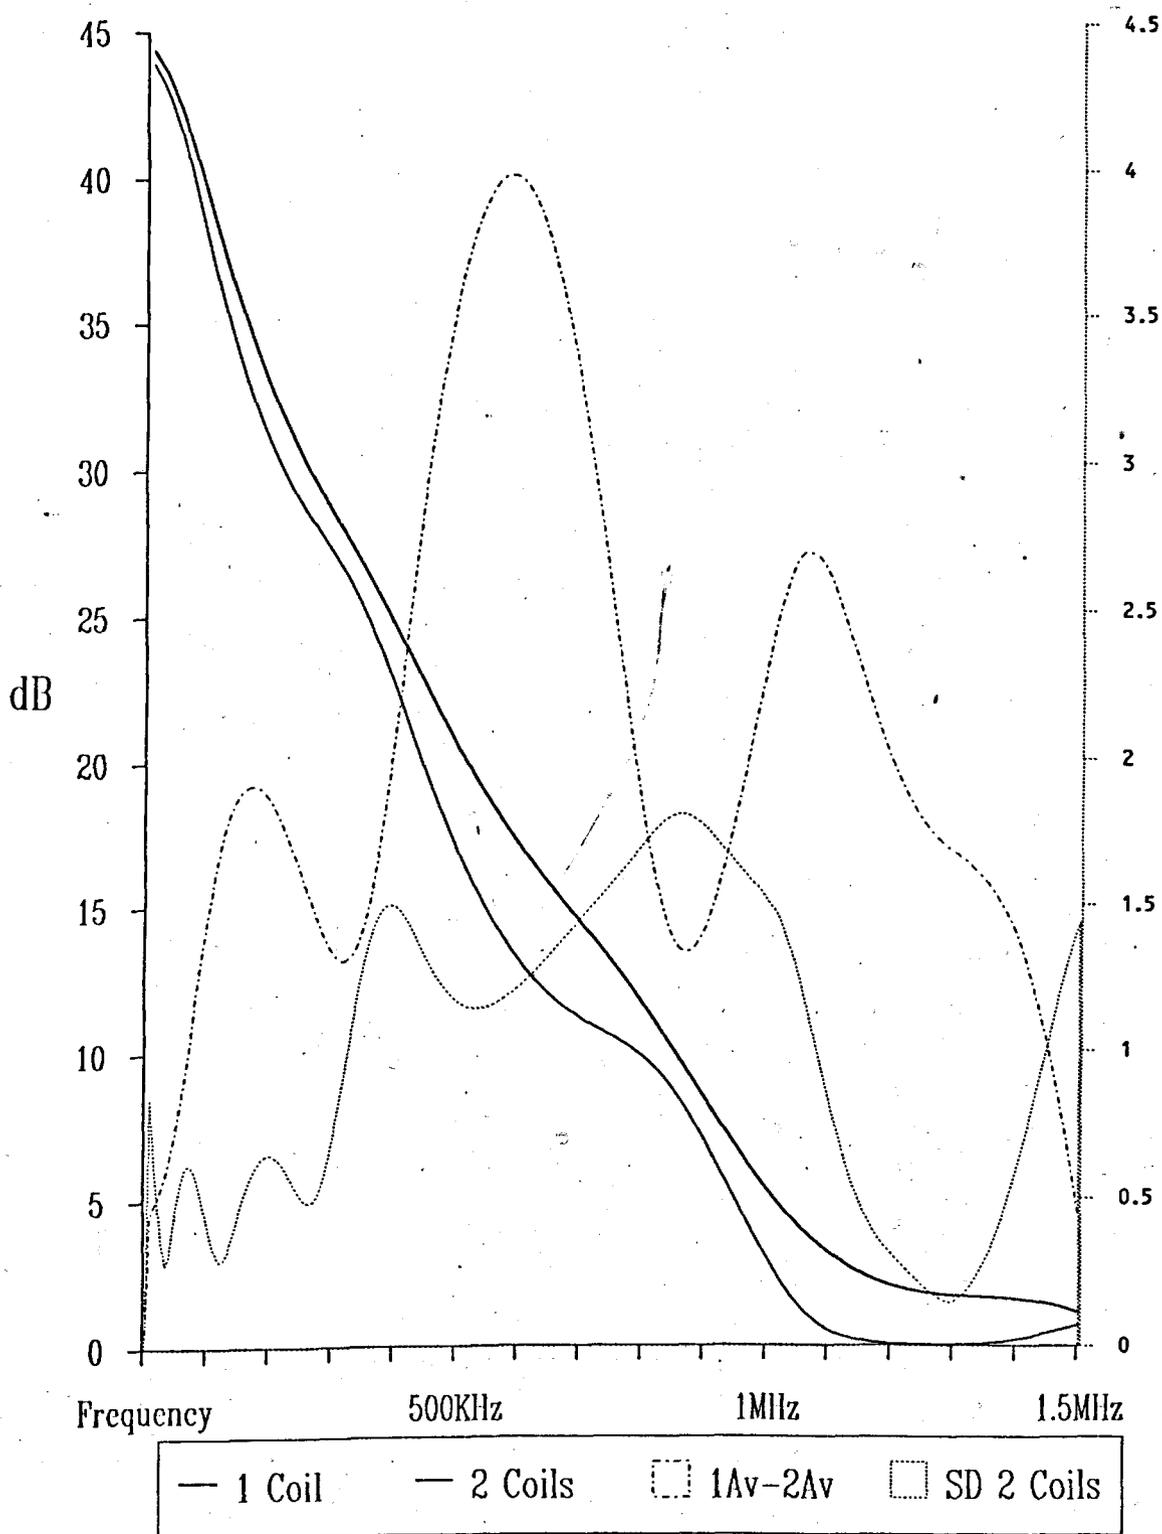


Figure 6.4 MEM Analysis for Coil 3 vs Coil 1

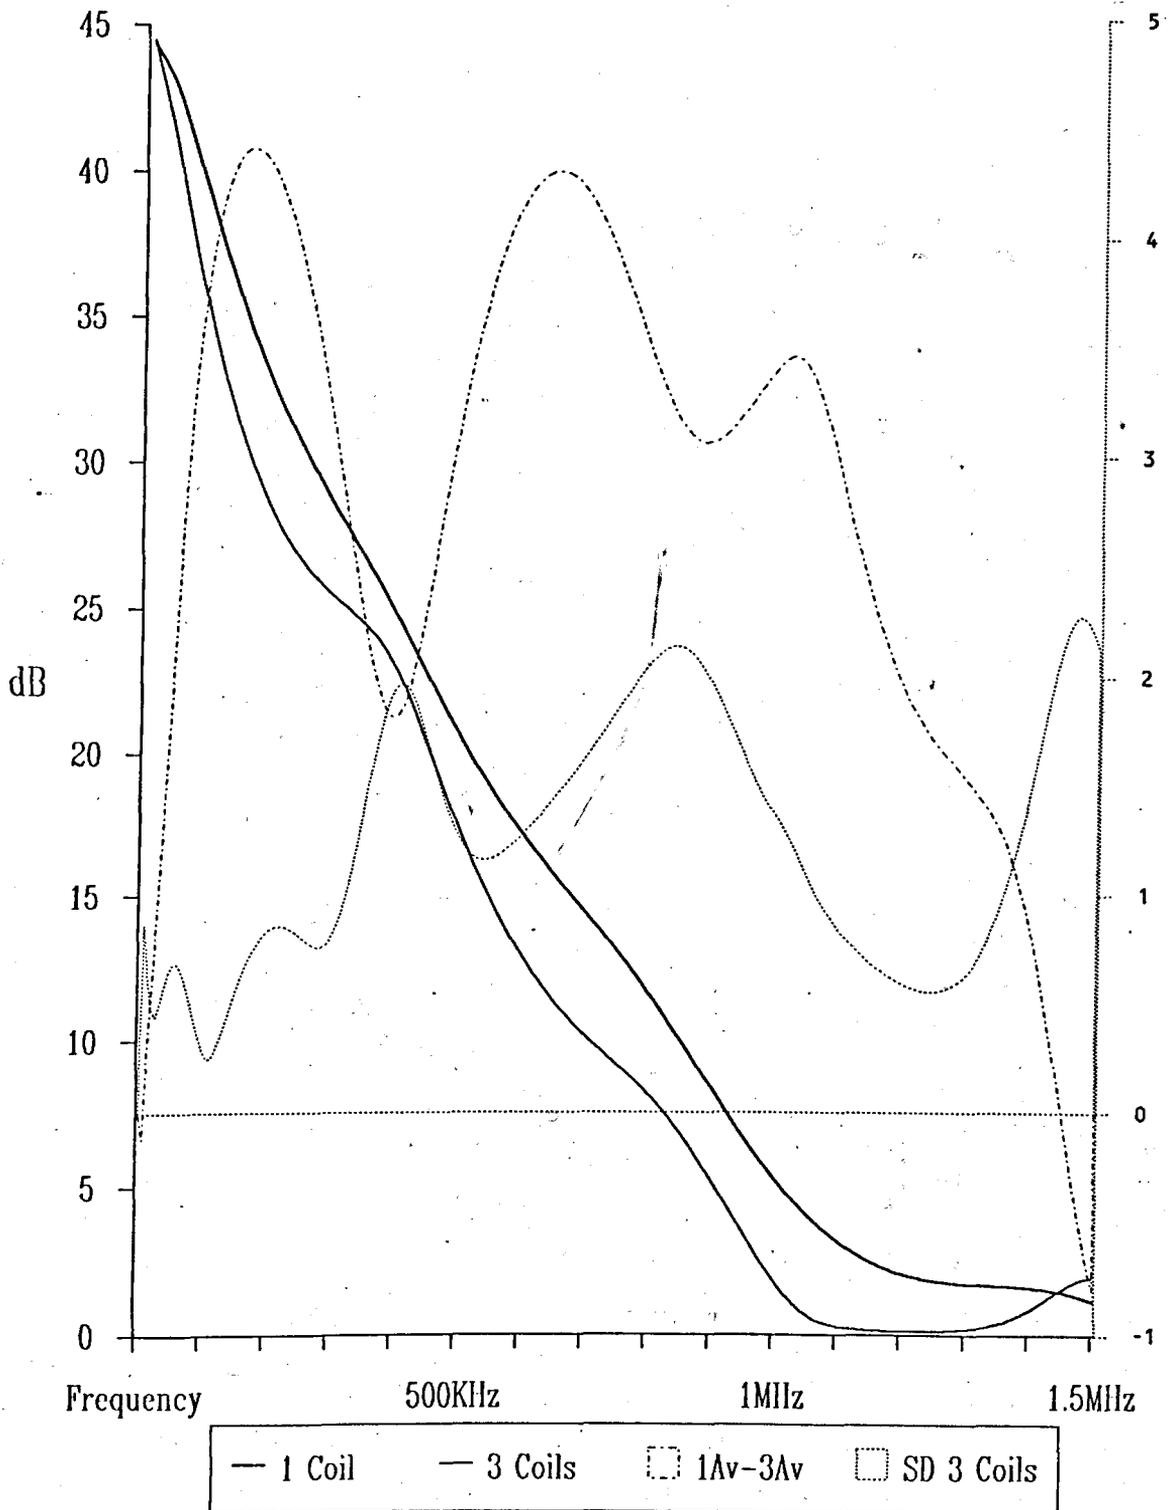


Figure 6.5 MEM Analysis for Coil4 vs Coil 1

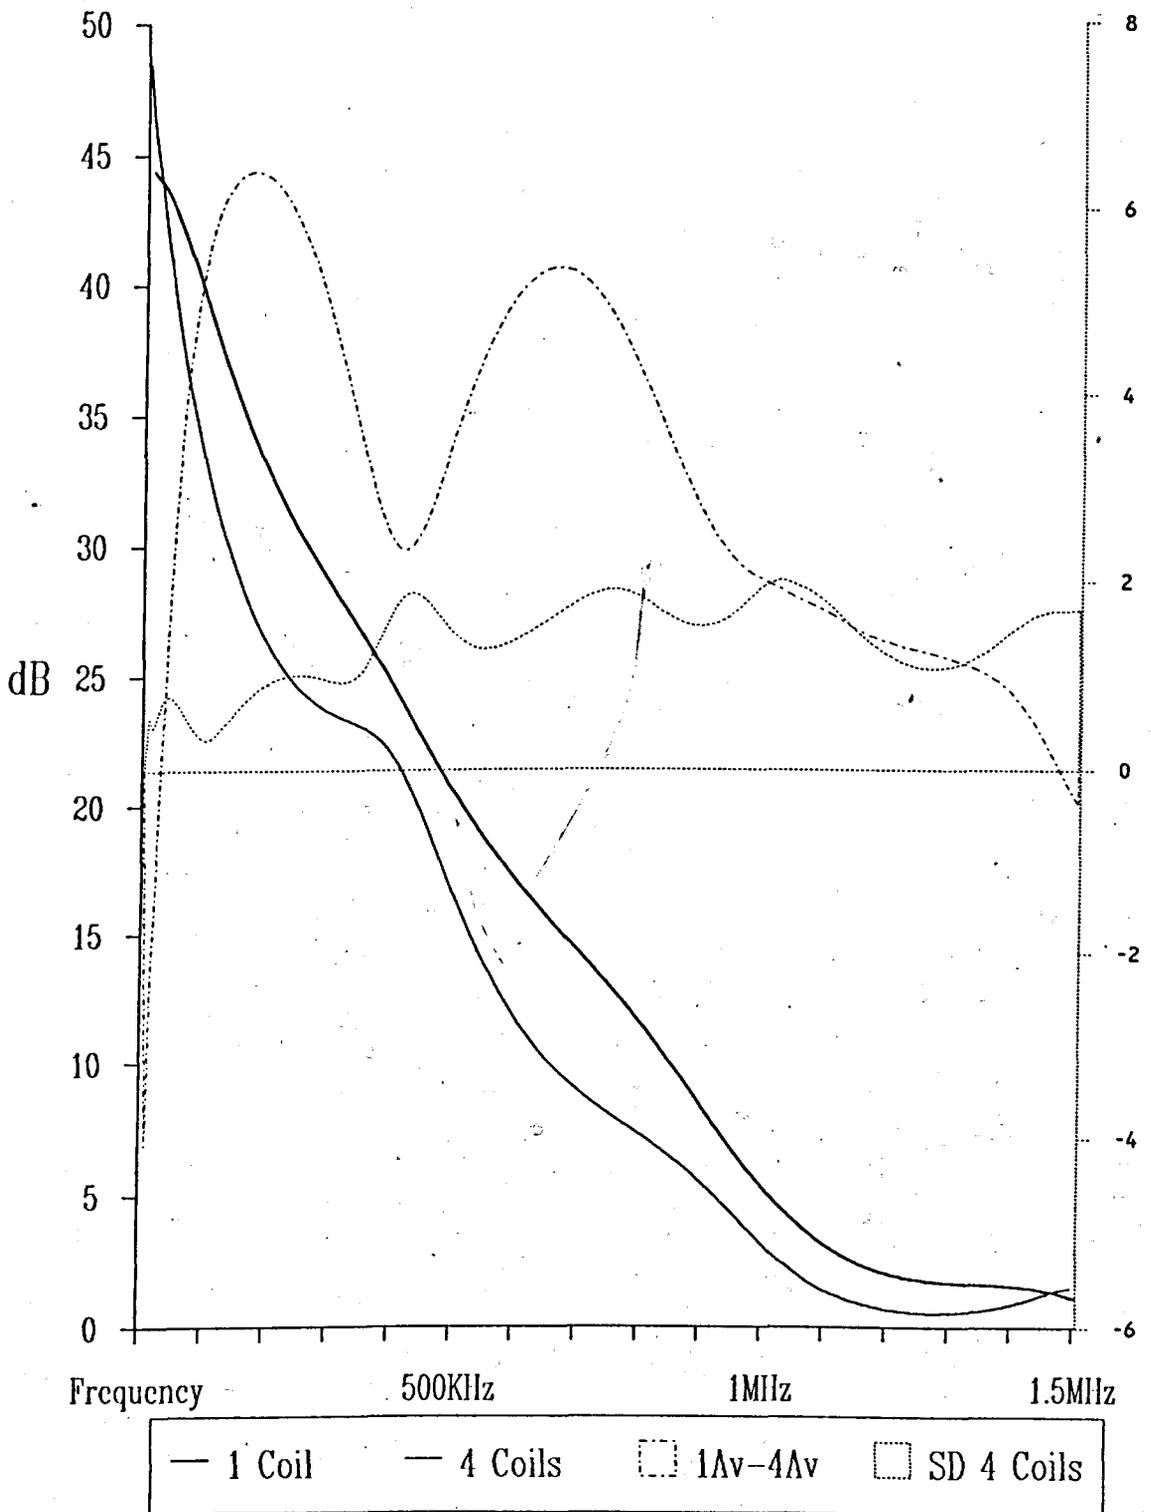


Figure 6.6 MEM Analysis for Coil 5 vs Coil 1

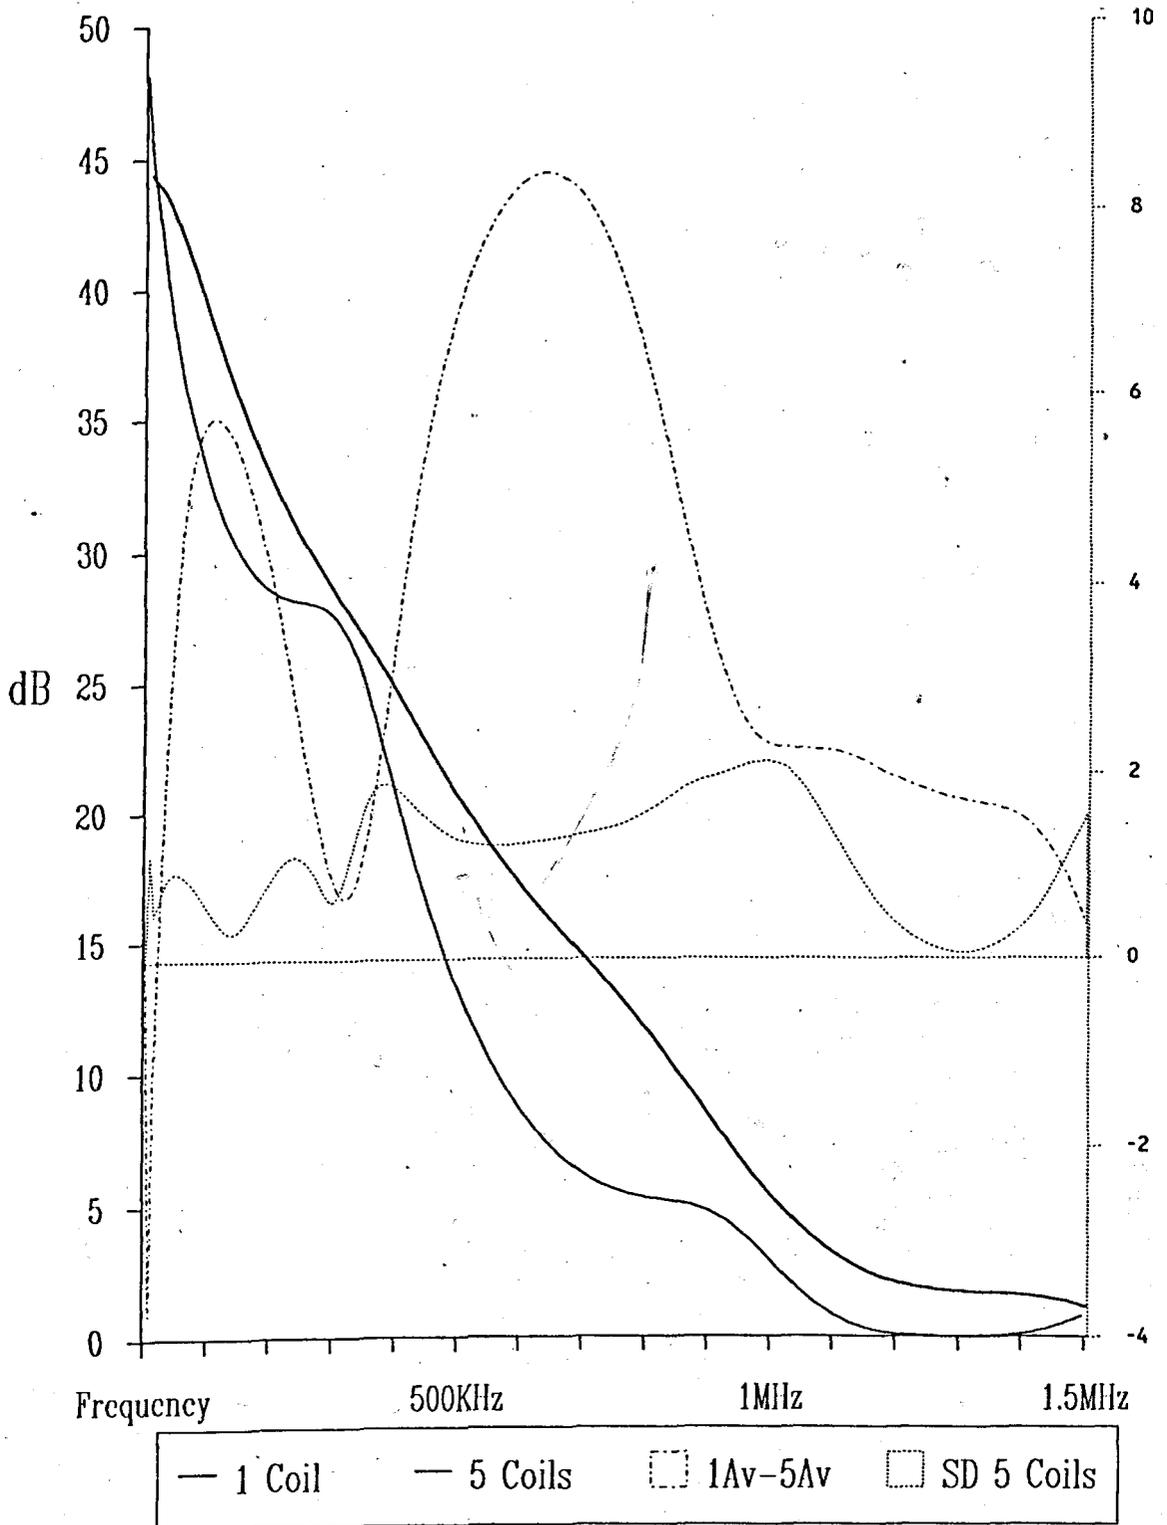


Figure 6.7 MEM Analysis for Coil 6 vs Coil 1

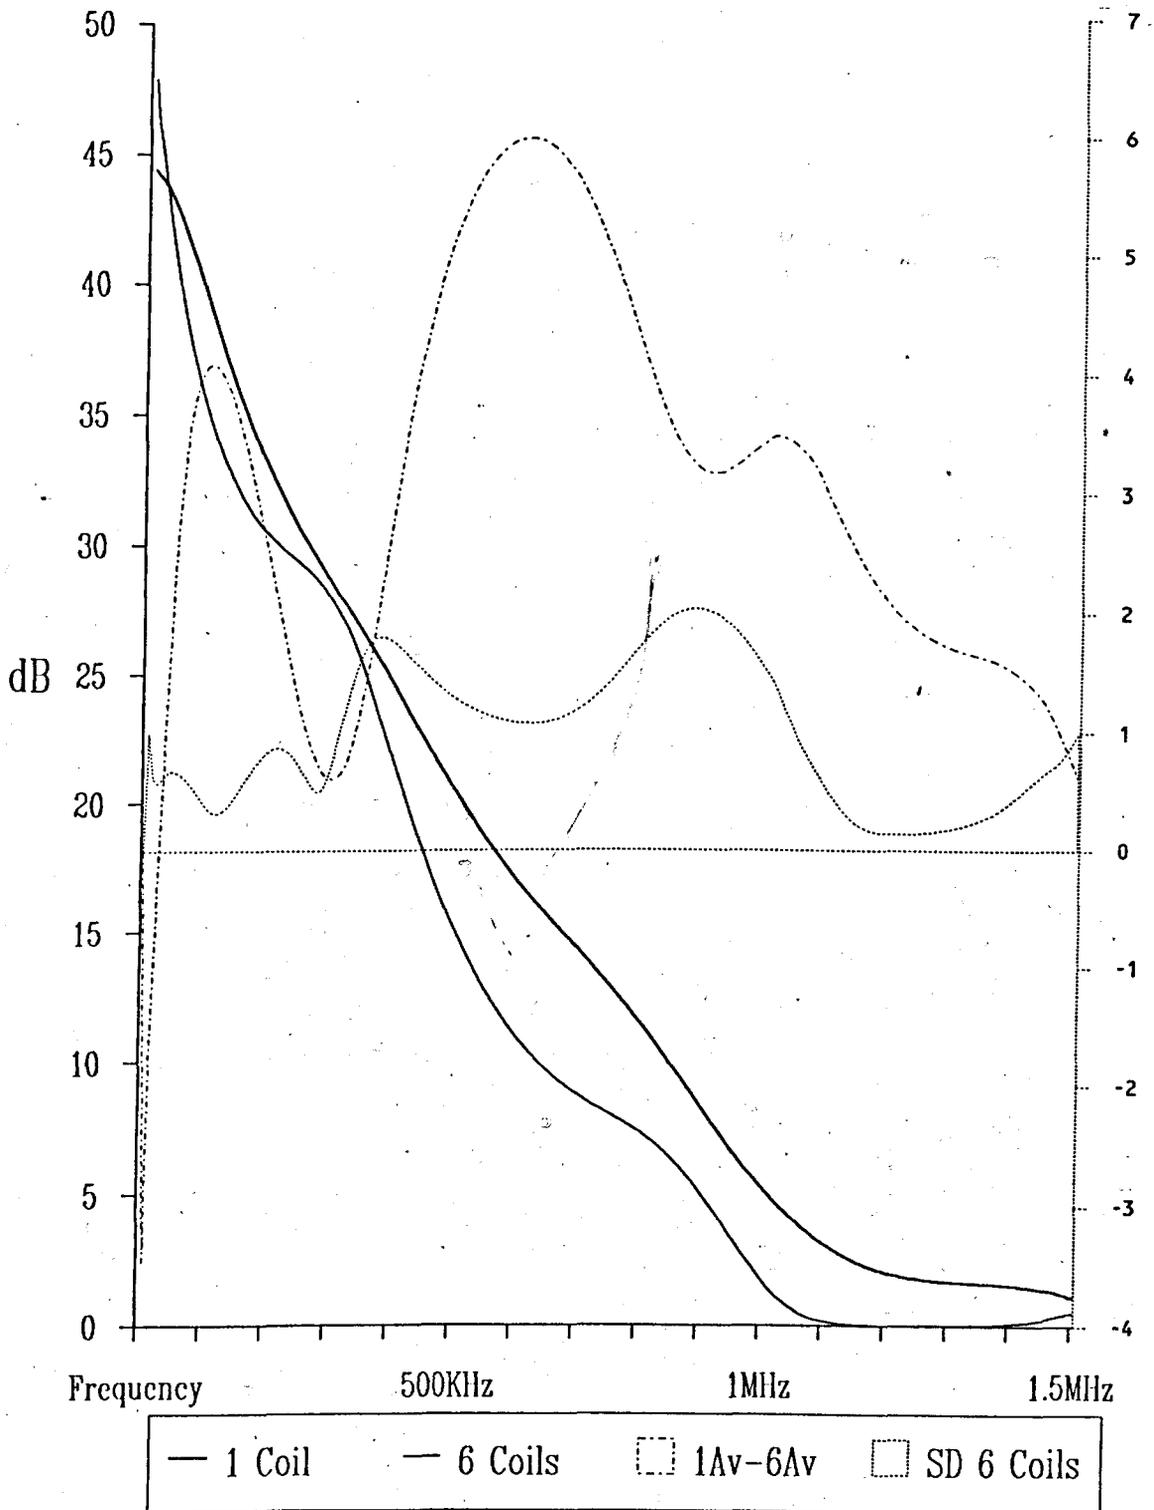


Figure 6.8 MEM Analysis for Coil 7 vs Coil 1

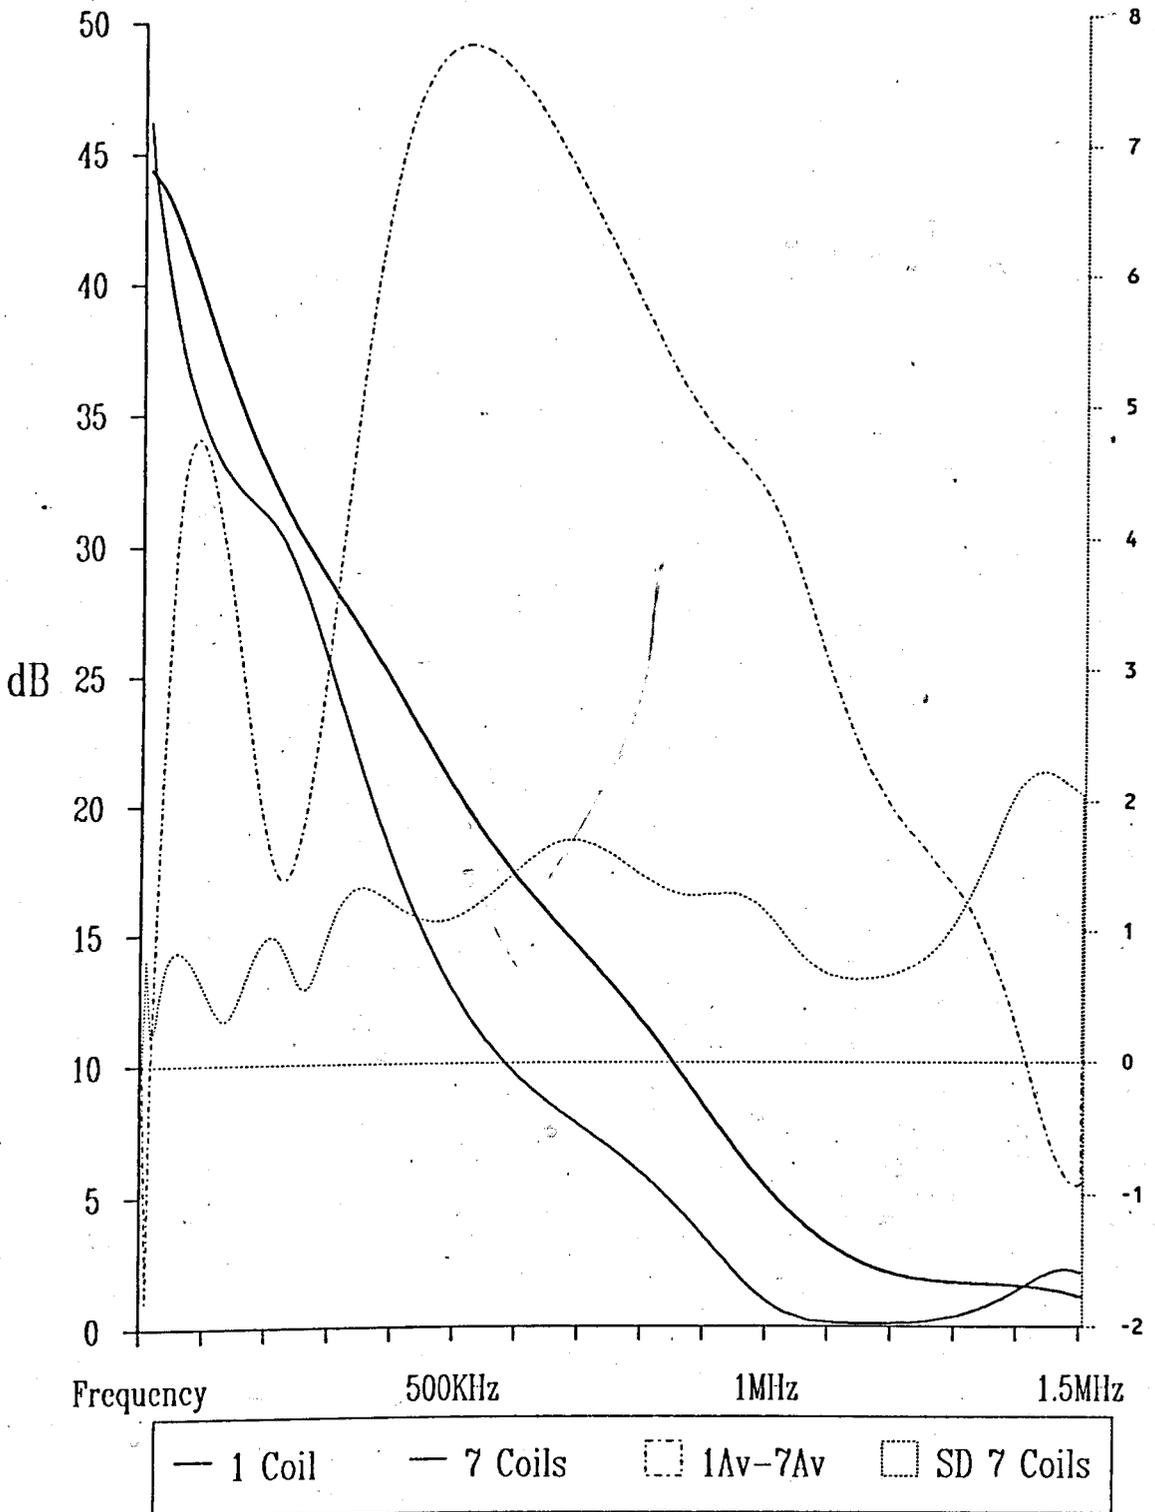


Figure 6.9 MEM Analysis for Coil 8 vs Coil 1

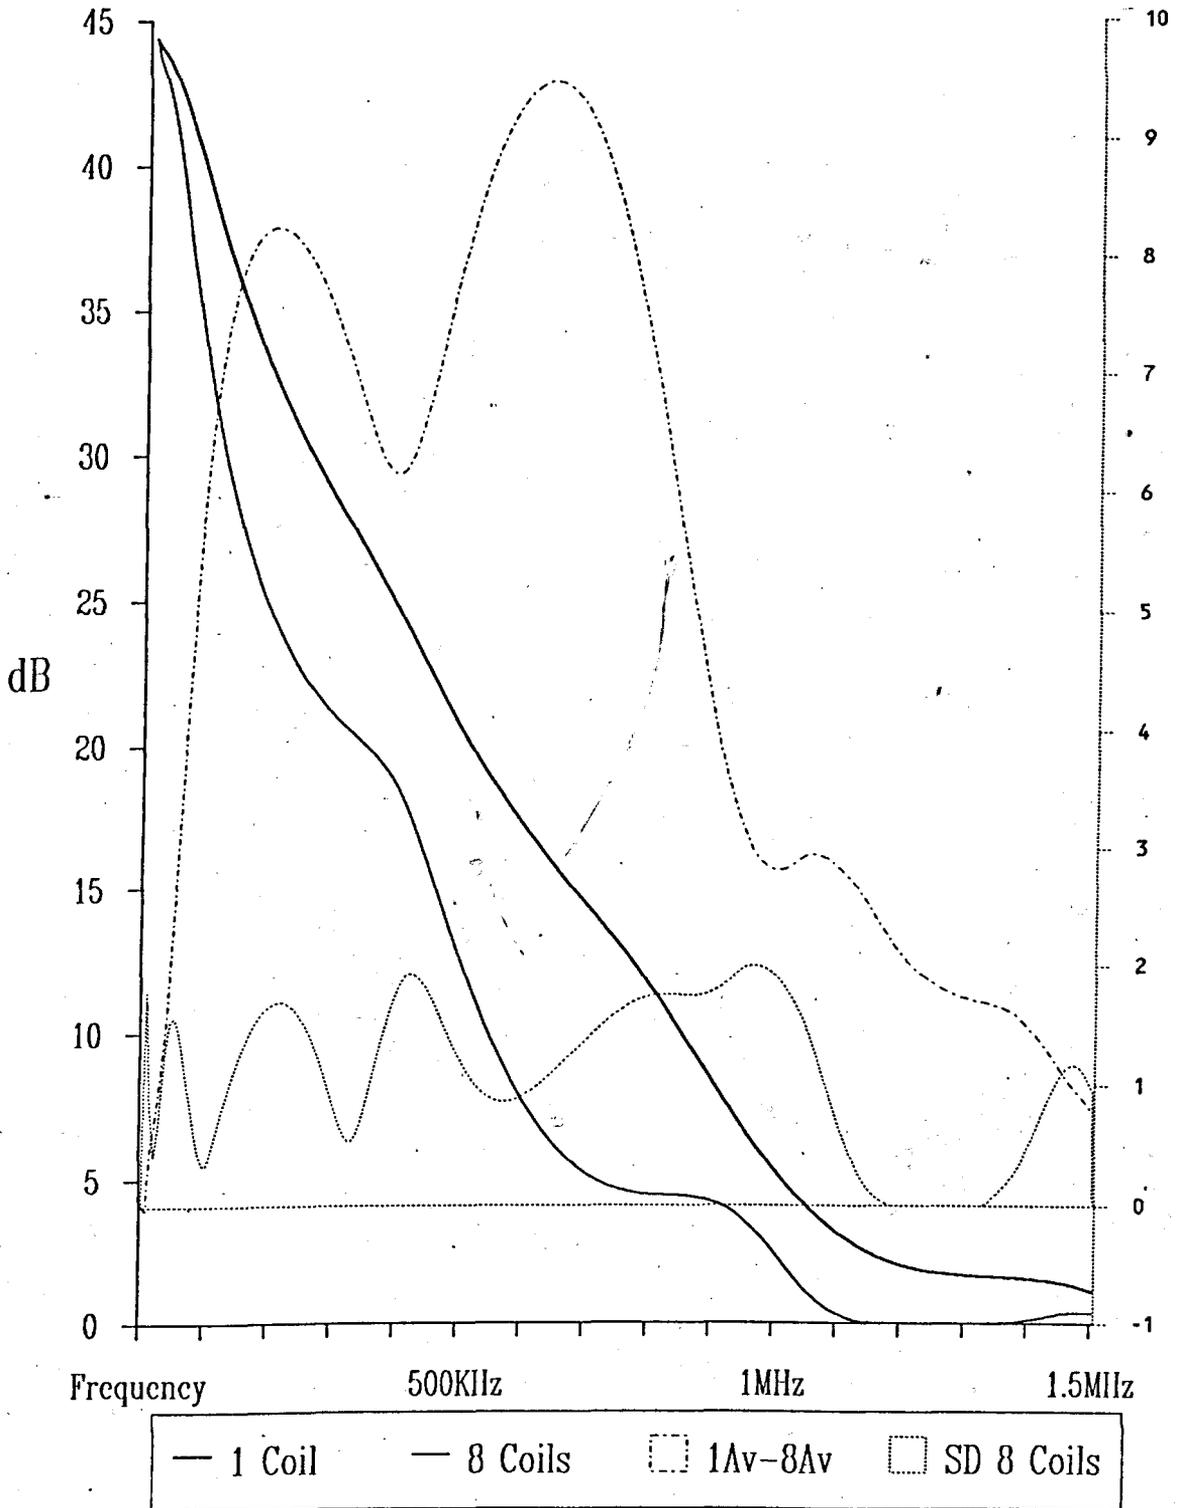


Figure 6.10 MEM Analysis for Coil 9 vs Coil 1

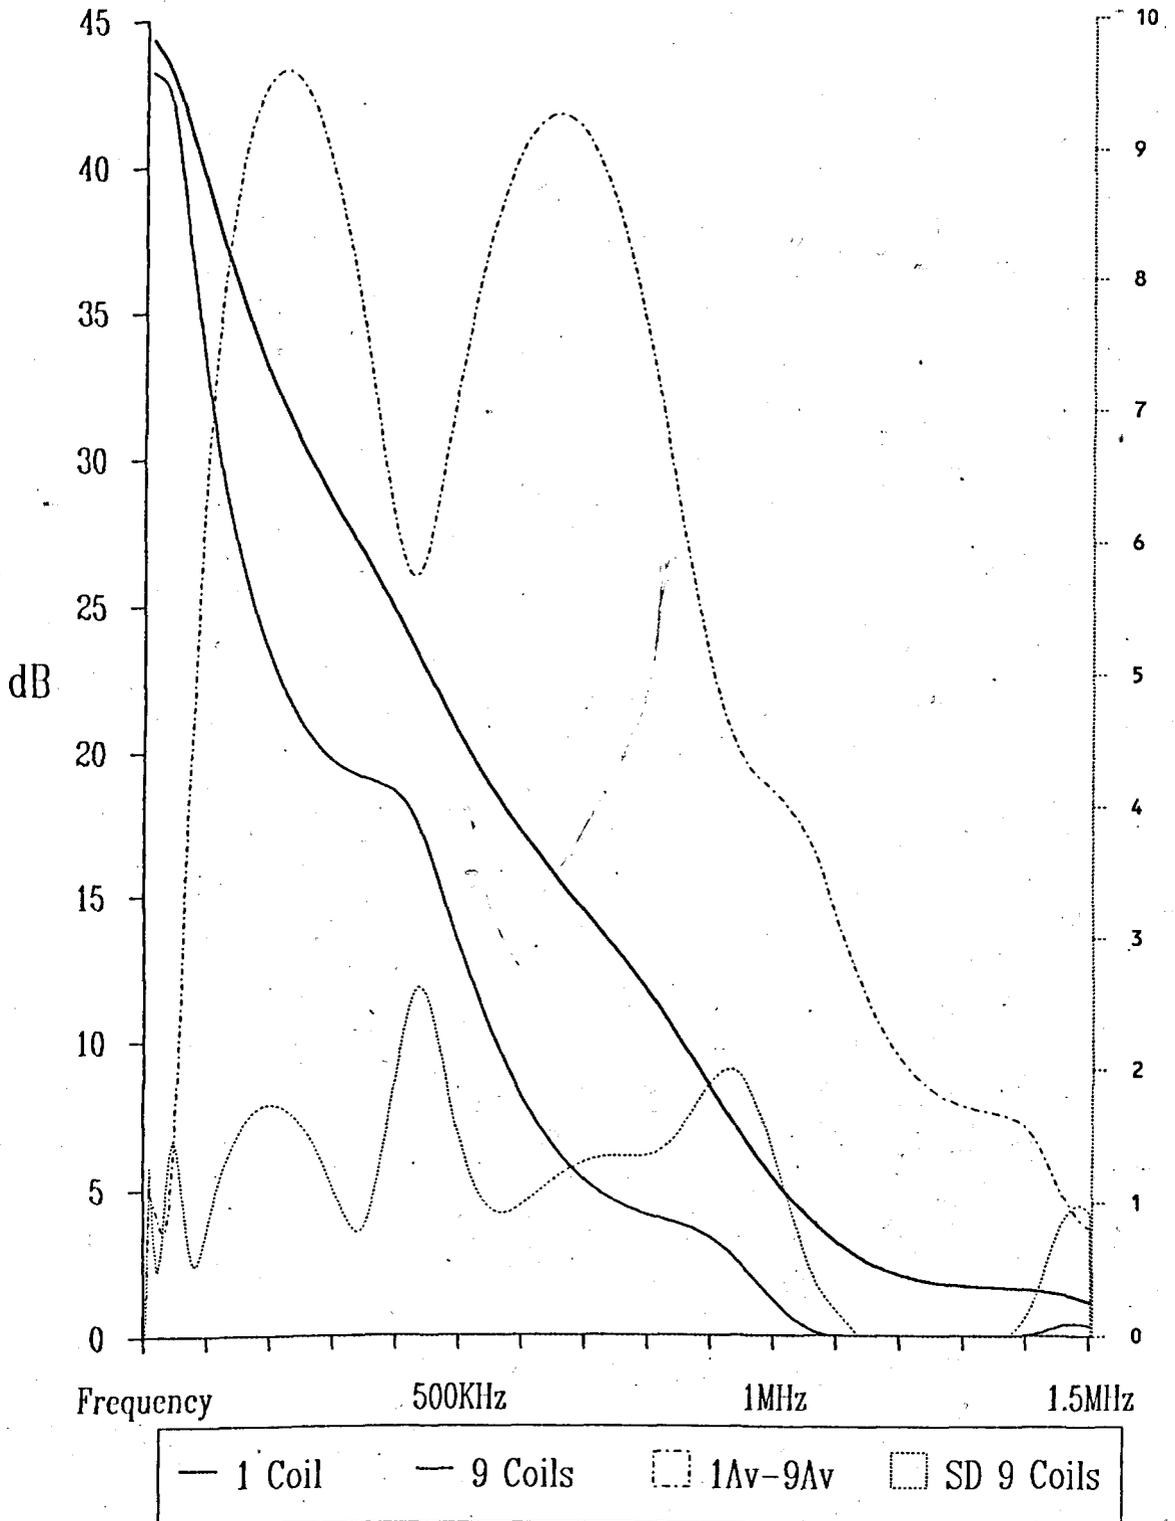
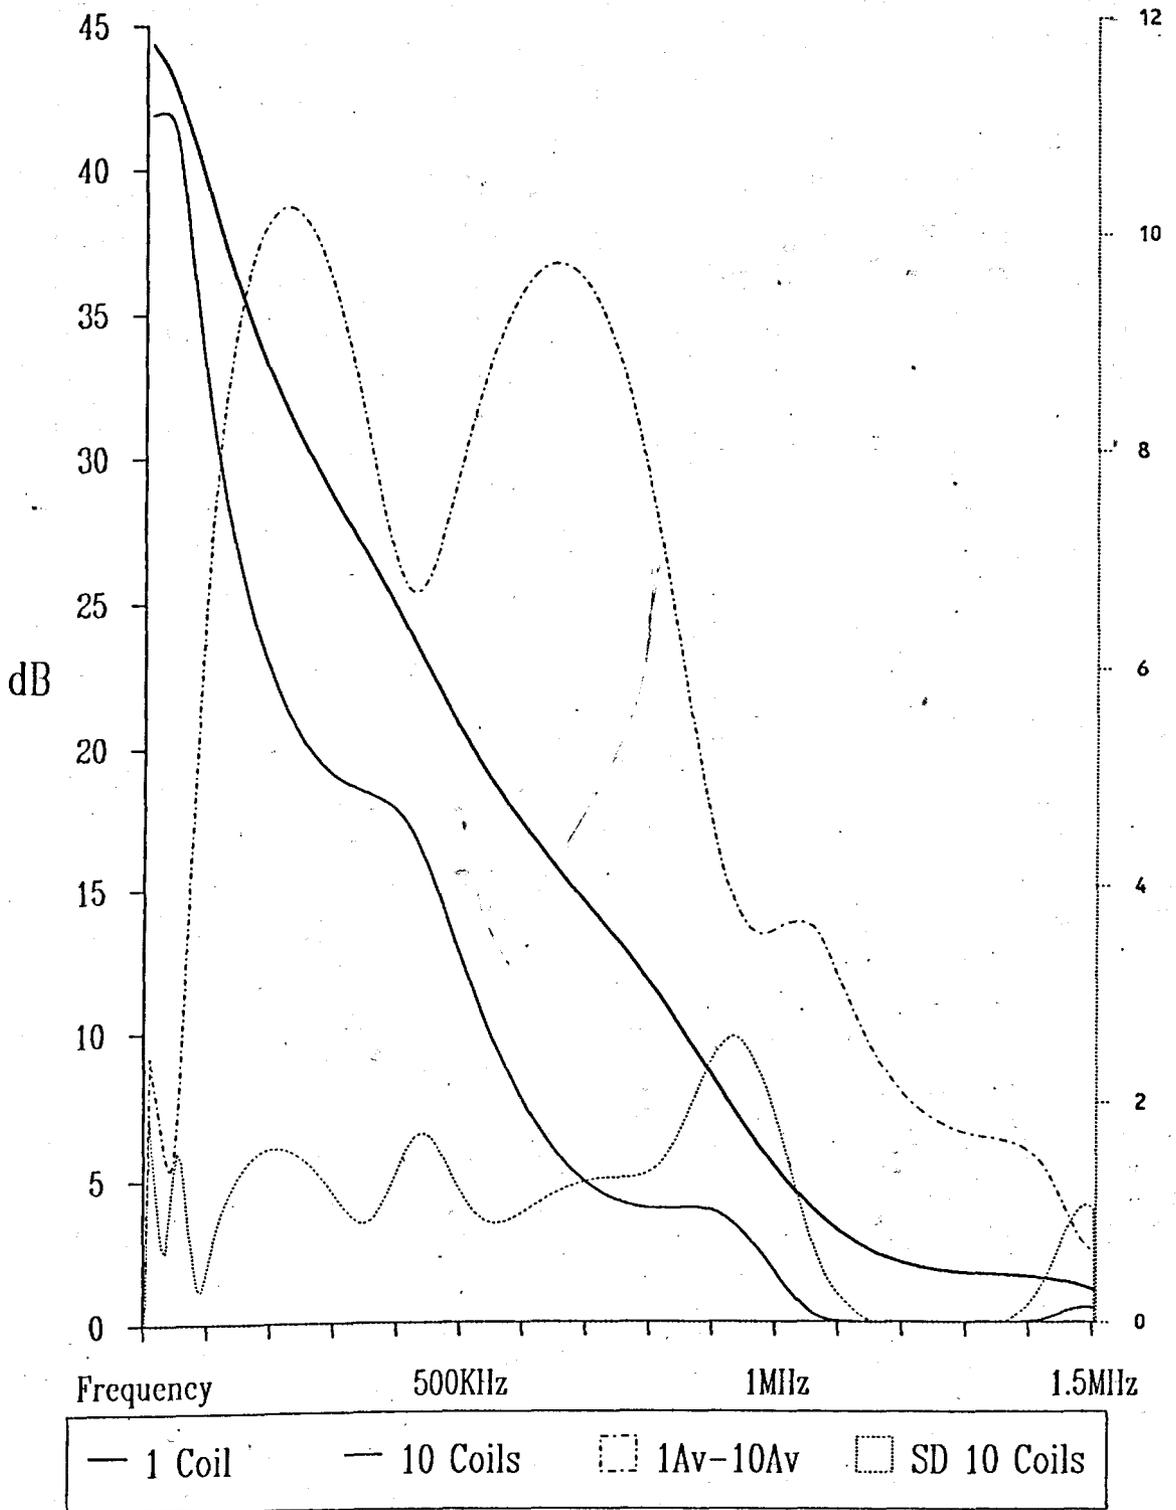


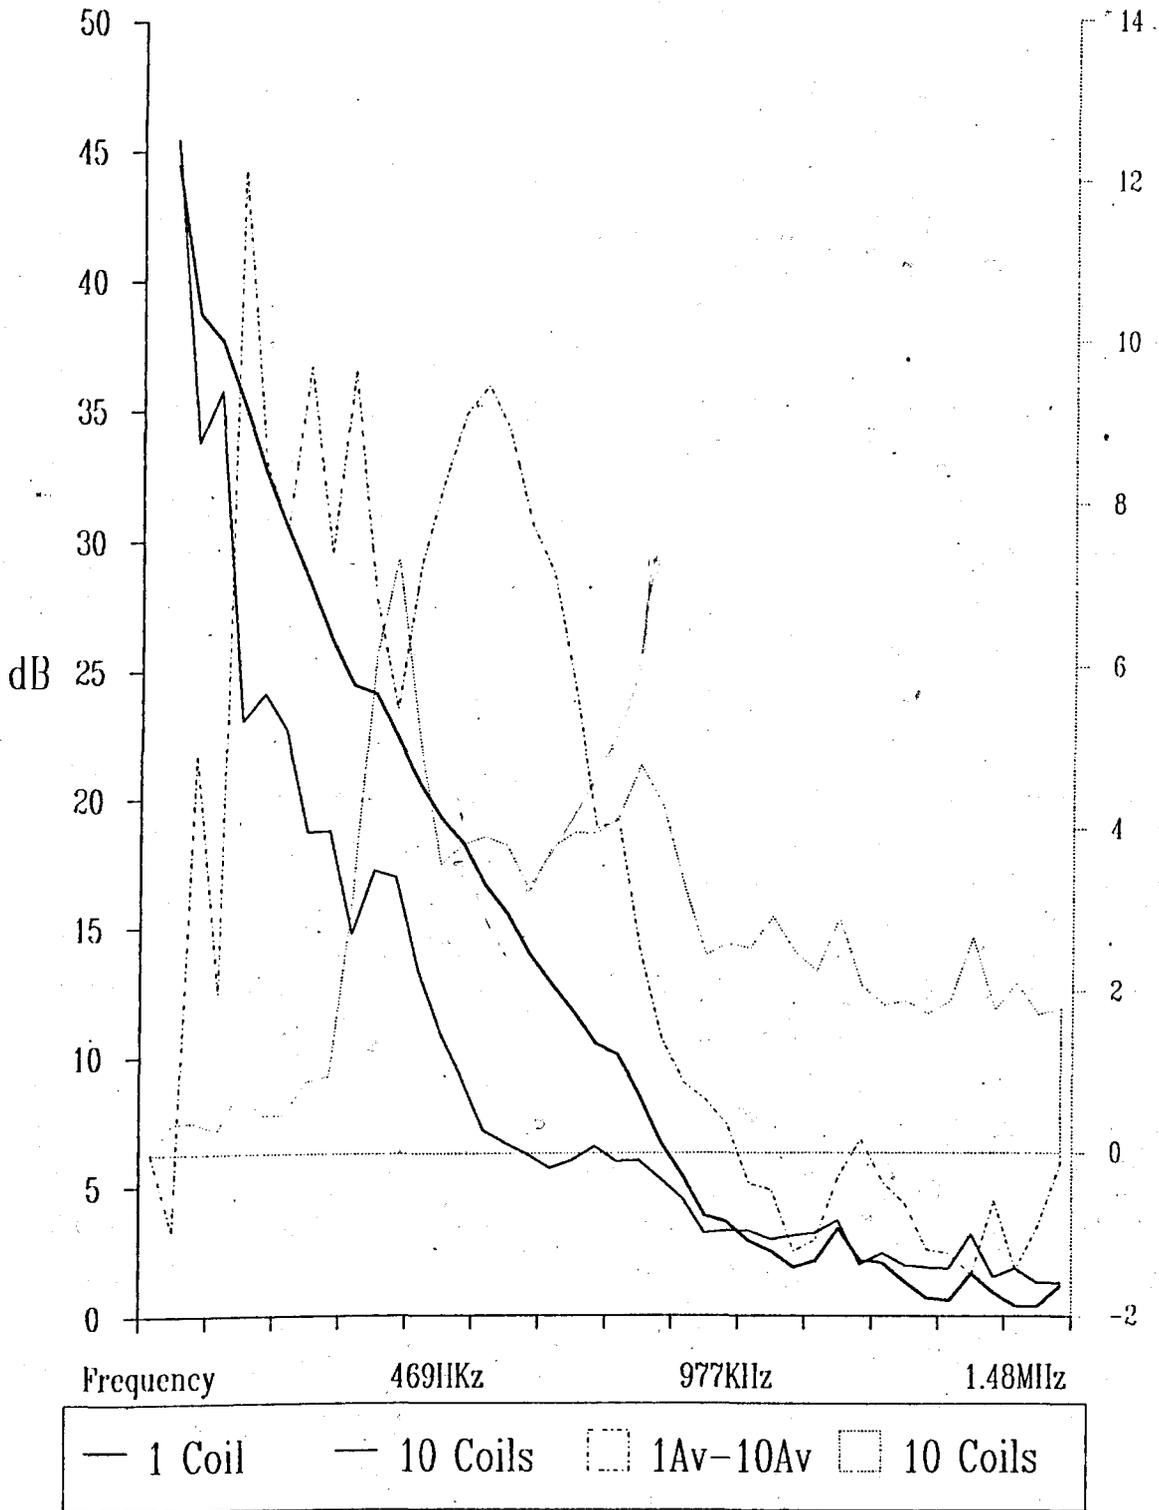
Figure 6.11 MEM Analysis for Coil 10 vs Coil 1



A similar spreadsheet was constructed using the FFT data and graphs plotted for each coil analysis: comparison of the resulting graphs showed that the MEM analysis method produced more precise results than that of the FFT. Figure 6.12 shows the graph, using FFT data for coil 10.

Comparing this with Figure 6.11, the same graph but using MEM analysis, it can be seen that the FFT result has a 'coarser' appearance, with sharp, narrow peaks and a higher overall standard deviation. This detracts from its usefulness for identifying specific frequencies that could be used for classifying individual discharges. As a consequence of this, efforts now concentrated solely on analysis of the Maximum Entropy Method results.

Figure 6.12 FFT Analysis for Coil 10 vs Coil 1



Examining Figure 6.3, it is apparent that there are three main peaks in the difference profile, though not very large. These arise whenever the mean frequency profile for coil 1 exceeds that of coil 2 by more than the standard deviation of coil 2. The peaks occur at approximately 200kHz, 600kHz and 1.1MHz. Considering these peaks, in conjunction with the standard deviation plot for coil 2, it can be seen that all three peaks are larger than the standard deviations associated with each peak, therefore, on the evidence of this graph, it may be possible to use those three peaks in the difference profile as a starting point for classifying pulses originating from coil 2.

Examining Figure 6.4, we see a larger peak, of about 4.5dB, in the difference profile at 200kHz. The second peak seems to have shifted up in frequency slightly, to 650kHz, whilst the third has shifted down slightly, and increased in amplitude. An increase in amplitudes of the difference profiles is to be expected as profiles from deeper coils are examined. However, comparing the difference profiles from Figures 6.3 and 6.4, we see no increase in the peak at about 650kHz, but a doubling at 200kHz: this could be used as a basis for discriminating between pulses emanating from coils 2 and 3.

Switching to the graph for coil 4, Figure 6.5, as expected the peak at 200kHz is seen to rise, to about 6.3dB, and the peak at 650kHz rises also. However, there is no longer a peak at 1.1MHz and, given the standard deviation at that point, statistically the level of the spectral lines there will be quite uncertain and could lead to anomalies in classifying pulses from coil 4: i.e. the difference profile at 1.1MHz is about $1.8\text{dB} \pm 1.9\text{dB}$ due to the standard deviation of frequency profiles from pulses injected at coil 4.

Moving on to Figure 6.6, a slight drop is noted for the peak of 200kHz and, in fact, the peak has moved down in frequency to about 150kHz. The 600kHz peak is seen to rise to 8dB, whilst the absence of a third peak is also noted.

Figure 6.7 shows quite unexpected results: the peaks at 150kHz and 600kHz drop, and a peak at 1.1MHz re-appears. This would normally be representative of pulses injected nearer the line-end of the winding, not 6 coils inside. The reasons for this

occurring are uncertain, but whatever the cause, the particular combination of frequency peaks shown, should enable classification of pulses emanating from coil 6.

The results, depicted in Figure 6.8, extend those shown in the previous diagram, in that the two peaks at 150kHz and 600kHz rise, and the peak at 1.1MHz is starting to disappear. These results indicate that the anomalies of Figure 6.7 are not due to experimental error, but some other property.

Further investigation into this spectral anomaly are outwith the boundaries of this particular project. However, speculating on a possible cause, it could be that the stator winding inductance and capacitance 6 coils distant from the line end are acting as some sort of tuned circuit whose resonant frequency is 1.1MHz or some sub-multiple thereof. If this were the case then a partial discharge behaving, to all intents and purposes like an impulse, would supply enough energy at the resonant frequency to amplify its affected spectral components whilst travelling through the region of the stator where this is occurring. Thereafter the travelling pulse would be attenuated as normal. If this explanation were true, then all other pulses emanating from deeper in the winding would also be affected. However, as they have already travelled some distance through the stator before reaching coil 6, their spectral energy at the excitation frequency may be so reduced that the resonance has little effect.

The difference profile of Figure 6.9 conforms to expectations, in that the two main peaks rise further, but the lower frequency peak shifts slightly up in frequency. Comparing this to the next plot, Figure 6.10, it can be seen that the lower peak is the only one that could be used to distinguish between pulses from coils 8 and 9, because the peak at 600kHz remains constant for both.

Finally, Figure 6.11 shows a slight increase in both main peaks, but, given the level of standard deviation, it seems unlikely that a classifier could distinguish between pulses emanating from coils 9 and 10. This ability is not too important as most partial discharges occur near the line-end of the stator winding, where the electric stress is highest.

Gathering together all of the difference profiles, an overview of how the peaks build up, as pulses are injected deeper into the winding, can be obtained, Figure 6.13.

Here, each plotted line represents the difference between individual profiles and the reference, and further subtracting the standard deviation.

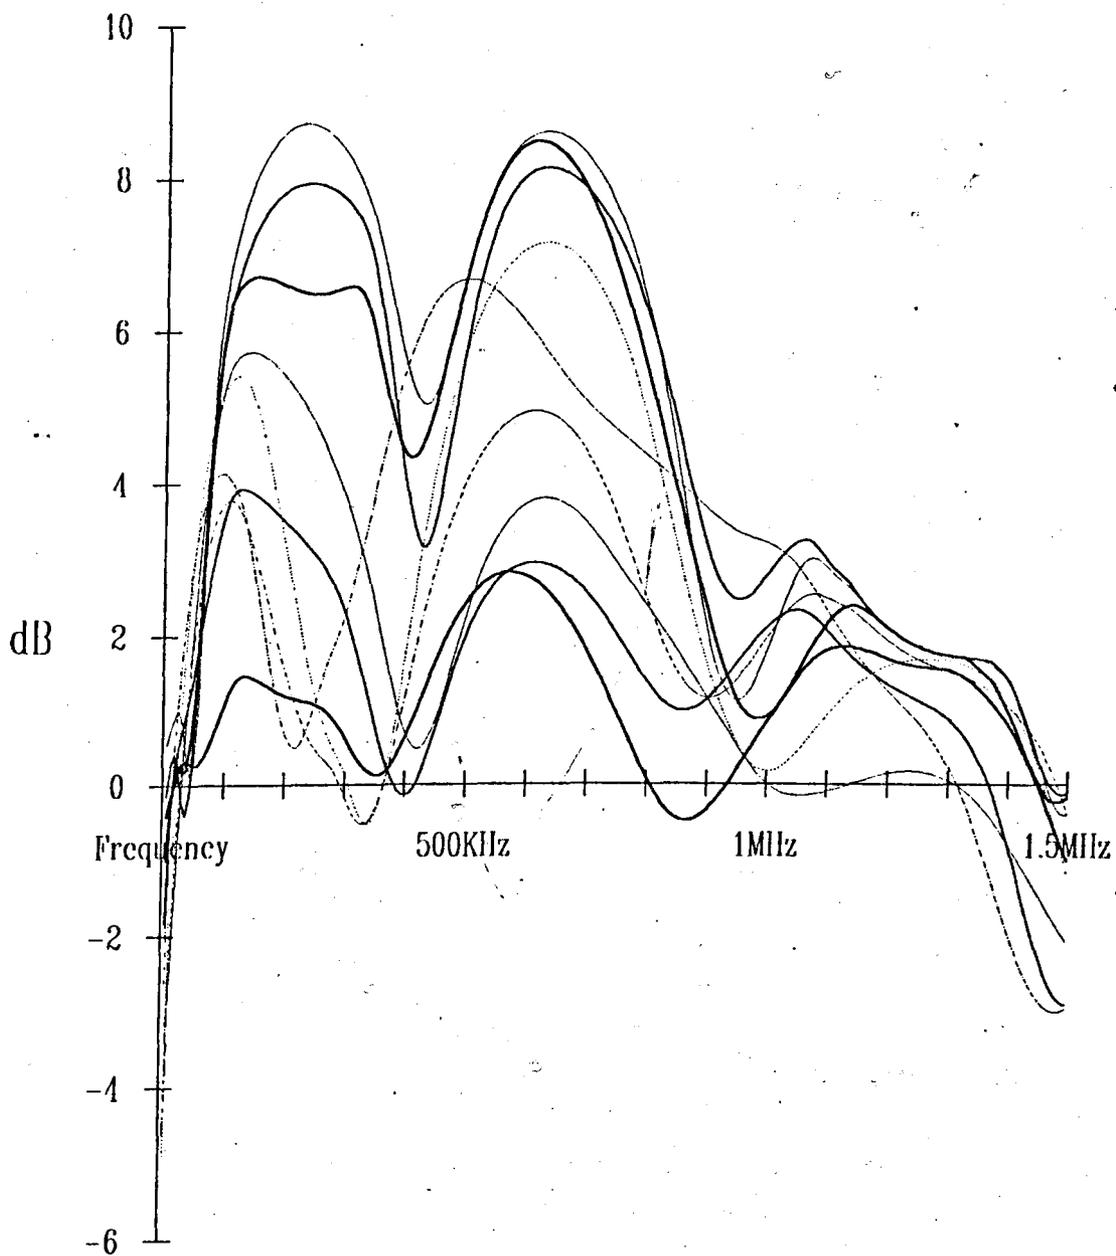
e.g. The line for coil 3 is:

$$\text{mean1} - \text{mean3} - \text{sd3}$$

This graph tends to highlight the frequency attenuation characteristics as pulses travel further within the winding, and can be used to visually identify potential frequencies to be used in an automatic classifier, capable of recognising where a pulse originates from.

Figure 6.13 is quite cluttered, but it does highlight the changes in the two main frequency peaks in the difference profiles. It can be seen that these two peaks tend to swing up and down the spectrum, and vary in width, as pulses travel further through the winding. However, these two peaks could well be used as a basis for an automatic classifier, therefore, it was decided to examine frequencies around 150kHz and 600kHz with a view to developing a rudimentary partial discharge location classifier, in order to investigate the feasibility of such a scheme.

Figure 6.13 Differences in Average Frequency and Standard Deviation of the Deeper Coil



— 1/2-2Sd	— 1/3-3Sd	— 1/4-4Sd	— 1/5-5Sd	— 1/6-6Sd
— 1/7-7Sd	— 1/8-8Sd	— 1/9-9Sd	— 1/10-10Sd	

6.3 A Rudimentary Partial Discharge Location Classifier

The main problem in designing a partial discharge location classifier is identifying specific features that are unique to pulses from individual locations. Examining the two previously identified frequencies, 150kHz and 600kHz, for all 10 pulse locations, Figure 6.14, it can be seen that the combination of the two selected frequencies change for every location within the stator. However, if an uncertainty factor is added, utilising the standard deviation, the situation changes somewhat, Figure 6.15.

Figure 6.15 takes the two frequencies of interest, 150kHz and 600kHz, and plots how they vary as pulses from different locations within the winding are analysed. Instead of displaying one value for each location, the diagram depicts a range, statistically within which the value should be, given analysis of an unknown pulse. The range displayed is calculated using the standard deviation values: i.e. On the whole, a analysed pulse will have its frequency component within the mean \pm one standard deviation.

Figure 6.14 Change in Specific Frequencies as Pulses Travel from Deeper Coils

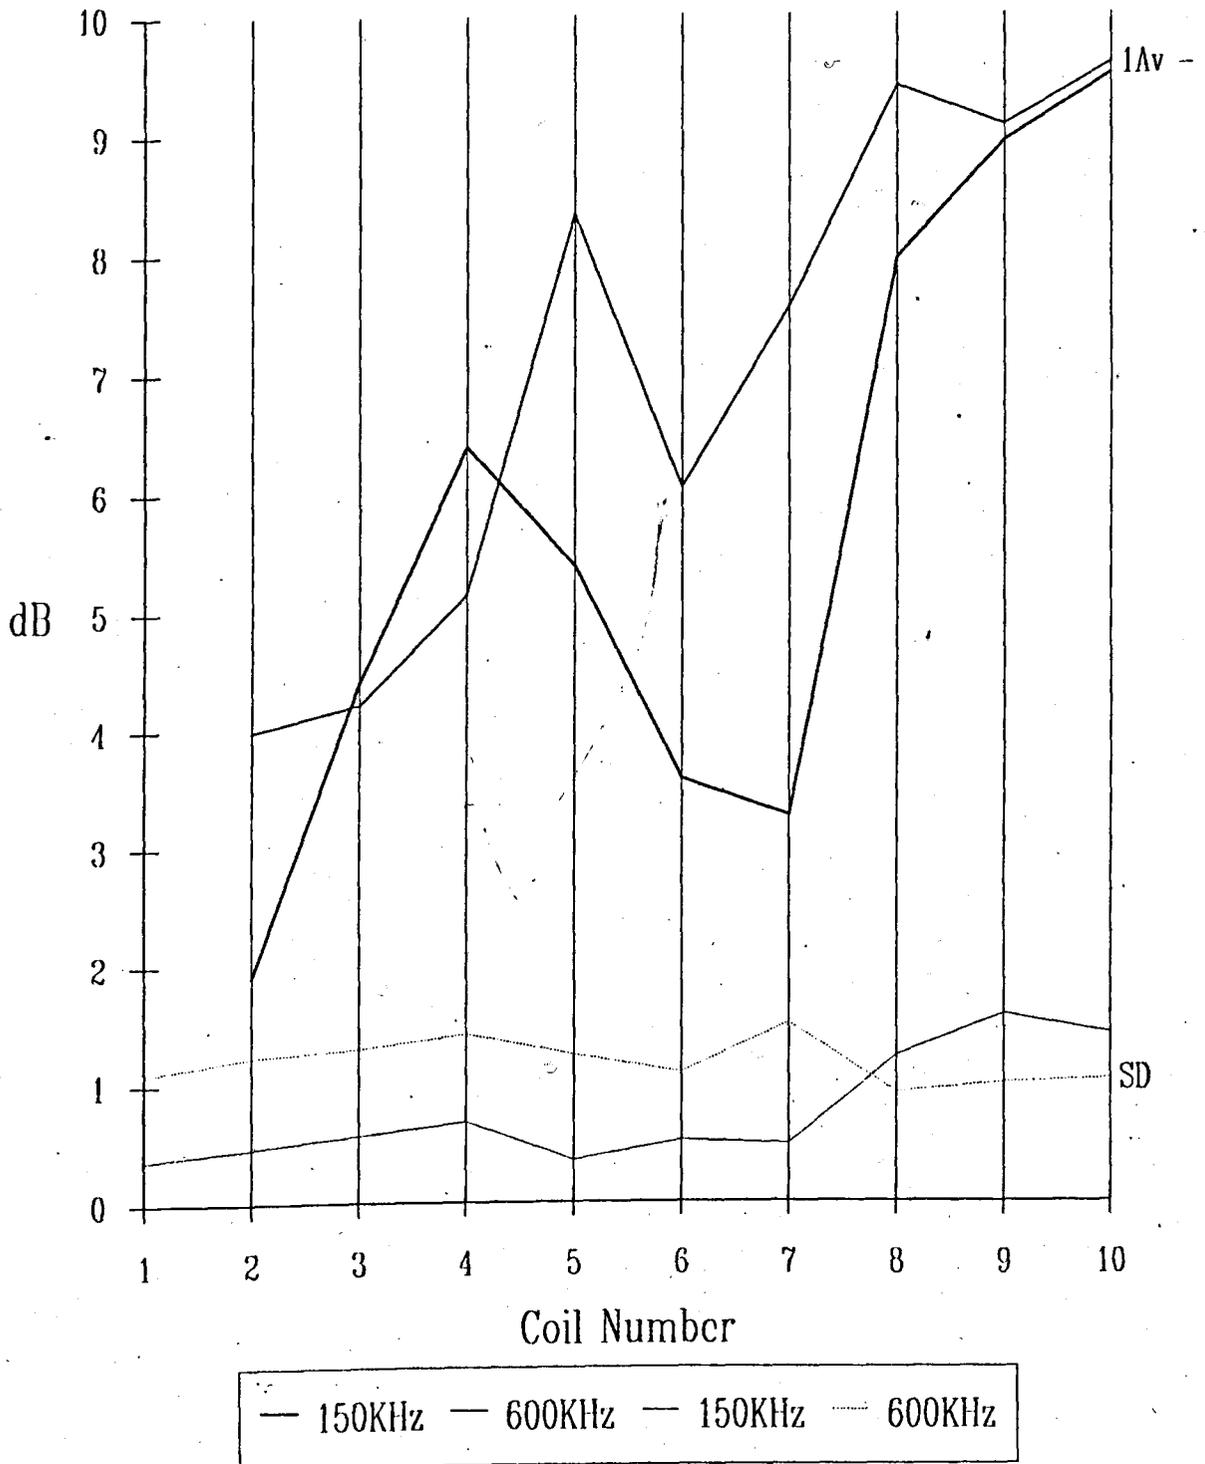
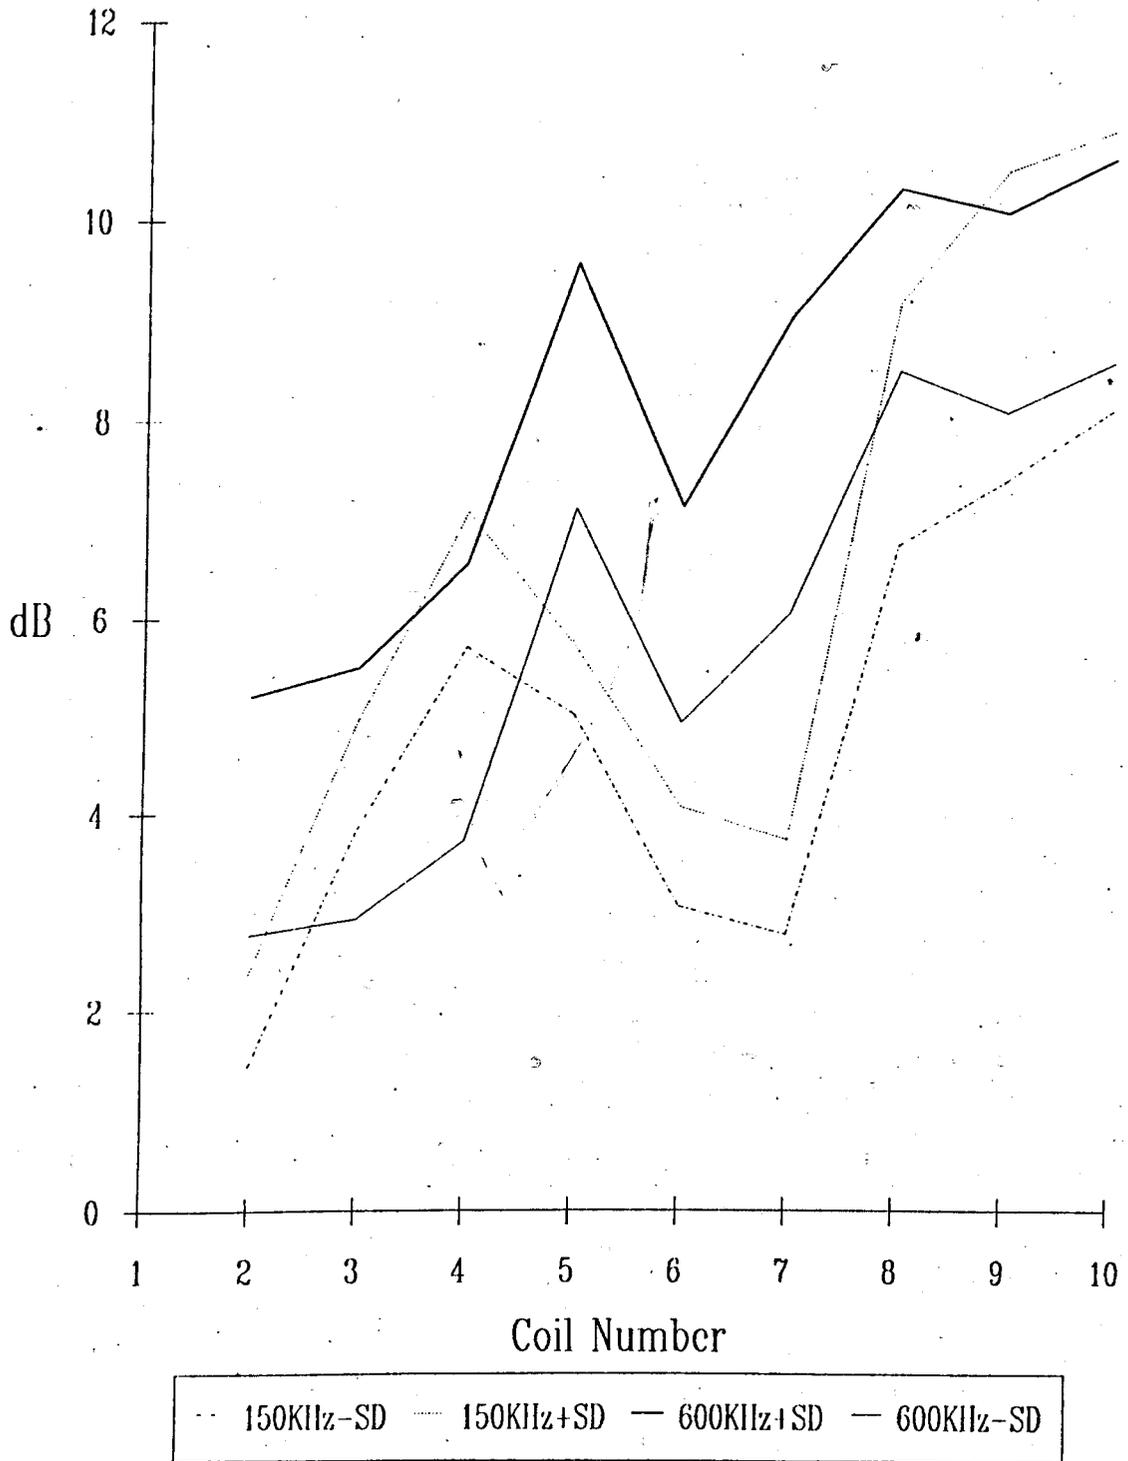


Figure 6.15 Frequency Difference Profile with Uncertainty Factor



Examining the situation for pulses coming from coils 8, 9 and 10, it becomes apparent that, for the two frequencies of interest, the range of probable values makes it very difficult, and sometimes impossible, to uniquely categorise a pulse as coming from coil 8, 9 or 10. Closer examination of the value ranges for the other coils reveals overlap regions, whereby it would be possible for certain analysed pulses to appear as if they could come from one or more locations.

e.g. If an analysed partial discharge pulse had a 150kHz component (when subtracted from the reference) of 4dB and a 600kHz component of 5dB, then examination of Figure 6.15 reveals that the pulse could be from coil 3 or coil 6!

This overlap effect, due to the variance of 50 analysed pulses from the same location, means that these two frequencies, 150kHz and 600kHz, are unlikely to be very successful if they are used as the only criteria in a partial discharge location classifier. Therefore, one or more other suitable frequencies have to be identified.

Re-examining the mean frequency profiles for all 10 locations, Figure 6.13, a frequency is required which changes quite a lot from location to location. Comparing the 300kHz components for coils 3 and 6, it can be seen that they vary substantially. Thus if a third frequency component, at 300kHz, were to be added to the location classifier, an improvement in diagnostic capabilities would undoubtedly result.

In order to demonstrate the feasibility of partial discharge location within a stator winding, a rudimentary classifier was developed. It is a very simple routine that just compares the amplitudes of the three frequency components of an unknown pulse, to that of the reference data from coil 1. Given the relative amplitudes of these three frequency components, the classifier then makes a decision as to where the unknown pulse originated from. A flowchart of the algorithm can be seen in Figure 6.16.

The partial discharge location classifier proved to be quite successful when used on the experimental data: it can successfully identify where 80-90% of unknown pulses come from. Considering it only compares 3 frequencies, and the decision thresholds are fixed and set by inspection, it ably demonstrates that a practical classifier could be readily developed.

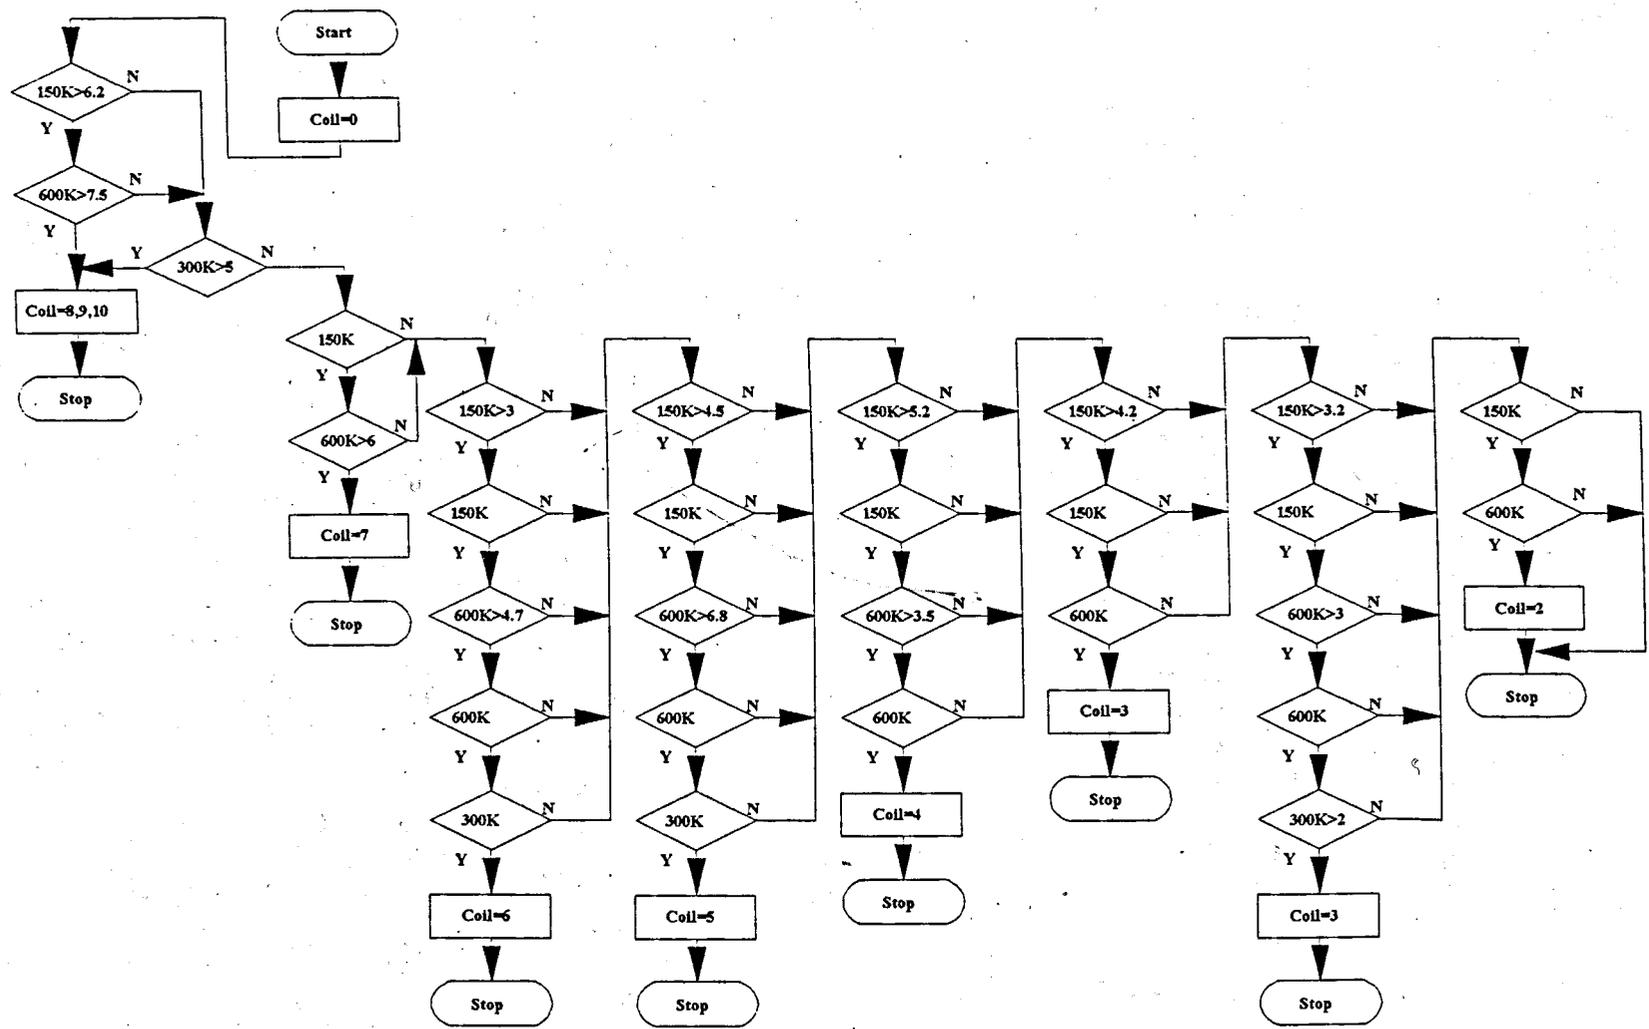


Figure 6.16 Classifier Flowchart

6.4 Summary

Analysis of the raw data using the features of a spreadsheet is far simpler than designing specific analysis algorithms. Spreadsheets allow the rapid manipulation of data into different "what if" scenarios. This makes them ideally suited to the investigation into the correlation between different data sets.

The differences between the spectral analyses were highlighted and utilised to build a simple location classifier. Whilst this classifier demonstrated the principle, it is a rather false example as it was tested purely on the experimental data that was used to build it. A large amount of further work would need to be done in order to produce any practical application of this method.

CHAPTER 7

SYSTEM APPLICATION AND FUTURE DEVELOPMENT

7.1 Major Achievements

At the moment, partial discharge analyses rely upon experts visually interpreting results and making judgements based upon their knowledge and experience. Techniques have been shown here that describe a completely new approach to the analysis of partial discharge data:

- A high speed analogue and digital electronic system was specified, designed and developed that is capable of directly digitising the extremely short duration partial discharges.
- Modular software routines were designed to control the electronics and display the digitised information for analysis.
- Additional software analysis routines were developed to process the discharge information in the time domain.
- A completely new technique was devised that allows the removal of interference signals from the monitored data that could otherwise prevent accurate analysis of the partial discharge severity.
- Different methods were investigated in order to measure the frequency content of partial discharges. This had never been attempted before, primarily due to the lack of suitable technology.
- Finally, a frequency analysis algorithm was devised and applied to experimental data in order to investigate the feasibility of using the frequency content of partial discharges to ascertain their inception point within the stator winding.

7.2 Monitoring Setup

The ultimate aim of this work was the development of an instrumentation system that would have the capability of processing the semi-random partial discharge pulse patterns and identifying the characteristics which have most relevance to diagnosis of insulation condition. To this end both dedicated hardware and appropriate software algorithms have been developed.

Present conditioning monitoring of the insulation health of machines uses a set-up similar to that depicted in Figure 7.1. This arrangement only shows one phase of a machine being monitored. In reality, all three phases are monitored and their signals recorded.

This monitoring technique relies upon an expert who can visually examine the changing pulse record and pick out the relevant large discharge peaks, correlating all three phases to see if any other unusual discharge effects are present.

The new system aims to directly replace the process, display and tape recorder sections of Figure 7.1, giving the expert access to diagnostic algorithms that are not directly available using the existing set-up.

The existing method of visually examining the partial discharge pulses is supported, using the various display options, as it is still the best way of quantifying discharge activity within a machine. But, the additional pulse, noise removal and frequency analysis algorithms can serve as a very useful supplement in the diagnosis of insulation health, especially when integrated into one complete instrumentation system.

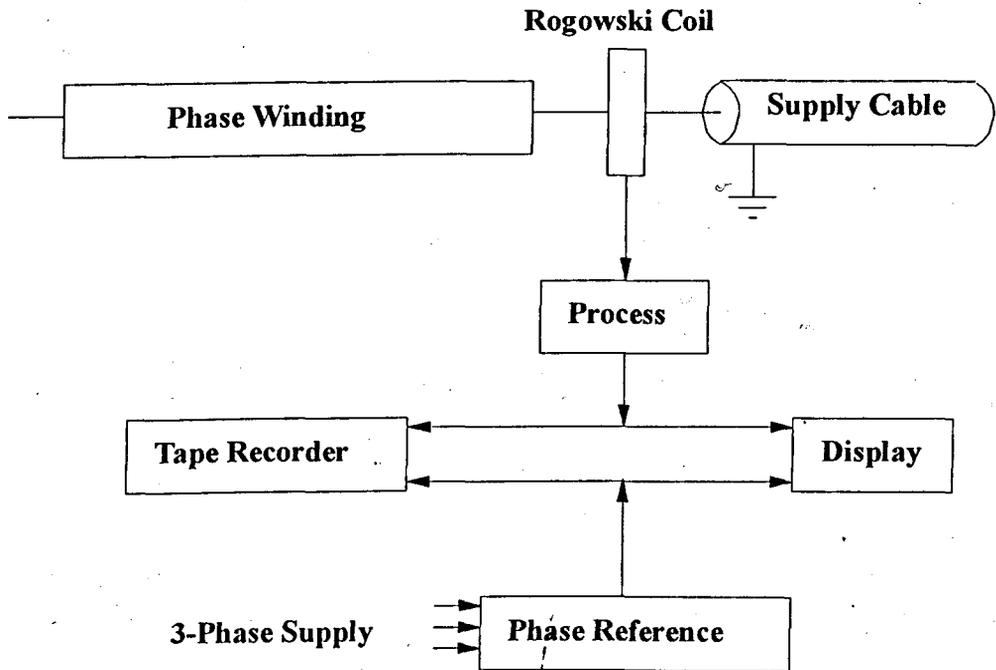


Figure 7.1 Present Monitoring Set-up

7.3 Time-Domain Processing

Time domain processing can be split into two areas: visual analysis of the digitised record and pulse processing. Each area only has a "snapshot" of all three phases, digitised for approximately 20ms (1 period of the 50Hz supply). This has the disadvantage that the digitised period may not be truly representative of the overall machine insulation condition: i.e. there may be large discharges occurring, perhaps once every 10 periods.

This drawback is inherent in any system designed for high-speed data acquisition; the high-speed sampling needed for practical frequency analysis produces a large amount of data that has to be stored and processed. However, this drawback can be minimised simply by repeatedly acquiring new data until the user judges that a true representation of the discharge profile has been obtained.

7.3.1 Visual Analysis

Visual analysis of the acquired data is presently the main weapon in the partial discharge monitoring arsenal. With it, an expert can estimate the level and severity of partial discharge activity within a given machine. In addition to this, he can inspect the discharge pattern for particular signatures that indicate previously encountered faults. Visual analysis is supported in a number of ways:

- The digitised data for all three phases can be displayed on screen at once and the entire captured record scanned. This enables the user to pick out large discharges and correlate the information from all three phases, identifying such phenomena as phase to phase discharges.
- The user can limit the display to one phase only, thus increasing amplitude resolution and, using a tracking cursor, make amplitude and duration measurements of individual captured discharges.
- The option exists to change the horizontal resolution of the display, effectively allowing more information to be squeezed onto the screen at once. This would allow the operator to quickly scan through the sample record.

These visual analysis routines effectively allow the system to be used as a specialised digital oscilloscope, albeit a very expensive one.

7.3.2 Pulse Processing

Visual analysis is supplemented with an integrated package of pulse processing routines that perform a number of analyses on existing pulses within the data record.

The current analysis options include:

- Pulse height analysis, whereby identified pulses are counted and categorised according to their amplitudes. This produces a graph which gives information on both discharge pulse amplitudes and their repetition rates within the data sampling period.

- Phase height analysis, whereby the entire data record is split into phase "windows". Each window corresponds to a number of data samples, within which the maximum positive and negative pulses are found and plotted. This results in a "map" of large discharges according to where they occur within the complete data record.
- Phase count analysis, where once again the data record is split into phase windows, but, this time the number of pulses which occur within each window is found. This produces a "map" of repetition rates for the partial discharges as the sampling progressed.
- Noise removal, whereby periodic interference superimposed on the acquired partial discharge data, commonly caused by thyristor converters, is identified and removed. If this were not done, it could lead to a faulty diagnosis, as the interference pulses are usually very much larger than partial discharges.

These pulse analysis options, when used in conjunction with visual analysis, amount to a powerful partial discharge monitoring tool. However, only time-domain applications have so far been discussed.

7.4 Frequency-Domain Processing

Frequency analysis of partial discharges potentially offers the capability of approximately locating discharge sources within the winding of a machine, by making use of the frequency dependant attenuation characteristics associated with pulse propagation in a stator winding.

Problems arise due to the nature of the signals being analysed. Partial discharges are transient events, typically lasting less than 10 μ s. This leads to complications when applying conventional frequency analysis algorithms, as these require a large number of data samples to estimate the spectral components with any precision.

After initial investigations into the practicalities of frequency analysis of partial discharges, a rudimentary partial discharge location classifier was developed. As it stands, however, this classifier remains impractical until it is developed further to overcome some outstanding problems:

- The classifier was constructed using repetitive pulses of the same amplitude. In practice, partial discharges have a diverse range of amplitudes.
- Discharges do not occur as single entities, but may be superimposed upon one another, altering the overall pulse shapes.
- Discharges characteristics will vary depending upon the type of insulation, the machine under test, as well as the type of fault occurring.
- Discharges occur along the entire length of the winding coils, not just at the specific measurement points that were used to develop the classifier.

All of these practical considerations underline the fact that this system was designed merely to investigate the possibilities, with further work building upon these findings to fully exploit their potential.

7.5 Future Development

The system developed here was primarily designed to investigate the processing possibilities for partial discharge monitoring of large electrical machines. Any subsequent system would greatly benefit from the lessons learned in this investigation and could be practically applied in an industrial situation.

7.5.1 Existing System Specification

The specification of the current partial discharge monitoring system is based upon the need to gather as much information from individual discharges as is possible, without sacrificing system practicality.

Investigations carried out on the prototype system indicated that for, laboratory-based measurements, initial amplifier stages of gain 100 was all that was required. Digitising speed was set at 10MHz due to the practical constraints of the available 256K Static RAM as well as the apparent over specification of the prototype system (i.e. System has to be capable of holding partial discharge information for 1 complete digitised 50Hz mains cycle to enable thyristor pulse interference removal).

The system hardware, shown in Figure 7.2, has three identical analogue/digital data acquisition units (one for each machine phase). The data acquisition units connect to a PC by means of a custom-designed hardware controller and PC interface card.

7.5.2 Future Hardware

The hardware for a future monitoring system will greatly benefit from the inevitable advances in technology, not only in computational processing speed, but also in the associated data acquisition electronics.

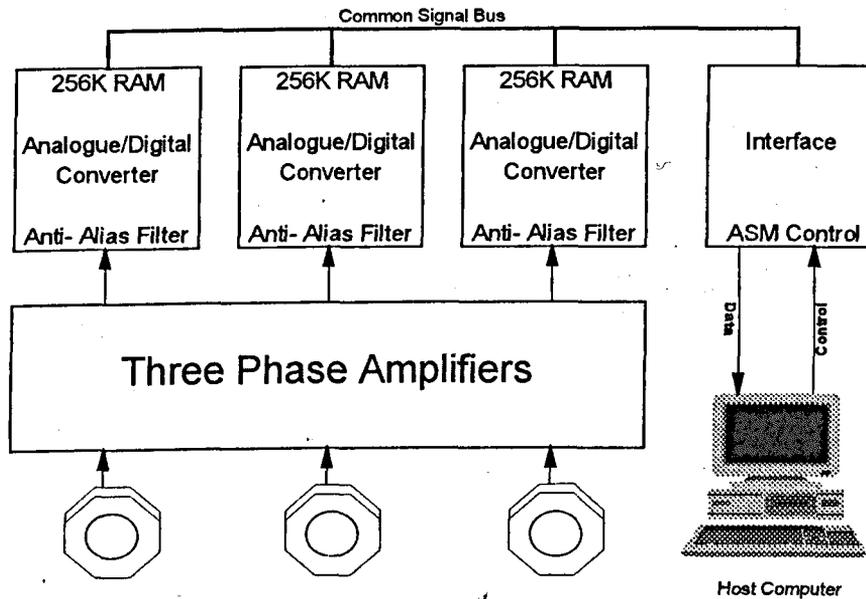


Figure 7.2 Existing System Hardware

This study has shown that a high bandwidth system is not necessarily needed, particularly if time-domain processing is all that is required. It is entirely feasible that all of the digitising electronics could be bought "off the shelf" and fitted in the expansion slots of a PC, leaving only the signal conditioning units to be designed or bought. The digitising speed of the system need not exceed 1-2MHz if pulse processing is all that is required. This can be amply illustrated by examining Figure 7.3.

Figure 7.3 shows pulses digitised at 10MHz. Obviously, there is a lot of scope to reduce the sampling rate further, whilst still acquiring the important data such as amplitude and duration.

To implement frequency-based processing, a higher bandwidth system would be required. Not necessarily as high as the prototype system: the spectral analysis plots

contain important components only up to 1.6MHz, but a digitising speed of around 5MHz would be a practical compromise. Using such sampling rates, it would still be feasible to install the data acquisition hardware in the expansion slots of a PC, greatly increasing the data transfer rate and hence processing speed. Additionally, it may even be possible to utilise the PC's expanded memory as the main data store.

The processing speeds of Personal Computers have advanced considerably since the prototype system was developed, giving future developers access to processing options that were considered impractical on the existing system. However, the additional processing necessary to implement a practical partial discharge frequency analysis algorithm would be considerable: Consider a typical partial discharge record for one phase of a machine under test: one period of the supply frequency could produce anything from 100-1000 discharges. A three phase system would therefore have to cope with up to 3000 discharges ever 20ms, too much for even the fastest '486 PCs to perform frequency analysis on in real time. A practical system would still have to capture the data and then process it later.

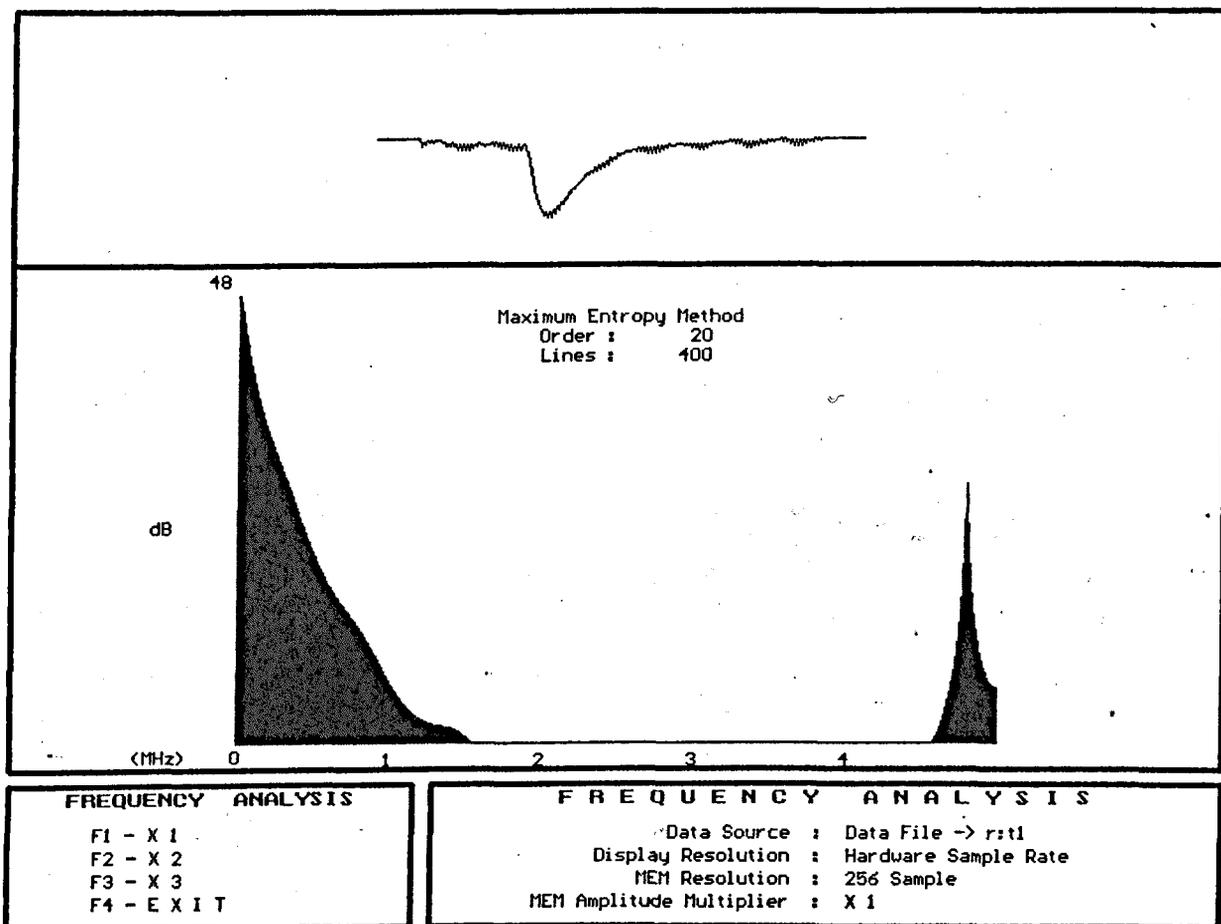


Figure 7.3 Screen Shot of Digitised Discharge

7.5.3 Future Software and Processing

There are a large amount of future processing options that can be explored; in fact, more than enough to merit another project.

The most obvious enhancement to the existing time-domain pulse processing options would be the provision of database support. This would allow subsequent sampled records to be compared with earlier versions, allowing the operator to chart the development of faults over time. The existing software routines could be enhanced to provide comparative pulse-analyses of the current monitoring record with those stored for the same machine in the database. The provision of database support could also be expanded to allow the operator to store notes along with the data records, containing information on the monitoring set-up, conclusions made, updates to the

maintenance schedule etc. In addition, automated reports could be set up to print the results of the monitoring.

Additional processing routines could be developed, once sufficient information becomes available, to recognise the different types of discharges that can occur (slot discharges, phase-to-phase discharges etc.) and then provide diagnostic information about each. Quantifying slot discharges may require some sort of pattern recognition software to be employed in order to distinguish them from other partial discharge activity. This, however, could prove difficult as each incident of slot discharge activity would be subtly different from the last, making it difficult to identify common features of slot discharges that could form the basis of identification.

Further research would be required before a frequency-based partial discharge location classifier could be practically implemented. Once perfected however, it would be a very useful addition to the suite of existing time-domain routines, in that the existing routines could be used to determine which machines require further examination in the frequency domain. These would then be examined using the frequency analysis routines, linked to the location classifier in order to ascertain where the main problems are occurring within the winding. Obviously the rules for the classifier would be different for each machine type and possibly, location. However, once the classifier had been configured for a particular type of machine it could then be used on other machines of the same type. The existing system could well be utilised to produce a more accurate classifier simply by analysing more experimental data and applying the results to the Data Solving application within a Spreadsheet program in order to calculate the relationships between the sets of figures.

A pictorial representation of the future system, using the above components, can be seen below in Figure 7.4.

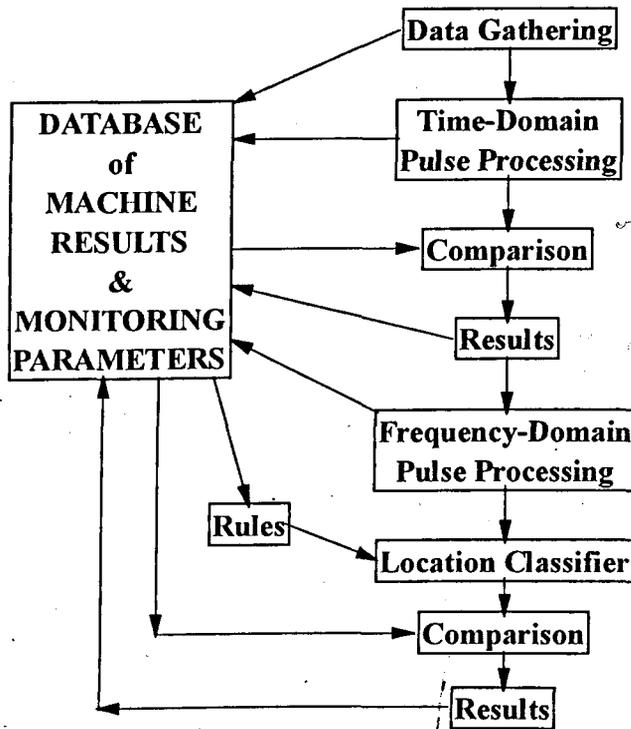


Figure 7.4 Software Components of a Future PD Monitor

Here, data is collected from a machine under test and stored as a record in the database for this machine. The data is then examined using the time-domain pulse processing techniques developed in this project, with the results stored in the database for future reference. Then, the results of the current test are compared with previous tests, stored in the database, to ascertain if any further degradation in the insulation has occurred in the time between tests. The outcome of the comparison can then be stored in the database as well as any comments or recommendations that the expert performing the tests thinks necessary.

After the pulse processing analysis, the expert can invoke the frequency analysis routines. The results of these would also be placed in the database for future reference, before invoking the location classifier to ascertain where the discharge activity is taking place. Note that the database would provide the classification rules for each test, as different machines and configurations may have an effect on the pulse shaping characteristics of their windings. The outcome of the location classifier

could then be compared with results obtained previously to see if any significant changes have occurred within the machine since it's last test, and to try and spot trends in the results in order to recommend a course of action. The results of these comparisons could then be stored in the database along with the expert's recommendations.

The provision of database support is the first step towards a fully integrated Expert System. In simple terms, once case studies of actual faults have been gathered and stored in the database, the faults can be diagnosed by a human expert and a course of action recommended in each case. This expert knowledge can then be transferred into a expert system shell where it can be classified into a set of rules that could then be applied to new untested situations. Whenever the system comes across new situations it would use its existing knowledge base (case studies) to determine a course of action or a set of conclusions. Obviously, the size and diversity of the knowledge base will be a major factor in the system successfully diagnosing new situations.

Whenever the expert system meets new situations, the results can be added to its knowledge base for future reference; in effect, the system is learning. If this is the case, some sort of checking needs to be in place to make sure that the system does not give wrong diagnoses based upon previous sets of results.

It may even be possible that some sort of heuristic system could be used to monitor a machine and, applying some expert knowledge and statistical information, refine its partial discharge locator classification rules to adapt to the particular machine under test (i.e. Given that the system has some broad rules governing frequency content of discharges as they travel through the winding, it can collect data over a period of time, categorising it according to the discharge frequency content, and, using statistical information such as the likelihood of discharges emanating from particular windings, adjust its rules according to the collected data).

Obviously this is a very simplistic view of a very large field, however, expert systems have already been successfully applied to machine condition monitoring [6] and there are no reasons why this cannot be the case here.

In order to obtain a fully integrated expert system, it would be necessary to develop additional software routines to automate the sampling and analysis of the data. This would allow repetitive sampling of the discharge information to collect data on less frequent discharge patterns.

CHAPTER 8

CONCLUSIONS

The research project described here is an addition to a well established and comprehensive programme of research, at The Robert Gordon University, related to on-line condition monitoring of large electrical machines. [1,2,3] The work commenced in 1980 and, in recent years, has been funded by research contracts with industrial organisations, including BP Petroleum Development Ltd.

The potential benefits of on-line condition monitoring are of greatest importance in high cost continuous process industries, such as oil production and electrical power generation. Techniques developed at The Robert Gordon University are already being applied in these industries, [4,5,6] in particular, to the detection of incipient fault conditions in rotors of induction motors, and in high voltage stator windings of both induction motors and synchronous generators.

The work described was directed towards non-invasive measurement and analysis of partial discharge activity within the stator of high voltage electrical machines. To that end, this investigation was primarily concerned with the development of hardware and software techniques that are capable of processing the semi-random partial discharge pulse patterns and identifying the characteristics which have most relevance to the diagnosis of insulation condition.

Electrical machine condition monitoring is a very diverse and complex field. Various techniques were outlined in Chapter 1 which are capable of monitoring specific faults as they develop. Such techniques are continually being adapted and improved as a competitive marketplace demands more accurate and reliable methods. A non-invasive technique, using Rogowski Coils was described which is used to monitor partial discharge activity within the stators of high voltage electrical machines. This technique was used as the basis for the further development of hardware and software to process the discharge pulse patterns and to investigate means with which to extract further useful information from the available signals.

Investigations were carried out into the exact requirements of a new computer-based partial discharge monitoring system. This included experiments with a laboratory stator winding to determine the necessary performance of the amplification stages, as well as setting system parameters such as ADC sample rate and subsequent memory storage capacity. These parameters, once set, allowed detailed design of the various system modules to proceed and were eventually combined into a coherent prototype.

The next stage in the prototype development involved examining the software design process necessary to operate the hardware and implement the data processing functions: An expanding menu system was described that allowed new features to be easily incorporated into its structure, as well as providing a basic user-friendly environment to interact with the operator. The prototype software modules encompassed data storage and retrieval, waveform display and rudimentary pulse processing algorithms. All of these facilities combined produced a prototype system that was capable of elementary discharge monitoring, for a single phase of a high voltage electrical machine.

The three-phase monitoring system, whilst retaining the "look and feel" of the prototype system, was considerably enhanced to incorporate the three-phase display and pulse analysis routines. A statistical noise removal scheme was described that can automatically identify and remove the unwanted thyristor switching pulses that can interfere with correct diagnoses of partial discharge activity.

Two different frequency analysis schemes were described in Chapter 6 that allowed investigations into the frequency content of partial discharges at different inception points within stator windings. The second scheme tried to address the fundamental problem faced when attempting spectral analysis of short-term transient signals such as partial discharges. This scheme, using the Maximum Entropy Method algorithm, was further utilised in conjunction with statistical analysis and a spreadsheet program to produce a rudimentary discharge location classifier. This classifier, whilst verifying the method, would be impractical for inclusion in a monitoring system at this stage; further research would undoubtedly have to be carried out into the spectral analysis of physically monitored discharges.

Finally, possible routes for future research in this field were mapped out. Recent advances in technology have made these schemes far more viable than was the case at the inception of this research project. The lessons learned, and the now proven techniques used in the development of the monitoring system have made an automatic partial discharge monitoring instrument an achievable goal.

REFERENCES

1. Thomson W T, Deans N D, Leonard R A, Milne A J; 'Monitoring Strategy for Discriminating Between Different types of Rotor Cage Faults', 18th University Power Engineering Conference, University of Surrey, April 1983, pp 241-246.
2. Cameron J R, Thomson W T, Dow A B; 'Vibration and Current Monitoring for Detecting Air-Gap Eccentricity in Large Induction Motors', IEE Proceedings, Vol 133, Part B, No 3, May 1986.
3. Buchan J, Edwards D G; 'A Technique for the Location of Partial Discharge in the Stator Winding of a High Voltage Motor', 5th BEAMA International Electrical Insulation Conference, Brighton, May 1986.
4. Thomson W T, Chalmers S J, Rankin D; 'On-Line Current Monitoring Fault Diagnosis in HV Induction Motors - Case Histories and Cost Savings in Offshore Installations', Offshore Europe Conference, Aberdeen, September 1987.
5. Edwards D G; 'Condition Monitoring of Electrical Insulation in High Voltage Motors', 2nd International Conference on Condition Monitoring, Condition Monitoring 1987, University College Swansea, March 1987.
6. Leith D; 'Expert Systems in A.C. Induction Motor Fault Diagnosis', PhD Thesis (CNA), The Robert Gordon Institute of Technology, Aberdeen, 1990
7. Tavner P J, Gaydon B G, Ward D M; 'Monitoring Generators and Large Motors', IEE Proceedings, Vol 133, Part B, No 3, May 1986, pp 169-180.

8. Simons, J S; 'Diagnostic Testing of High-Voltage Machine Insulation', IEE Proceedings, Vol 127, Part B, No 3, May 1980, pp 139-154.
9. Pinto C; 'The Discharge Erosion Problem in HV Machines', Electrical India, 15th November 1987, pp 5-11.
10. Kurtz M, Stone G C; 'Diagnostic Testing of Generator Insulation Without Service Interruption', CIGRE, Paris, 1980, Paper 11-09
11. Wright M T, Yang S J, McLeay K; 'General Theory of Fast-Fronted Interturn Voltage Distribution in Electrical Machine Windings', IEE Proc, Vol 130, Pt B, No 4, July 1983.
12. Wilson A, Jackson M A, Wang N; 'Discharge Detection Techniques for Stator Windings', IEE Proc, Vol 132, Pt B, No 5, September 1985, pp 234-244.
13. Thomson W T, Edwards D G, Jayram N D, Chalmers S J; 'Development of On-Line Condition Monitoring Techniques for High Voltage Machines', BP Contract 5680/Dyce, Report 2, 1986.
14. Edwards D G, Chalmers S J; 'Development of On-Line Condition Monitoring Techniques for High Voltage Machines', BP Contract 5860/Dyce, Report 3, 1987.
15. Austin J, James R E; 'On-Line Digital Computer System for Measurement of Partial Discharges in Insulation Structures', IEEE Trans. Electr. Insul. EI-11, 4, 129-139, 1976.
16. Tanaka T, Okamoto T; 'A Minicomputer Partial Discharge Measurement System', IEEE International Symposium on Electrical Insulation, 86-89, 1978.

17. Tanaka T, Okamoto T; 'Microcomputer Application to in-situ Diagnosis of XPLE Cables in Service', Int'l Conf. Prpr. Appl. Dielectric Mat., 499-502, 1985.
18. Okamoto T Tanaka T, 'Novel Partial Discharge Measurement -Computer-Aided Measurement Systems', IEEE Trans. Electr. Insul., EI-21, 6, 1015-1019, 1986.
19. Kay S M, Marple S L, "Spectral Analysis - A Modern Perspective", Proc IEEE, Vol 69, No. 11, November 1981.
20. Press W H, Flannery B P, Teukolski S A, Vetterling W T, "Numerical Recipes in C, Cambridge University Press, 1988
21. Mackoul J, 'Linear Prediction: A Tutorial Review', IEEE Proc, Vol 63, No. 4, 561-580, 1975
22. Linkens D A, 'Non-Fourier Methods of Spectral Analysis', In Digital Signal Processing, Jones N B, ED. Peter Peregrinus, pp213-254
23. Yule G U, "On a Method of Investigating Periodicities in Disturbed Series with Special Reference to Wolfer's Sunspot Numbers", Philisophical Trans. Royal Society, London, series 4, Vol 226, pp267-298, July 1927.
24. Walker G, "On Periodicity in Series of Related Terms", Proc Royal Society, London, Series A, Vol 131, pp518-532, 1931
25. Burg J P, "Maximum Entropy Spectral Analysis", 37th Annual International Meeting, Society of Exploratory Geophysicists, Oklahoma City, Oct 1967.
26. Parzen E, "Statistical Spectral Analysis (Single Channel Case)", Dept Statistics, Stanford University Report 11, 1968

27. Chalmers S J, Jamieson J W; 'Development of a Partial Discharge Monitoring System for HV Electrical Machines', 24th University Power Engineering Conference, Queens University Belfast, September 1989.

DEVELOPMENT OF A PARTIAL DISCHARGE MONITORING SYSTEM FOR HV ELECTRICAL MACHINES

S J Chalmers, J Jamieson

Robert Gordon's Institute of Technology

INTRODUCTION

In recent years the subject of condition monitoring of many items of plant and machinery has taken on considerable significance. In particular, methods of on-line condition monitoring which allows a process to continue undisturbed have received much interest.

Techniques developed at RGIT have concentrated on non-invasive on-line monitoring of the condition of electrical induction motors. Methods of detecting fault conditions in rotors and in stator insulation have been devised.

The work described in this report is directed towards the measurement of partial discharge activity within the stator of high voltage machines.

All high-voltage motors operating at voltages of 6.6kv or greater generate partial discharges within their stator winding. These discharges occur external to, or within small voids in the coil insulation structure, and are the result of high electric stress in the winding. The pulses produced are of relatively low energy (10^{-9} to 10^{-6} coulombs) and of very short duration (100's of nanoseconds). Under normal operating conditions the energy associated with the partial discharges is small. However, degradation of the insulation results in an increase in the number of partial discharges, and their associated energies. Periodical monitoring of partial discharge activity is therefore a potentially useful technique.

The basic transducer used to obtain information on partial discharge activity should ideally be completely non-invasive. The device which is used for this purpose is the Rogowski coil. This coil is essentially an air cored current transformer designed to have special characteristics. This results in a transducer which is optimised to pick up the high frequency partial discharge pulses but reject "interference" signals such as the fundamental mains frequency.

Previous work (1) has demonstrated the possibility of detecting and measuring partial discharges using a Rogowski coil. This paper details research into a computer-based system for the measurement and subsequent processing of these signals, the ultimate aim being the development of a dedicated instrument capable of fulfilling this function.

DESIGN REQUIREMENTS

The concept of a computer-based instrumentation system for the detection and processing of partial discharges is by no means a new one. Austin and James (2) demonstrated

the feasibility of such a system limited only by the available technology. Tanaka and Okamoto (3), (4), (5) with the use of superior technology considerably advanced the capabilities of the computer-based measurement system. However, all of the previous efforts have invariably relied on some sort of pulse stretching or detection circuitry in order to capture the discharges. The concept of directly digitising the partial discharge signals is a relatively new one. Couple this with the processing power of the new generation of microcomputers, and the result should be a system that is truly superior to any of its analogue counterparts.

The design of such a system, however, was not straightforward as there were a number of technical, as well as operational constraints to be taken into consideration.

Technical Constraints

The non-invasive method of using remote Rogowski Coil detection of partial discharges results in signals of very small amplitude (fractions of a millivolt to a few millivolts), which in turn places stringent design requirements upon the initial analogue processing stages, specifically the need for adequate amplification, whilst maximising bandwidth, and limiting noise.

The digitising process also has far reaching design implications associated with it. Most importantly, the speed and resolution of the Analogue to Digital Converter had to be a compromise between the bandwidth of the amplifier, the specification of the anti-alias filter, and the subsequent storage needs for the digitised discharge pulses.

Operational Considerations

Ultimately the monitoring system is destined to operate in an industrial environment, specifically on offshore installations, therefore the relevant working practices and conditions had to be taken into account during the design process. Special consideration had to be given to the effect that electrical interference would have upon the system. The most serious source of electrical noise on an offshore installation is a result of the switching performed by thyristor speed controllers of DC motors. These drives produce periodic pulses of large amplitude that are picked up by the Rogowski Coil and completely swamp any partial discharge activity for a short time during the 50Hz cycle. Unfortunately these pulses cannot be removed simply by filtering, or Automatic Gain Control circuitry, as they have a similar frequency content to the

partial discharges, and can occur at any point in the mains cycle. However, once the signal has been digitised, a number of possible removal techniques become practical. These will be discussed in a later section.

Having identified the important design constraints the system can be described in more detail.

SYSTEM OVERVIEW

The partial discharge monitoring system can be divided into 4 main modules:

- (i) The analogue processing stages incorporating a high-gain, low-noise amplifier and anti-alias filter.
- (ii) The digitising module, comprising of an 8-bit "flash" analogue to digital converter with associated circuitry.
- (iii) The primary data storage module and hardware controller, which serves as a temporary store for the digitised information, and co-ordinates the digitising process.
- (iv) A microcomputer that acts as the system controller and information processor, handles operator interaction, and displays the results.

A block diagram of the system can be seen in Figure 1.

The entire system was designed very much with future expansion in mind, which meant that the most obvious solution to a particular problem was not always adopted.

The next sections discuss each system module in more detail.

Analogue Signal Processing System

In order to appreciate the problems associated with the development of an adequate front-end for the instrument one has first to investigate the actual signals being applied to the system.

The amplitude and shape of partial discharge pulses varies both with the size and point of origin of the discharges within the stator winding. Obviously in order to extract the maximum amount of information from such pulses a high bandwidth system is needed, with adequate dynamic range and noise performance. Previous experiments using the Rogowski Coil technique (1) have shown that the upper frequency limit need not be as high as initially expected. This led to the amplification upper frequency limit being fixed at 4MHz which seemed a reasonable compromise between frequency resolution and subsequent design difficulty.

In determining the necessary gain of the amplifier, laboratory experiments were carried out on a 6.6KV stator winding setup as shown in Figure 2. Artificial discharges of known magnitude were applied to the line-end of the stator by means of a voltage generator E in series with a capacitor C, with the termination network constructed to mimic the surge impedance of the motor supply cable.

The resulting signals output from the Rogowski Coil were amplified and examined. The conclusions of the experiments indicated that an amplifier gain in excess of 500 was needed to ensure that sub 1000pC discharges could be detected.

Combine this requirement with the need for good noise rejection and the result is a high specification amplifier - 2GHz Gain-Bandwidth Product, high slew rate and Signal to Noise Ratio > 50dB.

The specification of the amplifier effectively precluded the use of operational amplifiers in the design, so various configurations of FETs and bipolar transistors were tried until a satisfactory combination, an adaptation of an ultrasonic amplifier, was found.

To prevent aliasing occurring in the digitised signals active low pass filters needed to be employed to limit the signal frequencies to half the sampling rate (Nyquist sampling criteria). These filter networks had to disrupt the desired signals as little as possible and therefore needed to have linear gain and phase characteristics in the pass-band. The class of filters chosen to accomplish the anti-aliasing task are Thompson-Butterworth or Compromise filters, so called because they combine the gain characteristic of a Classical Butterworth filter with the phase response of a Bessel filter.

Once the unwanted higher frequencies of the partial discharge pulses have been removed the signal is ready for application to the digitising stage.

The Analogue to Digital Converter Module

The specification of the analogue to digital converter affects the entire system as the digitising speed dictates the necessary response and size of the ensuing data storage module, as well as the bandwidth of the preceding amplifier and the response of the anti-alias filter. The resolution of the converter again determines the size of the data store, as well as the required noise characteristics and gain of the amplification stages. These decisions regarding the analogue to digital converter have important ramifications on both the cost and overall performance of the final system, so initially a balance was reached between cost, speed and design difficulty, so ensuring that the prototype had adequate performance, with the option of upgrading at a later stage.

The converter chosen is a 8-bit "flash" converter operating at 15MHz. The resolution of the analogue to digital converter, coupled with the level of amplification governs the precision of the discharge measurement process.

ie. If the system were to be calibrated at full scale for a discharge of 255000pC then the measurements would rise in steps of 1000pC from 0 - 255000pC. The speed of conversion places certain demands upon the data storage module both in response and size, because as the conversion process gets faster the total memory required to hold the data increases, and has to have a correspondingly smaller access time.

Hardware Controller and Data Storage Modules

The roles of the hardware controller and primary storage modules are to oversee and control the sampling process and provide a temporary storage facility for the digitised signal. The initial design problems revolved around choosing the size of the storage facility. This had to be a trade off between cost, speed and the amount of data that it was required to hold. The main criterion that this decision was based upon concerned the need to eliminate thyristor switching noise from the digitised signal. This noise is periodic in nature but its exact location in the 50Hz cycle will vary slowly with time, therefore in order to uniquely identify its position within the cycle a number of consecutive 50Hz cycles need to be stored. If one assumes that 3 cycles is the minimum amount needed to unambiguously ascertain the position of, and thus, eliminate the thyristor noise it is a simple matter to estimate the size of memory needed to hold the samples:-

period of 3 x 50Hz waveforms = 0.06 sec

$$\begin{aligned} \text{No of samples in 3 cycles} &= \text{time} \times \text{sample rate} \\ &= 0.06 \times 15 \times 10^6 \\ &= 900000 \text{ samples} \end{aligned}$$

This meant that a large amount of Static RAM had to be employed in order to cope with the amount of data being produced from the ADC. Data compression methods were considered in order to reduce the memory overhead, however a minimum figure could not be arrived at for the success of any scheme, as it would be dependant on the partial discharge activity. Also, the control circuitry necessary to implement the data compression scheme would need to be complex in order to work at a maximum of 15 Mbytes/sec.

The method chosen was therefore to build a bank of 1 Mbyte static RAM to act as a temporary sample store before transferring the data to the host computer for processing. The design of the digital control circuitry was greatly helped by use of Computer Aided Design and simulation packages, and the chip count kept as low as possible by implementing the discrete logic functions in programmable gate array devices.

The Microcomputer System

The host system had to accommodate a number of features:

- (i) It had to act as a software development system.
- (ii) It had to have considerable processing power to allow implementation of very complex digital signal processing functions in a reasonable time.
- (iii) It had to have a large memory capacity to cope with the amount of digitised data.
- (iv) It had to be relatively inexpensive, so ensuring a cost-effective system.

The microcomputer chosen was an IBM (AT) compatible machine with a 32-bit 80386 micro-processor, running at 20 MHz. It has 4 Mbytes of on-board RAM and as VGA graphics capability. IBM compatibility ensures that there are a large number of software tools available, as well as hardware accessories. The 32-bit central processor will allow it to perform complex calculations with relative ease, whilst the 4 Mbytes of RAM are more than sufficient to accommodate the sampled data, and any processing programs.

So far only the hardware of the condition monitoring system has been discussed, however the success or failure of the system rests mainly with the software that processes the discharge pulses. The next section will examine some processing schemes and the proposed development of the monitoring system software.

DISCUSSION OF DIGITAL SIGNAL PROCESSING SCHEMES

The first problem to overcome is that of the offshore DC drives superimposing thyristor switching noise on the partial discharge signals. Considering the periodic nature of this interference, several possibilities for its elimination are possible, the simplest of which is subtraction of successive 50Hz cycles, whereby the periodic interference will be eliminated and the pseudo-random partial discharges should remain unaffected. However, if discharges occur at the same points in successive cycles, these will also be removed, or more importantly, their shapes and amplitudes altered. Another scheme is to first locate the positions of the interference in the cycle, and then to window out that section of the cycle as the noise will have completely swamped any partial discharges in that section.

Having eliminated any contaminating signals, the next task is to process the discharge pulses. Here again a number of options present themselves. Initially pulse height analysis may be adopted whereby the discharges are counted and categorised according to their amplitudes, in order to give a quantitative indication of the health of the stator insulation. If this simple technique was coupled with a facility whereby the previous history of the machine was also known, then some degree of predictive analysis may be employed. An enhancement to this scheme may involve determining the frequency content of each pulse, in order to determine its point of origin in the winding and perhaps predict its original amplitude. It may even be possible to categorise the type of discharges.

The interaction with the user is another area that must be dealt with. Initially the systems objective is to develop the appropriate DSP algorithms and demonstrate their feasibility. But ultimately the aim is to produce a monitoring instrument, and therefore consideration has to be given to the user interaction. The graphics capabilities of the computer are such as to allow detailed displays should they become necessary, and it would be a relatively simple task to incorporate an elementary expert system or database management program into any software produced.

CONCLUSIONS

This paper discusses the design of a system for the reliable detection and processing of partial discharge signals. It makes clear the technical limitations, and practical problems associated with such a system. Following a general design specification, the design and implementation of the individual sections are discussed in detail. At the time of writing the hardware has been built, and it undergoing testing and integration, into a complete system. The report concludes by discussing the next stage of the project - software development, by setting out the direction in which this will proceed.

REFERENCES

1. Edwards D G, 1987, "Condition Monitoring of Electrical Information in High-Voltage Machines", Proceedings Condition Monitoring 87, Swansea, 788-797.

2. Austin, J., and James, R E., 1976, "On-line Digital Computer System for Measurement of Partial Discharges in Insulation Structures", IEEE Trans. Electr. Insul. EI-11, 4, 129-139.

3. Tanaka, T., and Okamoto, T., 1978, "A Minicomputer Partial Discharge Measurement System" IEEE International Symposium on Electrical Insulation, 86-89.

4. Tanaka, T., and Okamoto, T., 1985, "Microcomputer Application to in-Situ Diagnosis of XPLE Cables in Service", Int'L Conf. Prop. Appl. Dielectric Mat., 499-502.

5. Okamoto, T., and Tanaka, T., 1986, "Novel Partial Discharge Measurement - Computer-Aided Measurement Systems", IEEE Trans. Electr. Insul., EI-21, 6, 1015-1019.

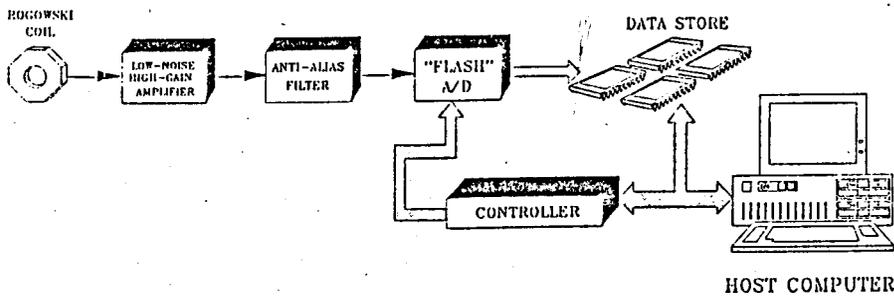


Figure 1. System Overview

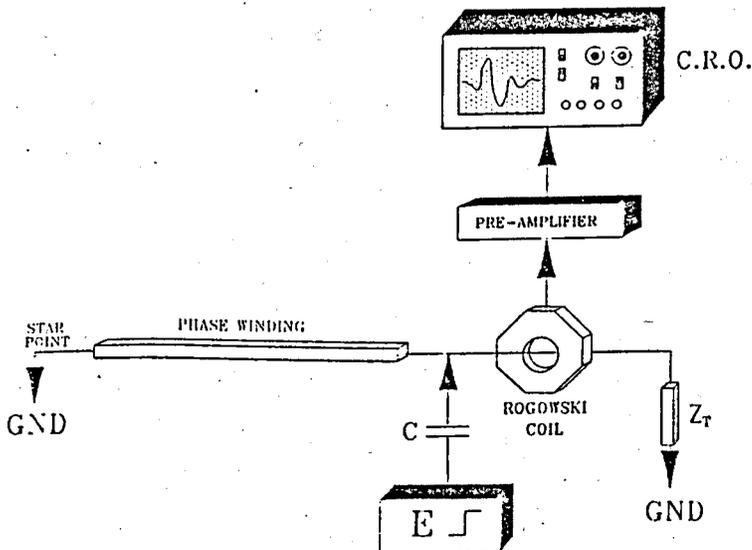


Figure 2. Experimental Setup