Fault Tolerant Drives for Safety Critical Applications

James Alexander Haylock

A thesis submitted for the degree of
Doctor of Philosophy

Department of Electrical Engineering
University of Newcastle upon Tyne

Abstract

The correct operation of adjustable speed drives, which form part of a larger system, is often essential to the operation of the system as a whole. In certain applications the failure of such a drive could result in a threat to human safety and these applications are termed 'safety critical'. The chance of a component failure resulting in non-operation of the drive can be dramatically reduced by adopting a fault tolerant design. A fault tolerant drive must continue to operate throughout the occurrence of any single point failure without undue disturbance to the power output. Thereafter the drive must be capable of producing rated output indefinitely in the presence of the fault.

The work presented in this thesis shows that fault tolerance can be achieved without severe penalties in terms of cost or power to mass ratio. The design of a novel permanent magnet drive is presented and a 'proof of concept' demonstrator has been built, based on a 20 kW, 13000 RPM aircraft fuel pump specification. A novel current controller with near optimal transient performance is developed to enable precise shaping of the phase currents at high shaft speeds. The best operating regime for the machine is investigated to optimise the power to mass ratio of the drive.

A list of the most likely electrical faults is considered. Some faults result in large fault currents and require rapid detection to prevent fault propagation. Several novel fault sensors are discussed. Fault detection and identification schemes are developed, including new schemes for rapid detection of turn to turn faults and power device short circuit faults. Post fault control schemes are described which enable the drive to continue to operate indefinitely in the presence of each fault. Finally, results show the initially healthy drive operating up to, through and beyond the introduction of each of the most serious faults.

Contents

List of Figures

List of Tables

Acknowledgements

A project of this nature requires input from many individuals and organisations, too many to fully credit in this space. The author would express his gratitude to EPSRC, as without their funding this work could not have been undertaken. The input of the industrial partners, LucasVarity, Rolls Royce IRD, Welwyn Micro Electronics and Ultra Electronics, has been invaluable and the individual contributions of Jeff Coles, Chris Weir, Jerry Watson and Nigel Collins have been particularly appreciated.

Thanks are due to my supervisors, Barrie Mecrow, Dave Atkinson and Alan Jack, for all their efforts, from their initial investigations and EPSRC grant application through to the guidance, help and encouragement provided over the three years during which I worked with them and for much more besides. Without them I could not have achieved so much.

Special thanks are due to Barrie for agreeing to act as my PhD supervisor. For his patience in reading through each section of my thesis and providing useful feedback, irrespective of the quality of the initial drafts, and for the advice, support and encouragement throughout the period of time when I slowly turned a collection of ideas and results into a finished thesis.

All the technicians in the department deserve a mention, especially John Noble for building the machine and for helping to commission the test rig, and Colin Scott for his work on the electronics. Thanks are also due to numerous small companies for helping out with the difficult and unusual jobs that we asked them to do, especially Gee Bee, Balancing Services, and the staff in the Physics Department workshop for their assistance in removing the remains of a bearing from our machine shaft after its untimely demise.

The UG lab has been a lively, motivating and, most importantly, friendly place to work, so thanks also to Adil, Andy, Bernhard, Chris French, Chris Maddison, Christian, Gavin, Hasan, Howard, Ibrahim, Jamie, Jowad, Ken, Oystien, Pete, Phil, Steve, Volker, Wanda, for making the lab an enjoyable experience, even when the work was not going too well.

Now there is just enough space left for a special thank you to my parents, who have helped in many ways, not least for their moral and financial support whilst I've been writing up, and especially to my father for his unsolicited offer to proof read this thesis and for his subsequent dedication to the task.

Symbols

Wherever practicable the meaning of each symbol used in this thesis will be given when the symbol is first used. The symbols in the following list are widely used throughout this work. Note that the use of italic print in certain circumstances is not intended to imply any difference in meaning. Unless stated otherwise, the symbols in the list below represent instantaneous values.

B	Flux density
e	Instantaneous phase back-EMF
E_f	Fundamental component of phase back-EMF
F	Percentage by which the measured change in current overshoots the expected current
i	Current
i^*	Current reference (demanded value of current)
I_a	Fundamental component of phase current
$\hat{I}_a$	Peak value of phase current
$i_{CtrlError}$	Maximum current error which could result from controller deficiencies
i_{dc}	Current supplied by the DC link
i_{error}	Current error
i_{error_RMS}	RMS value of current error
I_{FWCS}	Current seen by the freewheel current sensor
i_k	Current at sampling instant k
i_k^*,	Current reference at sampling instant k
$\|i\|_k^*$,	Current magnitude demand at sampling instant k
$i_{MeasError}$	Maximum current measurement error
$i_{Measured}$	Measured current
I_{PHASE}	Current seen by the phase current sensor
$\Delta\hat{i}_{Pulse}$	Peak value of a current pulse
Ki	Integral gain in a PI controller
Kp	Proportional gain in a PI controller
L	Inductance
L_{phase}	Phase self inductance
L_{turn}	Self inductance of a single turn in a phase winding

n	Number of turns per phase
P_{gap}	Air gap power
Q_m	Power converter volt ampere rating
R	Resistance
$R(t)$	Reliability
R_{ON}	Resistance used in a gate drive circuit to control IGBT turn on speed
R_{OFF}	Resistance used in a gate drive circuit to control IGBT turn off speed
R_{phase}	Phase resistance
R_{turn}	Resistance of a single turn in a phase winding
t	Time
Δt	Time period over which current error accumulates
t_{error}	Error in the time for which voltage is applied to the phase to drive current
t_k	Time at sampling instant k
t_{k+2}	Time at sampling instant k+2 ($t_{k+2} = t_k+2T$)
T	The sampling period
Tc	Computation time of the controller
Td_k	Dead time corresponding to the current i_k
Te	Sampling period (Te=T)
T_e	Instantaneous electromagnetic torque
v	Instantaneous voltage
V	Fundamental component of voltage
V_{LINK}	DC link voltage
V_{CE}	Collector-emitter voltage across an IGBT
V_{CESAT}	Collector-emitter voltage across an IGBT, when the device is turned on
V_{error}	Error in DC link voltage measurement
Vf_k	Forward voltage drop across a power device corresponding to the current i_k
W'	Co-energy
X_q	q-axis reactance
X_d	d-axis reactance
X_s	Synchronous reactance
z^{-1}	A delay of T in the z domain
$z(t)$	Failure rate
δ	Load angle, angle between the phase back-EMF and the phase voltage
ϕ	Phase angle, angle between the phase current and the phase voltage

γ	Current angle, angle between the phase back-EMF and the phase current
λ	Time averaged failure rate
θ	Rotor angle (radians, electrical)
θ_k	Rotor angle at sampling instant k (radians, electrical)
τ_0	Sum of the small time constants in a control system due to the process itself, the feedback transducers and the delay of the ADC
τ_S	Equivalent small time constant of a discrete time controller
ω	Rotor electrical frequency (actual)
ω_c	Calculated rotor electrical frequency
ψ	Flux linkage
ψ_f	Flux linkage resulting from magnet flux
ψ_a	Flux linkage resulting from phase current
$1\rightarrow$	Used on oscillograms to indicate the zero level of trace 1
$2\rightarrow$	Used on oscillograms to indicate the zero level of trace 2

Chapter 1

INTRODUCTION

Widespread use of adjustable speed drives has been made possible through the development of power electronic devices and increasingly sophisticated electronic controls. Unfortunately the increased complexity of adjustable speed drives, as compared to the systems which they replace, gives rise to fears about reliability in certain applications. Adjustable speed drives would give benefits in a wide range of aerospace applications but the safety critical nature of these applications demands extremely high mean time between failures. The mean time between failures for a system can be dramatically improved by adopting a fault tolerant design. This begs the question: 'Can an adjustable speed drive be designed to be fault tolerant without severely compromising the design?'.

The discipline of condition monitoring attempts to predict the onset of faults to enable maintenance to be carried out before drive failure. If the condition monitoring function could be incorporated into the drive electronics then the resulting drive could self monitor, at least until the point at which it reported its own impending failure. This could be termed a 'self monitoring drive'. However some faults are very serious and will propagate quickly, resulting in the immediate failure of the drive. Thus condition monitoring alone is not sufficient to achieve the goal of fault tolerance but it will be shown that 'self monitoring' is a requirement in a fault tolerant system.

If a drive were able to survive any likely fault and continue operating with a minimal reduction in performance then it could be termed a 'fault tolerant drive'. Historically the subject of fault tolerant drives has received little attention but recently the level of interest has increased substantially to a point where a number of publications have appeared. It is becoming apparent that the development of fault tolerant drives might be achievable rather than just desirable. Whilst it is probably true to say that the subject of fault tolerance still isn't generally recognised as an area in its own right in the field of

electric drives, self monitoring and fault tolerant drives are emerging as new topics for research.

The principle objective of this thesis is to investigate the possibility of designing a drive capable of surviving any single point failure, detecting the failure and continuing to operate. Every effort will be made to achieve this goal with the minimum penalty in terms of increased complexity, parts count, cost, weight and size. To develop a more detailed set of aims it will be necessary to examine exactly what is meant by fault tolerance and how it can be achieved, and to choose a suitable application area on which to focus the work. These issues will be addressed later in the chapter.

1.1 Thesis Overview

This thesis is divided into four parts. The introduction gives important background information, explaining the aims and objectives and examining what is meant by fault tolerance. Chapters 2 to 4 describe the drive design, control and the resulting fault-free performance. The end result of this second section is a high performance drive which provides some degree of inherent fault tolerance without unduly compromising the fault-free performance. Chapters 5 to 9 deal with each of the elements of fault detection and post fault operation. The result of the work presented in these chapters is a drive which is capable of continued operation in the presence of any one of the most likely failures. As might be expected, the final chapter highlights the major achievements and shortcomings of the research and examines the extent to which the aims have been fulfilled.

1.2 What's New

This section outlines areas of work presented here, which the author believes have not previously been published by others.

- Review of the emerging field of fault tolerance in electric drives (Chapter 1).
- Testing of the first permanent magnet (PM) machine designed for fault tolerance.
- Application of a novel current control technique, giving near optimal transient performance (Chapter 3).

- The use of a ψ-i diagram to design a current profile to minimise the power converter kVA/kW requirement (Chapter 4).
- Examination of new fault sensors (Chapter 6).
- Development and hybridisation of an inverter fault current sensor (Chapter 6).
- A technique for using the control electronics to detect shorted turns in the machine (Chapter 7).
- Application of on-state voltage sensing to the detection of faults within the power electronics (Chapter 8).
- Experimental verification of the capability of the fault tolerant demonstrator drive to detect all the most likely faults in the power converter and in the machine (Chapter 7 and Chapter 8).
- Experimental verification of the capability of the demonstrator drive to continue to operate throughout the introduction of any of the most likely faults with appropriate post-fault control (Chapter 9).

1.3 Terminology

1.3.1 Electric Drives

The field of electric machines and drives is changing very quickly, and unfortunately the terminology is not universally standard. Within this thesis the term 'drive' is used to refer a complete electric drive including the machine as shown in Figure 1. The term power converter is used to refer to a set of power electronics used to convert electrical power from one form to another. The term control electronics refers to a set of signal

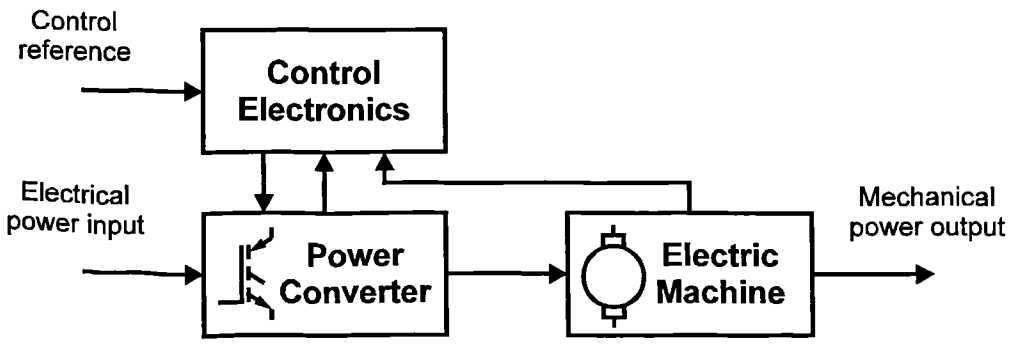

Figure 1 : Components of an Electric Drive.

electronics used to control and monitor the operation of a power converter and an electric machine. Note that this naming convention deviates a little from the norm, as the IEC definition for a 'power electronic converter' includes the signal electronics. This deviation is required to permit separate discussion of the power converter and control electronics when discussing fault tolerance.

1.3.2 Reliability and Availability

In applications which are safety critical or where a very high reliability is required for other reasons, techniques have been developed to formalise the assessment of system dependability. One such technique is to perform an ARMS (or a RAMS) study [1]. This attempts to assess the four elements; availability, reliability, maintainability and safety. Some of the definitions used by Caplen [2] are relevant as follows:

Failure	A failure refers to a deviation from the component or system specification.
Failure rate	The instantaneous failure rate, z(t), is defined as: $$z(t) = \frac{\text{Number failing per hour at instant t}}{\text{Number surviving at instant t}}$$ Note that for electronic components, z(t) is taken as a constant, equal to the time averaged failure rate, λ.
Failures in 10^9 hours (FITs)	The units often used to measure failure rate.
Reliability	Expresses how prone an item is to failure. Reliability is the probability that an item will still be working after operating for a given time. Reliability is a function of time, so when asking about the reliability, it is necessary to know the time frame of interest. For a batch of components on test, component reliability, R(t), may be expressed as:

$$R(t) = \text{Reliability} = \frac{\text{Number surviving at time t}}{\text{Total number at time t} = 0}$$

Mean Time Between Failures (MTBF) The average time for which a unit will operate before it fails. MTBF is sometimes given the symbol θ and can be calculated from $\theta = 1/\lambda$. MTBF is time invariant which makes it a more convenient figure than reliability. Thus MTBF is commonly used in requirement specifications.

Mean Time To Repair (MTTR) The average time from a fault occurring in a unit to the repair of the unit. In some applications this time will include the time taken to return the unit to a servicing point.

Availability The proportion of time for which the unit is operational. This is expressed as a percentage or a fraction of unity and is calculated from:

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

Clarke and Cornwell [3] use the term MTBF for system failures and the term fault rate (synonymous with failure rate) for component failures in there paper on fault tolerant power converter systems.

1.4 Why Fault Tolerance?

There are at least two good reasons for the current increasing interest in fault tolerant drives. The first is possibly the reason for the original industrial support for the work reported here. To use the words of Kastha and Bose [4]: 'A common fear that power electronics equipment is not sufficiently reliable is preventing its widespread application....'. It is true that in the past power electronic apparatus has developed a reputation for poor reliability. Whether or not new products continue to suffer from poor reliability remains to be seen.

The second and arguably more important reason for looking at fault tolerance is not limited to the field of power electronics or drives. This is the present trend towards greater dependence on increasingly complex systems. In many applications the failure of a system could compromise human safety, or result in a loss of production time, work in progress or computer data. If human safety is dependent on the operation of the system then it is described as 'safety critical' [5]. Sometimes it is sufficient to make a safety critical system 'fail safe' but in situations where failure of a system to operate will be either very expensive or life threatening then fault tolerance should be the goal.

1.5 What is Fault Tolerance?

Much can be learnt from a paper by White [6] on fault tolerant power converters. This paper provides an excellent introduction to the subject of fault tolerance and is especially relevant here. Coming from a telecommunications background, White is concerned about converting mains (utility power) into a high availability DC power supply (to feed electrical equipment in a telephone exchange). In the broad sense an electric drive is also a power converter, converting electrical power into mechanical power.

While the term 'fault tolerance' seems fairly self explanatory, it is actually taken to mean different things by different people. Writing about fault tolerance in power systems, White gives a definition of fault tolerance which is adopted here:

'no single failure will cause the system to malfunction'

The term 'fault tolerance' is erroneously used to refer to a continuous spectrum of possibilities ranging from high availability systems through to strongly fault tolerant systems. Levels of system availability can be improved by a design which permits continued operation in the presence of some faults. A system capable of delivering full performance, even in the presence of multiple faults, would be termed 'strongly fault tolerant'. The spectrum of possibilities is best understood by reference to Figure 2.

It is important to note that no system can be made completely tolerant of all possible fault combinations. There is always some combination of events and failures that will lead to the disruption of a system. This is true because it is difficult to either predict the consequences of a multi-point failure or to provide the necessary level of redundancy required to continue operation with more than a small part of the system damaged.

ORIGINAL IN COLOUR

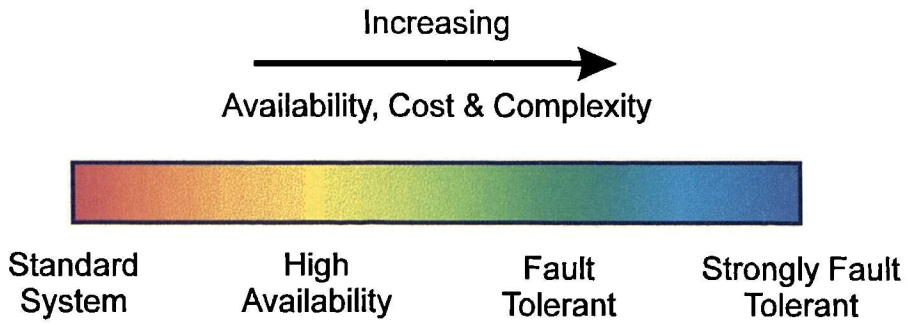

Figure 2 : The Spectrum of Availability

1.5.1 The Cost of Fault Tolerance

It may appear that the requirement for additional parts makes fault tolerant systems inappropriate for most applications on cost grounds. It is true that to continue to deliver full performance in the presence of any fault requires redundancy and that the resulting system will be correspondingly more expensive. However, simply partitioning a system does not necessarily add to the cost, and can allow a non-redundant system to continue to operate in the presence of a fault - albeit with reduced performance.

Kastha and Bose examine the possibilities for operating a faulted induction motor drive without significantly adding to the system cost. This work is obviously of considerable interest, since induction machines are rugged, low cost and widely used in adjustable speed drives. However the constraints of using a standard power converter and machine topology mean that power device short circuit failures or failures within the machine are serious problems which have not been adequately addressed, [7], [8].

In larger power converters each power switch is composed of several parallel and/or series power devices to achieve the required current and voltage ratings. The presence of large numbers of devices provides possibilities for continued operation in the presence of a single device failure as demonstrated by Sinha et al [9].

Krayagina [10] examined the commercial benefits of employing fault tolerant control electronics in CNC machine tools and machining centres. In doing so, the trade off between reduced costs of down time and increased capital cost of the system was recognised.

1.5.2 The Relationship Between Fault Tolerance and Reliability

Introducing fault tolerance into a system increases the number of components. The reliability of the system, calculated by summing the failure rates of the components, decreases in proportion to the increased parts count. Clearly the chance of a fault in the system has increased and the MTBF has decreased.

However, introducing redundancy into a fault tolerant system ensures that the system continues to deliver the required level of performance, even in the presence of the first fault. If the faulted system is repaired promptly, the chance of a second failure occurring will be very small. Thus the reliability of the system in terms of the mean time between system failures will be dramatically higher than the reliability of a similar system without redundancy [6].

1.6 The Foundations of Fault Tolerance

White [6] observes that there are four requirements in the achievement of fault tolerance.

- Partitioning and Redundancy
- Fault Isolation
- Fault Detection and Annunciation
- On Line Repair / Survival until Service

Each of these requirements will now be examined in more detail.

1.6.1 Partitioning and Redundancy

Partitioning is crucial to achieving fault tolerance. It is the act of dividing the required performance of a system between one or more separate units. Simply dividing the performance up in this way minimises the impact of the failure of a unit. To ensure that the system is capable of delivering the required level of performance in the event of one or more failures, one or more redundant units must be added to the system.

Active Versus Passive Redundancy

Two forms of redundancy are identified by Prasad et al [11]. Termed as active redundancy and passive (or inherent) redundancy, these are two different means to the same end. Both approaches require that the system be partitioned and that some level of

redundancy be provided. The difference between the two techniques is in the way in which redundant resource is managed. In a system with passive redundancy the redundant units are normally operating and are indistinguishable from other units. If a unit fails the loading of all healthy units increases to make up for the resulting performance shortfall. In a system with active redundancy the redundant units are off line and are switched into the system to replace faulty units when the fault is detected.

In general active redundancy will be less robust, relying as it does on an additional intelligent unit to detect faults and switch in the appropriate redundant resource. However this route is often chosen, as it can require lower levels of redundancy.

The principle of using two or more parallel systems to provide fault tolerance has been practised in the fields of computing and control for some time. A typical example is the triplex arrangement used in aerospace control systems. A triplex control system has three sets of electronics, any one of which is capable of controlling the system. Two healthy channels are always required to vote out a faulty channel. This is an example of passive redundancy. In contrast ABB [12], [13] have recently employed a system with active redundancy to control a 13 MW turbo compressor. The system includes two sets of control electronics. In a drive of this size the control electronics make up only a small proportion of the complete cost, yet ABB claim that the two channel technique reduces the probability of breakdown by a factor of ten. The application of fault tolerance here does not extend to the power electronics or electric machine, so the claim should really be that the probability of breakdown due to failure of the control system is reduced by a factor of 10 - still a worthwhile achievement.

1.6.2 Fault Isolation

It is essential that any faulty units in a fault tolerant system do not disturb the operation of the remaining healthy units. This translates into a requirement for isolation between units. It is also important to ensure that a failed unit does not significantly disturb any inputs or outputs that it shares with other units.

1.6.3 Fault Detection and Annunciation

Fault detection is so important that Chapters 6, 7 and 8 are devoted to it. At a system level fault detection is required so that the system can be repaired or replaced in a timely manner, thus minimising the chance of a second failure in the faulted system.

1.6.4 On-Line Repair / Survival until Service

White gives the last requirement of fault tolerance as 'on line repair'. This is appropriate for a telephone exchange, where it is unacceptable to turn off the exchange to replace faulted equipment. In an aircraft however, the equivalent requirement is that further deterioration of the system should not occur whilst the aircraft is in flight. i.e. the system must operate until the aircraft can reach its scheduled destination. This is termed 'survival until service'.

1.7 The Route to a Fault Tolerant Drive

In a fault tolerant drive the four basic requirements of fault tolerance lead to investigation of the following areas:

1. Assessment of the likely faults.
2. Design analysis and simulation for each failure mode.
3. Development of techniques for imposing faults on the drive.
4. Imposition of faults on a working drive to validate the analysis and simulation results.
5. Development of sensors to aid fault detection.
6. Development of schemes for fault detection and identification.
7. Development of strategies for continued operation of the drive.
8. Analysis, simulation and testing of the drive performance for each fault.

An academic publication will typically address just one or two of these elements for a given system. However all of the elements must be addressed if the level of fault tolerance provided by the system is to be properly understood.

1.8 Significant Developments in the Fault Tolerant Drives Field

This section will present a broad outline of the important developments in the field of fault tolerant drives. The intention is to introduce as much of the significant work as possible, giving the reader an overview of the people involved and the historical context of their contributions. More detailed discussion of the work introduced here will be presented in later chapters where it is relevant.

The first developments in the field appear to be those published by Jahns [14] in 1980. Jahns investigates how drive 'reliability' could be improved by supplying each phase of an AC machine from an independent inverter.

About a decade later a series of publications started to appear, discussing drives which have been used in a manner described as fault tolerant. In some of these the fault tolerance has been achieved by driving an actuator from two separate machines in such a way as to ensure that the failure of one machine will not prevent the remaining healthy machine from driving the actuator. In the space shuttle remote manipulator system each joint actuator is driven by a pair of servo motors, [15]. Boyer [16] presents an actuation system for an aircraft flight control surface using two separate actuators, either of which is capable of controlling the surface in the event of the failure of the other. It is true that in some situations doubling up the electric drives will give the fault tolerance required, but there are disadvantages to this approach. The dual drive system will be twice the size, twice the weight and twice as expensive as a single drive. Having twice as many parts will also halve the MTBF. Finally the system designer must still ensure that the failure of one of the drives does not seriously affect the healthy drive, i.e. each drive must still be designed to be fail safe.

A number of publications have discussed the possibility of fault tolerant or high availability drives for ship propulsion. In this application it is vital that some drive capability remains in the event of a failure, but a reduction in output power is acceptable. The drives are large and of high value and this makes them more amenable to the application of fault tolerance. Furthermore, the high power level of a ship propulsion drive means that the power converter is naturally partitioned into a number of separate single phase inverters. Nerowski et al [17] use this natural partitioning to develop a ship propulsion drive using separate inverters arranged in groups, so that if one group is shut

down the drive will continue to operate from the remaining groups at a reduced power level. The same principle is used by Olivera et al [18] in a ship propulsion drive demonstrator. In 1991 Olivera et al reported on the faulted operation of this demonstrator showing that it is indeed capable of delivering useful power with one or more phases inactive. Finally, in 1995, Cosulich et al [1] take a step back to look at ship propulsion drives from a higher level and compare different approaches to partitioning the prop. shaft power requirement. A system using a single machine with a double winding and two power converters is compared to a system using a single 3 phase machine. The possibility of using two separate drives to power the single prop. shaft is also considered. Unlike the applications discussed by Boyer and Wu earlier, each machine would be designed to deliver only half the total shaft power required, so the mass and cost of the system would not be significantly greater than for a single drive. However, in the event of a fault, it is still important that the faulted drive does not unduly disturb the operation of the prop. shaft.

In the following subsections, important developments in the fault tolerant drives field have been organised by drive type.

1.8.1 Fault Tolerance in Switched Reluctance Drives

The switched reluctance machine, SRM, is widely acclaimed as being fault tolerant. In 1988 Richter [19], [20]. published work which was strongly supportive of the merits of switched reluctance machines for operation in harsh environments and in safety critical applications, citing fault tolerance as one of the major benefits of this machine type. By 1994 Richter [21], [22] had published work describing the development of an integrated switched reluctance starter generator at General Electric Aircraft Engines. In addition to the fault tolerant qualities of the SRM, the power converter in the starter generator system is split into 2 independent parts, each of which feeds a separate winding group comprising 3 phases in the machine.

It is true that the SRM has a degree of inherent fault tolerance. This arises because most SRM designs naturally provide complete electrical isolation and some physical and magnetic isolation between the machine phases. Furthermore the SRM is typically fed from a power converter which uses a separate power electronic circuit to feed each machine phase, Figure 3. It is often claimed that the unusual unipolar inverter used to

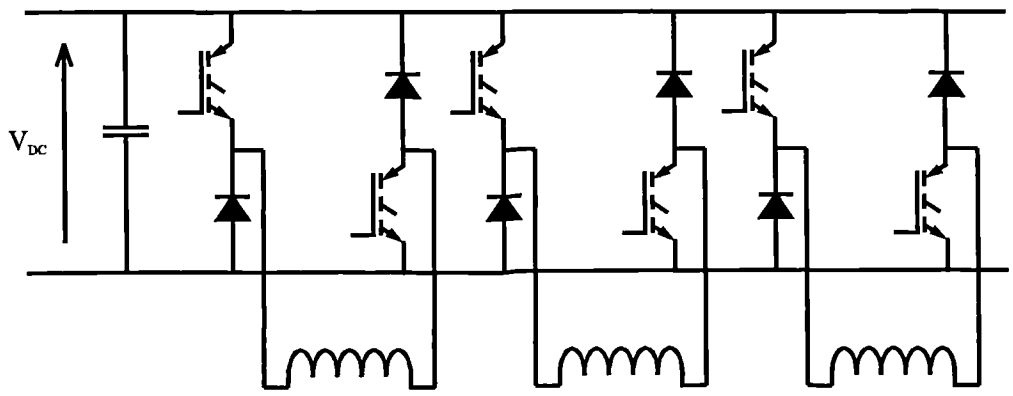

Figure 3 : SRM Power Converter.

feed the SRM cannot suffer from shoot-through faults, Radun [23]. Certainly it is true that if both power switches in a phase turn on simultaneously then the rise in current is limited by the inductance of the phase winding.

Stephens [24] examines fault detection and management systems for SRMs. His publication in 1991 probably marks the first serious effort to develop the fault tolerant potential of the SRM. Only recently has further work been presented, this time by Miller [25] showing that some of the most likely faults (turn to turn and power switch short circuits) can result in large fault currents and/or large unbalance forces, thus having a very serious impact on the machine's operation. One point that has apparently been universally overlooked is that the single phase unipolar inverter is still vulnerable to a short circuit failure in one of the freewheel diodes.

Fault tolerance in the SRM can be compromised, as some SRM winding configurations introduce significant mutual coupling between phases in an effort to improve certain aspect of machine performance. Examples of such machines include the fully pitched SRM presented by Mecrow [26] and the SR starter generator presented by Richter [22].

1.8.2 Fault Tolerant Induction Drives

The work by Jahns [14] on independent phase drives used an induction machine, but he acknowledged that coupling between phases was a problem. In 1985 Palmer [27] used an induction machine with bifilar windings, each winding being supplied by an independent unipolar bridge, to eliminate the possibility of shoot-through caused by transistor faults.

The induction machine fed by a six switch inverter is an extremely poor candidate for applications demanding fault tolerance, since it naturally possesses strong coupling between phases. Moreover, the six switch inverter offers only very limited possibilities for faulted operation, as will be shown in Chapter 2. However this is the most common drive configuration in many applications so the benefits of continued operation, even with limited performance, under fault conditions would be widespread. Short circuit faults are very serious and prevent further operation of this drive type, so all the effort has been concentrated on open circuit failures.

Kastha and Bose [28] started by examining the fault modes in a voltage fed inverter system for an induction motor drive in 1992. This was soon followed by work published by Lui et al [29] and Elch-Heb et al [30], [31] which examines single phase operation (using 2 machine phases) of induction motor drives. Most recently Kastha and Bose [7], [8] have examined techniques for reducing the pulsating torque which results from single phase operation.

1.8.3 Fault Tolerance in Permanent Magnet Drives

In 1987 Nyamusa and Demerdash [32] investigated the effect of a partial armature short circuit on a brushless permanent magnet (PM) machine. This work is really about simulation methods, but is significant here because it deals in detail with an important machine fault. A more general study of fault modes in a brushless DC drive is presented in a simulation study by Wallace and Spee [33] in 1988. By 1990 Wallace and Spee [34] had considered the fault consequences and presented 'remedial strategies' for various drive failures. However, for the drive configuration considered, no remedial strategy is available for the short circuit failure of power devices or machine windings.

Prompted by interest from Dowty Fuels Ltd, now part of Lucas Aerospace, Jack and Mecrow began investigating the possibility of fault tolerant PM drives. By 1994 they had published some simulation work which suggested that a PM machine could be designed to withstand a shorted machine winding, [35]. The work presented in this thesis continues directly from that point and further publications by the author are listed in Chapter 10, [36], [37], [38], [39], [40] and [41].

In the course of the research presented here the other significant developments in the field of permanent magnet synchronous machine (PMSM) drives have all been in the area of fault analysis and simulation, first by Ronkowski and Szczesny et al [42], [43]. Most recently Bianchi et al [44] have examined PMSM drive failures in the flux weakening regime.

1.8.4 Fault Tolerant Variable Reluctance Motor Drives

Very little attention has been given to the use of variable reluctance machines for fault tolerant drives. The notable exception is a recent paper by Qin et al [45] which examines the possibility of driving the flux linking the faulted phase to zero through control of the 2 remaining healthy phases. This approach is very effective in preventing current from flowing in the faulted phase. Unfortunately, controlling the flux in the faulted phase to zero reduces the drive torque capability to 0.13 p.u. Even if the flux in the faulted phase is controlled to 1.0 p.u. (allowing 1.0 p.u. current to flow in the fault) the torque output is still only 0.3 p.u. Thus at present this machine type does not seem attractive for applications which demand fault tolerance.

1.8.5 Fault Detection in Power Converters

A number of publications have appeared recently proposing the use of an expert system to diagnose faults in a six switch inverter. Renfrew and Tian [46] first suggested diagnosis and monitoring of AC drives by expert systems in 1991. Debebe [47] presented a similar knowledge based system for fault diagnosis in AC drives in 1992. Then in 1994 Aris, Zhang and Hulley [48] proposed a scheme for fault diagnosis of inverter circuits for AC drives using a knowledge base and fuzzy logic. Later they developed their ideas to provide an on-line fault diagnosis system, [49]. These schemes tend to be too slow for use as part of a fault tolerant drive, where severe faults require rapid detection if fault propagation is to be prevented.

Craig [50] approaches the subject of power circuit fault diagnosis from an angle which is more helpful in developing a fault tolerant drive. His paper of 1991 presents a detection scheme for identifying faults in a single phase bridge. The technique works in real time and is capable of detecting short circuit device failures before fault propagation occurs.

1.9 The Application

The idea of a fault tolerant drive clearly has a wide range of possible applications. To provide a focus for the work a target application has been chosen. As will be seen in Chapter 2, the exact requirements of the application will influence many of the decisions which are made in the course of the design process. Whilst this will result in a drive which is matched to a particular application, it will be seen that the resulting drive is also appropriate for other functions.

1.9.1 'More Electric Aircraft' and Electric Fuel Pumping

The application chosen comes from the aerospace industry. There is currently a great deal of energy being directed by manufacturers of airframes, aero-engines and aerospace components at the concept of the 'all electric aircraft'. The idea is hardly new. Authors such as Cronin [51] were extolling the virtues of the improved efficiency that would be gained in an 'all electric aircraft' back in 1980, whilst at the same time citing work published as early as 1965. Perhaps because of the ambiguity of the phrase, it is increasingly common to see the 'All Electric Aircraft' referred to as the 'More Electric Aircraft' - there is certainly no intention to replace aviation fuel with electricity.

The first stage en route to the 'all electric aircraft', 'fly by wire', is well advanced with several aircraft relying on signal wiring as a replacement for mechanical linkages to make the controlling connections between systems in the aircraft. The second stage is termed 'power by wire'. The aim of the research into 'power by wire' is to assess the possible benefits in moving away from aircraft using mechanical linkages, hydraulic systems and pneumatic systems to distribute power around the airframe. In a 'power by wire' system the power would all be distributed by an electric power bus. In the initial enthusiasm for power by wire, the specification for the power distribution bus was set at 270 V DC [51]. In a paper on power management and distribution, Weimer [52] implies that the 270 V DC bus is a widely accepted standard for the 'more electric aircraft'.

The aerospace industry is naturally slow to adopt new technology and rightly so, given the potentially serious consequences of equipment failure. Nevertheless the flow of publications recording the progress and benefits of the 'more electric aircraft' continue to appear, with notable papers from Przybylko [53], Cronin [54] and Quigley [55]. Whilst progress is slow the 'all electric aircraft' seems eventually to be inevitable.

The development of 'power by wire' is most advanced in the small helicopter market with Jones and Jarvis [56] reporting the development of a full set of 'all electric engine' accessories, in 1990. The reported system included a generator, fuel pump, lubrication oil pump and variable geometry electric actuator. Several of the electrical accessories had been developed to the point were they had completed several hundred hours of operation on an engine test bed. In contrast Richter's examination of accessories appropriate to large civil aircraft engines, [19], was only at the stage of a paper study at that time.

Examples of electric accessory development away from the aircraft engine are mostly concerned with actuation systems. Even the primary flight control actuators could be powered electrically as demonstrated by the work of Jahns et al [57] and Crowder [58] on 2 electro-hydrostatic actuators. Boyer [16] presents the design and test results from a redundant electro-hydrostatic actuation system for a primary flight surface on a thin wing military aircraft. Some electric actuation has already appeared in passenger aircraft, albeit in secondary applications. One such example is supplied by Vickers Systems to lower and retract the undercarriage and to control steering and braking functions in a nine-seat private aircraft. The system is driven by an electrically powered hydraulic power pack, with hydraulic power derived from a variable displacement pump, driven by a permanent magnet brushless DC machine [59].

Whilst the introduction of the 270V DC bus on aircraft in service still seems some way off, a number of research establishments are working on prototype electrically powered accessories and perhaps most significantly GE has built a prototype starter generator, [22], [30] and [21]. This 250 kW machine would be at the heart of a 'power by wire' aircraft, combining the aircraft engine starter and electrical power generation functions into one electric drive.

One of the largest loads on the electrical power system would be the electric fuel pump [19]. Aircraft engines use a fuel pump which is at present mechanically coupled to the engine. The pump turns at a speed proportional to engine speed. The pump speed is arranged so that it is always capable of delivering sufficient fuel to the engine. Typically the largest amounts of fuel are required during takeoff when the engine is rotating slowly, in the relatively dense atmosphere. The supply of fuel to the engine is regulated by controlling a bleed valve which allows any fuel that is not required by the engine to be

re-circulated. The aircraft spends most of its operating time cruising at high altitude. Here the engines are running close to maximum speed to provide thrust in the thinner atmosphere, yet the fuel requirement is fairly low. The result is that the greatest proportion of the fuel is re-circulated. The action of pumping the fuel around the re-circulation path repeatedly is wasteful of energy and results in raising the fuel temperature unnecessarily. This latter point is especially significant since the fuel is used as a coolant on the aircraft and may already be close to its maximum temperature by the time it arrives at the fuel pump. Replacing the existing system with a fuel pump coupled to an adjustable speed drive, would provide a system capable of pumping only the amount of fuel actually needed and would give a number of other advantages, Radun [23]. Some of the principle benefits would be:

- Reduced mass due to the elimination of the mechanical linkage between the engine and the fuel pump.
- Higher efficiency.
- Easier maintenance as mechanical linkage is replaced by an electrical connector.
- Easier thermal management due to reduced heating of the fuel by unnecessary fuel circulation.

The operation of the of the fuel pump is obviously safety critical and this immediately raises the question: 'Can an electric fuel pump be designed to be sufficiently dependable?'

1.9.2 The Fuel Pump Drive Requirement

All drives for aerospace applications share three key requirements:

- High power to mass ratio, termed 'specific power'.
- High reliability.
- Solutions must be cost effective but are less cost sensitive than most applications.

The fuel pump application has the following characteristics:

- Load torque verses speed profile is close to a square law.

- Maximum torque is required at full speed.
- Some drive torque ripple can be tolerated in normal operation.
- Increased torque ripple can be tolerated during the short period of post fault operation required for an aircraft to complete its journey.

To help guide the project Lucas Aerospace supplied a specification which is representative of the fuel pump requirements for a small to medium sized civil aircraft. The drive has been designed around this target specification, shown in Table 1.

Parameter	Value	Units
Shaft power	16	kW
Maximum speed	13 000	rpm
Maximum temperature rise in the motor	80	°C
DC link voltage	270	V
Rated torque	11.75	Nm

Table 1 : Fuel Pump Drive Requirement

1.9.3 Project Scope

The aim is to design a drive to the specification provided. The drive should be capable of surviving any single point failure, detecting the failure and continuing to operate. When compared to a drive without fault tolerance, the fault tolerant design should involve no extra parts and no extra cost (and in this instance no extra size or weight). As the arguments are developed, the extent to which these objectives can be achieved and the trade-offs which must be made will become apparent.

The scope of the faults considered is limited to electrical failures in the electric machine and power converter. Providing fault tolerance in the control electronics is most likely to be achieved by a multiple channel system, with voting as described by Belmont [61]. However, the control electronics and the best form of electrical interface are beyond the scope of this project.

The sensors used to measure quantities such as shaft position and phase current are beyond the scope of the work here. The possibility of sensor failure forces consideration of sensorless operation. Hofer [62] discusses how a state observer may be used to estimate motor currents and voltages and hence motor speed. The estimation of these

quantities is the common aim of sensorless operation, however Hofer uses the state observer in a system with sensors to provide redundancy and permit continued operation in the event of sensor failure.

Chapter 2

DRIVE DESIGN

2.1 Objectives

This chapter investigates the possibilities for applying the first two requirements of a fault tolerant system to an electric drive. The first requirement demands the development of a drive design which has been partitioned into a number of units, each providing an equal part of the required performance. Redundancy must be added to the drive to ensure that it can continue to deliver the required performance if any unit is disabled. The second requirement, that of fault isolation, requires attention to be given to fault isolation and the prevention of fault propagation, to ensure that a fault in one unit does not affect the operation of remaining healthy units. It is important that the faulty unit does not unduly disturb the common input or output. This will be considered briefly here, but will not be considered in detail until Chapter 5.

Once a design has been finalised it will be characterised with some simple tests to provide a basic set of measured machine parameters. These parameters will be useful both to verify that some of the design aims have been achieved and as a starting point in designing the control scheme for the drive. Finally, the penalties resulting from each aspect of the design associated with fault tolerance will be examined in terms of drive mass, cost and complexity.

2.2 Drive Topology

The first task in designing a fault tolerant system is to introduce partitioning. Perhaps the simplest way of partitioning the drive requirement is to use two or more drives to provide the mechanical power required. In this application a mechanical coupling could be used to connect each drive to a single fuel pump or each drive could power one of several smaller pumps. The next stage in the design process is to introduce redundancy by adding one or more additional units. Let us assume that it is sufficient to add one

additional drive. Reasonable combinations might be '1+1 redundancy', with the total pumping power requirement supplied by 2 drives, either one of which is capable of delivering the full pumping requirement in the event of the failure of the other. Another possibility is '2+1 redundancy', in which the total pumping requirement is split in half and supplied by 3 pump drives such that any 2 could supply the full pumping requirement in the event of the failure of the other one. It is easy to see that either of these systems could provide fault tolerance and a very high level of availability but it is at a high cost in terms of increased mass, cost and complexity.

In this chapter, the possibilities for introducing a single fault tolerant drive will be examined. Again the exercise of partitioning must be performed. It is natural to partition a machine and power converter into phases, with each phase fed from a common DC bus and supplying mechanical power to a common shaft, whilst remaining independent in all other respects. It appears that Jahns [14] was the first to examine the possibility of partitioning a drive in this way.

Much of the research into faulted drive operation focuses on the continued operation of a star connected three phase machine fed by a six switch inverter, as shown in Figure 4. This is problematical because the short circuit failure of a power device leaves one terminal of the machine effectively connected to one side of the supply. Wallace and Spee [34] show that it is no longer possible to get any useful torque from the resulting configuration.

Note that whilst a device short circuit makes it impossible to control bipolar currents, it will still be possible to drive unipolar current in the machine. Although the drive can no longer produce any useful output power, it is still capable of producing braking torque. If a synchronous (slip free) machine is used then the drive will be capable of fail-freeze operation, i.e. of opposing changes in shaft position. This modest faulted capability could be very useful in some actuation systems, where locking the system in position is the preferred action in the event of a failure.

In Figure 4 it is clear that the inability to withstand the transistor short circuit fault in this configuration results from the existence of a star point connection. Similar arguments apply to the delta connected configuration. Wallace and Spee [34] examine the benefits of adding an additional phase leg to control the star point under various fault conditions.

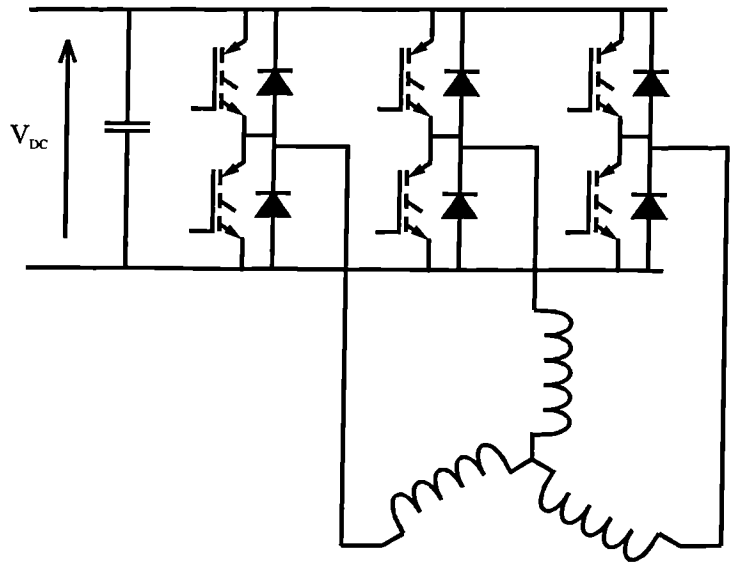

Figure 4 : A Three Phase Star Connected Machine Fed by a Six Switch Inverter

The 2 additional transistors are held off in normal operation but controlled in the case of an open circuit phase failure, allowing the drive to produce more torque than would otherwise be the case in the event of this fault. However even with the ability to control the star point, the drive is still unable to deliver any useful torque in the event of a shorted power device. To understand the reasons for this, consider a power device short circuit which results in one terminal of the machine becoming permanently connected to the positive rail of the DC link. Any switching state which results in the connection of any other terminal of the machine to the negative rail will also result in a progressive build up of current in the faulted phase. This current cannot be driven back down by the faulted power converter and the only way to prevent the build up of current in the faulty phase whilst continuing to operate the remaining healthy phases is to isolate the faulty phase from the remainder of the drive.

Isolation of a faulty phase from the remainder of the drive could possibly be achieved by fuses but this would require the fault to result in excessive current flow in order to cause the fuse to rupture. Furthermore, fuses are not favoured in this application as they are perceived as having a low reliability. A more elegant solution is to bring both ends of each phase out of the machine and supply each phase of the machine from a separate single phase power converter.

2.2.1 Phase Number

The choice of the number of phases into which to partition the drive involves a trade off. On the one hand a low phase count minimises complexity and the likely fault rate. On the other hand a high phase number minimises the impact of a failure in a single phase unit. Once a phase number has been chosen, redundancy must be added to ensure that the drive continues to provide the required level of mechanical power in the event of a failure. Ultimately the number of phases chosen and the number of additional phases will vary according to three principle factors:

- The mean time between failures for a single phase unit - MTBF calculated for a unit.
- The mean length of time for which the faulted drive must continue to operate - the MTTR.
- The mean time between failures required for the drive - MTBF required for the system.

The choice of phase number will also be affected by cost considerations which will permit larger phase numbers for larger drives, where the cost of the control for an additional phase is proportionately smaller.

The analysis presented assumes that each phase is equally able to contribute to the total torque output of the drive. This is only true if the torque is averaged over a full electrical cycle. If smooth torque output is required or if full torque is required at standstill then the situation becomes more complex. In the case of an open circuit failure the remaining phases must fill in the gap in torque at the angles at which it occurs. For an electric drive, the torque available from each phase varies with rotor position. Thus the number of phases chosen will affect the torque ripple resulting from a given current profile.

The choice of phase number for the demonstrator drive was made without the benefit of reliability data or a firm requirement for system MTBF. The mechanical power requirement was partitioned into 5 and it was assumed that adding 1 additional phase would provide the required system MTBF. The result is a '5+1 redundant', 6 phase machine. The resulting drive topology shown in Figure 5 may seem strange but is commonly used in large drives. For example Nerowski et al [17] describe a 1.1 MW PM synchronous drive which uses this topology. Partitioning the total power requirement

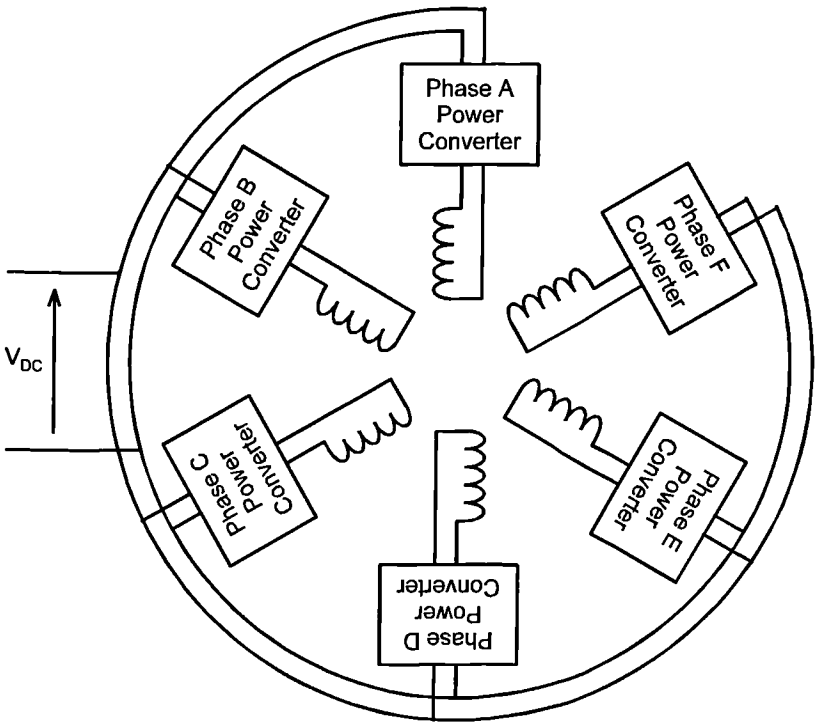

Figure 5 : Fault Tolerant Demonstrator Drive Topology

into many smaller phases makes the power levels more manageable and introduces the possibility of fault tolerance. Nerowski claims an improved level of availability due to this partitioning which is likely to be true. However the paper does not address the important areas of fault isolation, detection and reporting, so the drive could not claim to be fault tolerant. Other possible topologies which exploit the redundancy naturally present in a multiphase drive system are explored by Cosulich et al. [1].

2.2.2 System Failure Rate

Based on the phase number chosen, Lucas Aerospace undertook some analysis work which compared the MTBF of the 5+1 redundant topology presented here to the MTBF for a standard 3 phase drive system.

Specific requirements for achieving acceptable safety will be influenced by the intended application, e.g. for civil vs. military, number of engines, etc. The 180 minutes Extended range Twin engine OPerationS (ETOPS) requirements define the acceptable safety levels for twin engine civil passenger aircraft to be permitted to fly over significant distances (180 minutes flying time) from the nearest suitable airfield. Typical civil aviation safety requirements were used as a target for the safety critical analysis below. These included:

- No single fault is permitted to cause hazardous failure (e.g. unprotected over-speed) on the application.
- Any single fault which will cause or contribute to an in-flight shut down (IFSD) must have a failure rate of less than 1×10^{-7} per engine hour.
- Undetectable single faults which, in combination with a subsequent fault, could cause an IFSD must have a failure rate of less than 1×10^{-8} per engine hour.

Table 2 shows that the projected system failure rate for a standard 3 phase drive was 67.5×10^{-6}. To meet the ETOPS requirement the fault rates would need to be 3-4 orders of magnitude lower. The same fault rates were used to calculate the MTBF for two phases to independently fail in a 5+1 redundant drive. The resulting system failure rate was substantially lower than for the standard 3 phase drive.

Cause of failure	Failure rate allocated per phase (per hour)	No. of phases	Total failure rate (per flying Hour)
Short circuit between phases	6.7×10^{-6}	3	20.1×10^{-6}
Open circuit in windings / joints	13.4×10^{-6}	3	40.2×10^{-6}
Short circuit on connections (2 off)	1.0×10^{-6}	3	3.0×10^{-6}
Open circuit on connections (2 off)	1.0×10^{-6}	3	3.0×10^{-6}
Open/short circuit (misc.)	0.4×10^{-6}	3	1.2×10^{-6}
Total for all faults			**67.5×10^{-6}**

Table 2 : System Failure Rate Calculation for a Standard 3 Phase Drive

Although the analysis is based on rough figures, it demonstrates 3 significant facts:

- It appears that a 3 phase drive is not capable of providing a sufficiently high MTBF to satisfy the ETOPS requirement.
- Adopting a 5+1 redundancy provides an improvement in system failure rate of the order of 10^4.
- The resulting MTBF for the 5+1 redundant drive is sufficient to satisfy the ETOPS requirement.

2.3 Design Requirements for a Fault Tolerant Drive

One of the prerequisites for a fault tolerant system, that of drive function partitioning has been addressed. Now it is necessary to examine how the second prerequisite for fault tolerance, fault isolation, might be achieved in an electric drive.

2.3.1 The Machine Type

As described in Chapter 1, the SRM naturally possesses certain qualities which lend it a degree of fault tolerance. These qualities ensure the independence between phases. Typical permanent magnet synchronous machines (PMSMs) do not possess these qualities. In particular PMSMs often feature overlapping phase windings, stator slots shared by conductors of more than one phase, strong mutual coupling between phases. Moreover in the event of a short circuit fault the magnet flux will drive large currents into the fault, limited only by the typically low phase inductance and resistance.

Whilst the SRM appears to possess a natural advantage in terms of machine topology, it appears that the PMSM is capable of delivering as much as twice the power output per unit mass at the power levels considered here (5 to 20 kW), [37]. High specific power is an important consideration in many applications and especially in aerospace. If the fault tolerant features of the SRM could be embodied in a PMSM design, without compromising the high specific power of the PM machine, then the resulting 'fault tolerant PM machine' would be a strong contender for aerospace applications.

2.3.2 Single Phase Power Converter Topology

A voltage source inverter (VSI) has been chosen in preference to a current source inverter or a resonant converter. This choice is made largely because of the simplicity and low parts count of this type of inverter, although other advantages are gained too. The current can be profiled and simpler inverter topology makes fault analysis easier.

In a system supplied from a DC link there are 3 possibilities for the resulting voltage source inverter, as shown in Figure 6. Of these the H-bridge shown in Figure 6(c) has been selected for its simplicity.

2.3.3 Fault Containment and Fault Isolation by Design

A 6 phase synchronous PM drive has been designed to provide effective isolation between phases and limit the fault current resulting from a phase short circuit at the machine terminals. As shown in Figure 5, each phase comprises a single phase inverter supplying a phase winding in the machine from a common 270 V DC bus. The key attributes of the drive, which ensure effective fault isolation and help to prevent fault propagation, are:

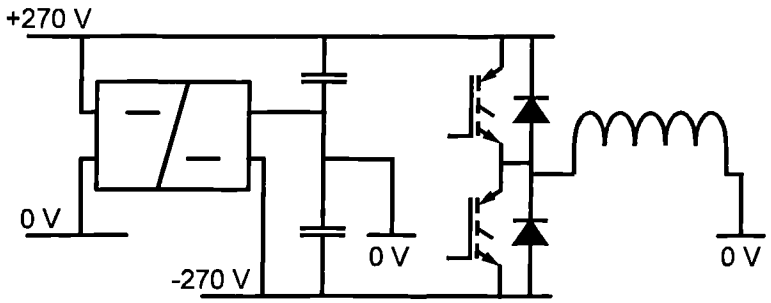

(a) Single Phase Leg with a DC/DC Converter to Generate a Negative Rail

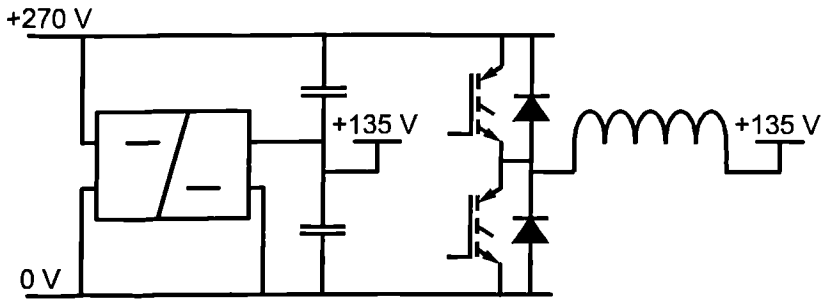

(b) Single Phase Leg with a DC/DC Converter to Generate a Centre Rail

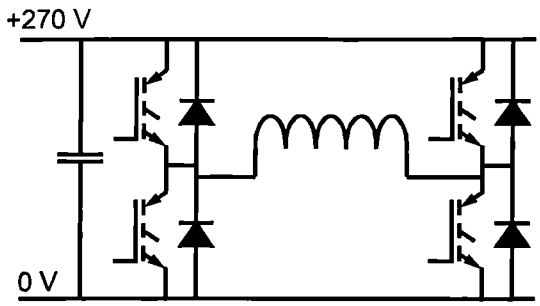

(c) Two Independent Phase Legs Forming the Familiar H-Bridge

Figure 6 : Possible Topologies for Single Phase Bipolar Power Converters

Electrical Isolation between Phases

Electrical isolation between phases has been shown to be an essential requirement if the drive is to be capable of continued operation in the event of a power device short circuit fault. Electrical isolation between phases is most easily achieved by supplying each phase winding from a separate single phase bridge. The use of an independent H-bridge to drive each phase doubles the number of power devices but only marginally increases the total volt-ampere rating, as each device need only withstand the phase voltage rather than the line voltage for star connected systems.

Physical Isolation between Phases

A phase to phase fault is especially serious, since it will disable 2 phases. Each phase winding is placed around a single tooth and only alternate teeth carry phase windings. Thus all phase windings, including end windings, are physically separated, thus virtually eliminating the possibility of a phase to phase fault.

Effective Thermal Isolation between Components

Thermal separation between phase windings and between power devices provides a barrier to fault propagation. By arranging adequate cooling for the outer surface of the stator the dominant temperature rise in the machine is confined within individual slots. This, coupled with physical separation of the phases, leads to effective thermal isolation between phases. In the power converter it is necessary to separate even those power devices in the same phase, as post fault control depends on failures being limited to a single device. This is achieved by a similar combination of adequate direct cooling and physical separation.

Magnetic Isolation between Phases

Without magnetic isolation fault currents in one phase will induce large voltages in other phases, preventing adequate control of them [35]. In surface-mounted magnet designs a substantial amount of the phase inductance arises due to cross-slot leakage flux. If this is to remain solely a self-inductance component then each slot should contain the conductors of one phase only.

The component of armature reaction flux which crosses the air gap will always contain an element which links the other phases. However, if a surface mounted magnet design with a non-magnetic retaining sleeve is employed, then the presence of the sleeve combined with relatively deep magnets greatly reduces the air gap component of the armature reaction field, so that in effect the mutual coupling is insignificant.

Implicit Limiting of Phase Winding Short Circuit Currents

Short circuits can occur within the machine and thus cannot always be protected against by fuses. The fault current which flows in the event of a single phase short

circuit, applied at the machine terminals, is limited only by the phase impedance. The short circuit current can be limited to rated current by designing the machine with a phase inductance of 1.0 per unit. By adopting this approach it is possible to operate indefinitely with a phase short circuit fault without excessive temperature rise. Torque ripple resulting from the short circuit current in the faulted phase will inevitably be present unless special measures are employed.

Producing an effective armature inductance of 1.0 per unit is generally considered difficult to achieve in a small permanent magnet machine. However in the machine design presented here it is relatively easy to achieve this inductance simply by manipulating the depth and width of the slot opening.

Implicit Limiting and Control of Shoot-Through Currents.

Power device short circuit faults in an H-bridge will result in a shoot-through fault when the other power switch in the faulted leg is turned on. If the DC link is not to be unduly disturbed by these faults then the fault current must be limited and promptly extinguished.

The choice of insulated gate bipolar transistors (IGBTs) as the main power switches in the bridge ensures that the fault current is limited to approximately 10 times rated current by device gain. Typical IGBTs are designed to survive this short circuit condition for a period termed 'the short circuit withstand time', typically 10 μs. Post fault control depends on faults being confined to a single device within the phase, so fault propagation must be prevented. Clearly some means of extinguishing the fault current within the short circuit withstand time of the devices is required both to prevent fault propagation and to prevent undue disturbance of the DC link.

One possibility is the use of fast acting fuses. However, fuse rupture results in a rapid and uncontrolled reduction of current. The induced voltage resulting from the high di/dt is likely to cause the breakdown and destruction of other power devices. In the authors experience, a shoot-through fault which is extinguished by the eventual open circuit failure of one IGBT typically results in the failure of a further two IGBTs in the bridge. Fault propagation of this nature is very serious as fault tolerance is more difficult to achieve in the presence of multiple faults. Fuses are also regarded as

unreliable circuit elements and it is difficult to arrange for sufficiently fast rupturing to protect the power device. Since it is the gain of the power device which limits the fault current, much of the DC link voltage is dropped across the power device. Thus the power dissipation in the power device is much higher than in the fuse, causing the failure of the power device before the fuse.

The preferred method of extinguishing the fault current is to turn it off with a healthy IGBT. There are always two series connected IGBTs in the shoot-through current path and provided action is taken before fault propagation occurs one of these IGBTs will always be capable of turning off. Many authors (e.g. Chokhawala et al [63]) have shown that it is possible to detect when an IGBT is turned on into a short circuit and for the IGBT to turn off the fault current safely within the short circuit withstand time of the IGBT.

2.3.4 Phase Failures and the DC Link

The requirement for fault current limiting ensures that the DC link is not disturbed, provided that the fault current duration is limited sufficiently to ensure that the link capacitance is not discharged significantly. The DC link is also susceptible to the failure of the DC link reservoir capacitance. The techniques for protecting the DC link from reservoir capacitance failure will be discussed when the drive design is presented later in this chapter.

2.4 Design Realisation

2.4.1 The Electric Machine

The details of the machine design are summarised in Table 3.

Parameter	Value	Units
Operating speed	13000	rpm
Poles	8	
Turns per phase	50	
Phases	6	
Active mass	3.227	kg
Target torque with 6 phases operating (6/5 * 11.75 Nm)	14.1	Nm
Target rated current (rms. sinusoidal)	21.1	A

Table 3 : Machine Design

The machine is wound with a single coil per phase. Each coil is wound around an alternate single tooth. The general configuration for a 6 phase, 8 pole machine is shown in Figure 7.

The stator for this machine is shown in Figure 8. The machine stator is wound in a manner which provides physical, thermal, electrical and magnetic isolation between phases, thus enabling the machine to continue operating with a faulted phase. This topology naturally results in a relatively high per unit reactance and with careful electromagnetic design a 1.0 per unit reactance may be achieved.

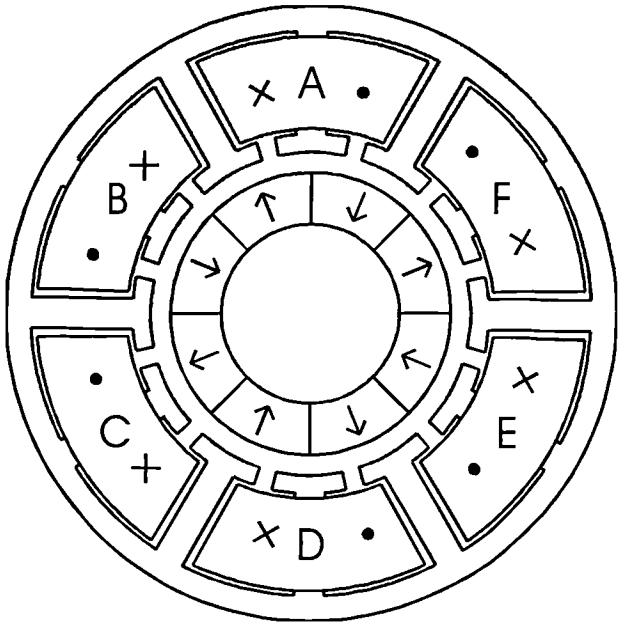

Figure 7 : Machine Topology

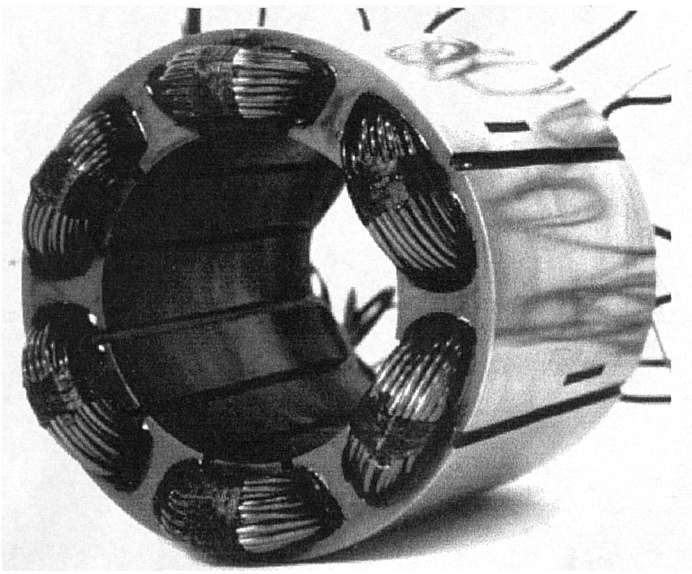

Figure 8 : A Stator Unit Wound with the Fault-Tolerant Winding Arrangement

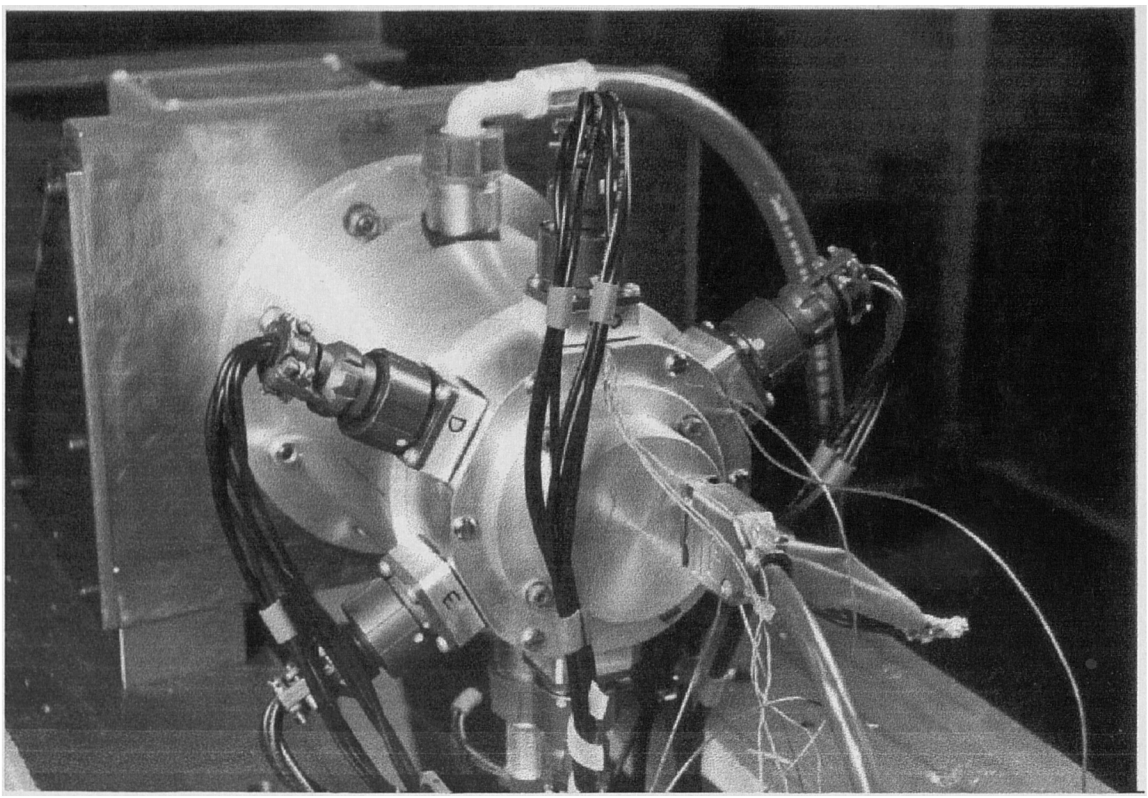

Figure 9 : The Demonstrator Machine Mounted on the Test Rig

The rotor assembly has 8 samarium cobalt magnets assembled onto a shaft and retained by an inconel sleeve with a wall thickness of 1.5 mm. The rotor and stator were assembled into a purpose built frame with a brushless resolver to provide rotor position information. The aluminium frame has a channel around its periphery to permit liquid cooling. The 6 independent phase connectors, the coolant inlet and outlet lines and the resolver connection harness are all clearly visible in Figure 9. Further details of the machine design are given in Appendix A.

2.4.2 The Power Converter

The inverter design comprises six independent single phase bridges supplied from a common 270 V DC link. Each bridge is capable of delivering 41.3 A continuously at a switching frequency of 10 kHz. Further details of the bridge design are given in Appendix B.

2.4.3 The Reservoir Capacitor

The DC link capacitor is recognised as one of the principle causes re in drive systems [64]. This is partly due to the low life of electrolytic capac' in aerospace

application it is probable that multi-layer ceramic capacitors will be used to provide the necessary capacitance, Radun [23]. Multi-layer ceramic capacitors have the advantages of increased reliability, no ripple current limitation, lower sensitivity to both temperature and pressure variations and longer life. Although multi-layer ceramic capacitors provide seven times lower capacitance per unit volume compared with their electrolytic equivalents, this is partly compensated for because ceramic capacitors are not ripple current limited.

Assuming that the reservoir capacitor is an essential component, some form of fault tolerant reservoir should be incorporated. The reservoir capacitance is naturally partitioned into several parallel capacitors. Redundancy can be achieved by simply providing more parallel connected capacitors than the drive actually requires. A difficulty arises as a short circuit capacitor failure is possible and each capacitor is connected directly across the DC link. Fusing each capacitor individually may be an acceptable solution providing that the remaining healthy capacitors are capable of supporting the link for the time required for the fuse to operate. However, a more reliable solution is to arrange for each parallel connected capacitor to be composed of two or more individual capacitors in series such that, if one capacitor in the series connection fails short circuit, the remaining capacitors in that series connected group are capable of supporting the link voltage. The drawback of a series connected capacitor solution is that adding series connected capacitors increases volume whilst at the same time reducing the effective capacitance of the series connected group.

Time constraints prevented the incorporation of 'fault tolerant reservoir capacitors' into the power converters used. Instead each phase is decoupled by a pair of 1000 μF electrolytic capacitors in parallel with a pair of 0.47 μF polypropylene capacitors. The capacitor size is fixed by the ripple current limit on the electrolytic capacitors.

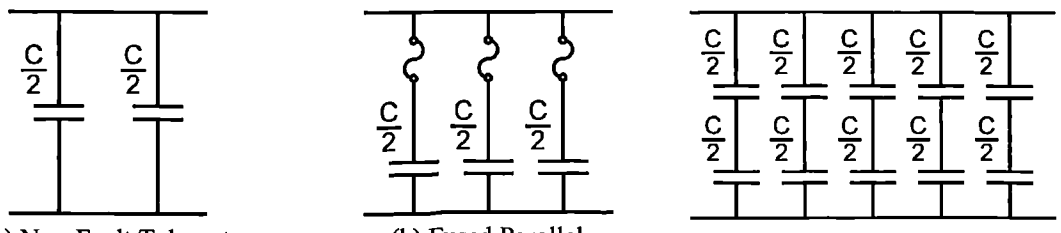

(a) Non-Fault Tolerant (b) Fused Parallel (c) Series / Parallel

Figure 10 : Possible Topologies for a Fault Tolerant Reservoir Capacitor

2.4.4 The Control Electronics

A set of control electronics has been developed to control the machine and power electronics. The control electronics comprise a TMS320C31 DSP card and several I/O cards mounted in a 19 inch 6U rack. The analogue to digital conversion of the phase current and DC link voltage measurements is performed by a set of electronics in a separate screened enclosure adjacent to the power converter. The digital signals are then fed to the controller in differential serial form down twisted pairs to provide a high level of immunity from electromagnetic interference. More details of the control electronics, including a functional block diagram, are given in Appendix C.

2.5 Initial Machine Characterisation

2.5.1 The Drive Test Rig

One important part of the fault tolerant drive demonstrator has not yet been fully described, the test rig. This comprises a DC power supply and a high speed motor loading rig.

A simple DC power supply has been built to supply the DC link required by the inverter. The DC voltage is derived by rectifying a 3 phase 50 Hz supply. The power supply is housed in the same enclosure as the inverter.

The motor loading rig allows the demonstrator machine to be tested under load. The loading rig includes a torque meter which provides shaft torque measurement from speeds of 13000 rpm down to standstill.

Further information about the test rig is contained in Appendix D.

2.5.2 Measured Parameters of the Demonstrator Machine

The parameters of the machine have been measured and are presented in Table 4. Also shown in the table are per unit (p.u.) values. The per unit system used throughout this work is based on E_f and I_a, not on V and I_a which are often used by others.

The phase inductance is a little higher than intended at 1.275 mH, representing 1.04 p.u. The phase to phase coupling is only 0.028 mH, which is 0.023 p.u. and 2.2 % of the phase inductance, demonstrating that the objective of effective magnetic isolation between phases has been achieved.

Parameter	Value	Units	P.U. Value
Peak phase back-EMF at operating speed	198.9	V	1.00
Target rated current (rms. sinusoidal)	21.1	A	1.00
Phase resistance	0.156	Ω	0.023
Phase inductance	1.275	mH	1.04
Mutual coupling between phases	0.028	mH	0.023
Short circuit current	18.4	A	0.87
Temperature rise in a winding due to losses in that winding	1.21	°C/W	-
Temperature rise in a winding due to losses in an adjacent winding	0.064	°C/W	-

Table 4 : Measured Machine Parameters

Thermal tests show that the temperature rise of a winding due to losses in that winding is 1.21 °C/W. The thermal coupling between phases may be expressed in terms of the temperature rise in winding due to loss in an adjacent winding. The thermal coupling between phases is 0.064 °C/W. As a percentage of thermal coupling from each phase to the coolant, the thermal coupling between phases is 5.3%. Thus a substantial degree of thermal isolation between phases has been achieved.

The peak back-EMF of the machine in the table above is calculated from back-EMF data measured at 4000 rpm. This back-EMF data is plotted in Figure 11. Note that there is good symmetry between pole pairs. The back-EMF is essentially sinusoidal, but with a visible third harmonic component. The third harmonic is present because each phase is independent rather than star connected.

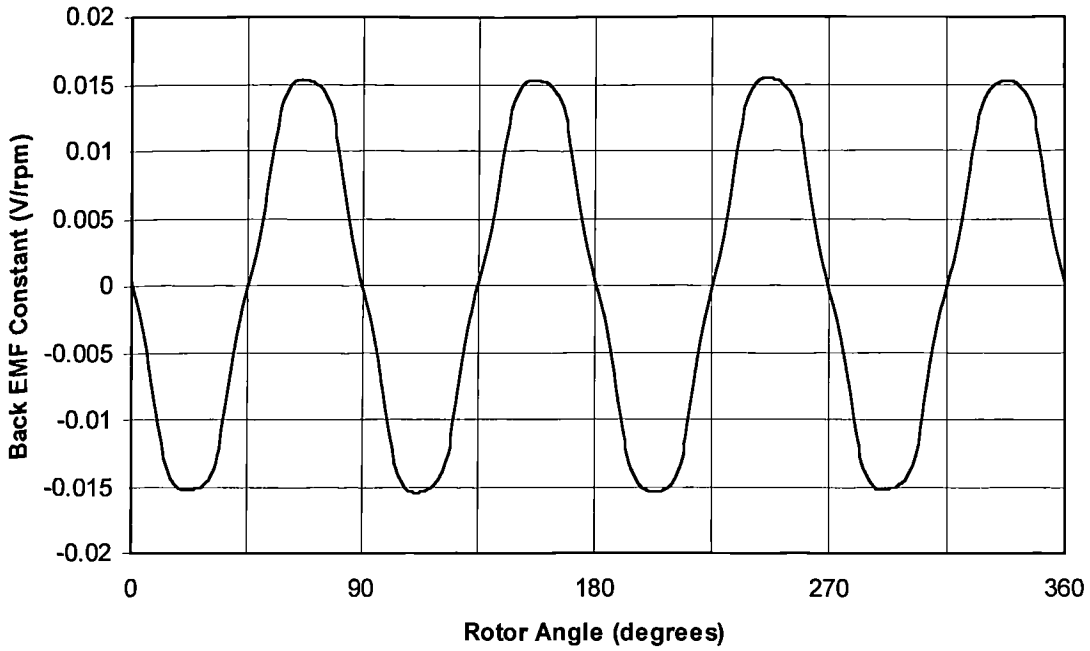

Figure 11 : Back-EMF Data Measured at 4000 rpm

The static torque characteristic of the machine was measured using the torque meter on the rig. The resulting torque vs. angle curves, which include the cogging torque, are shown in Figure 12. Note that the cogging torque shown is the total cogging torque for the machine and appears more severe than is really the case, as it is superimposed on the torque for a single phase. The peak cogging torque is 0.27 Nm, 2.3% of the rated torque

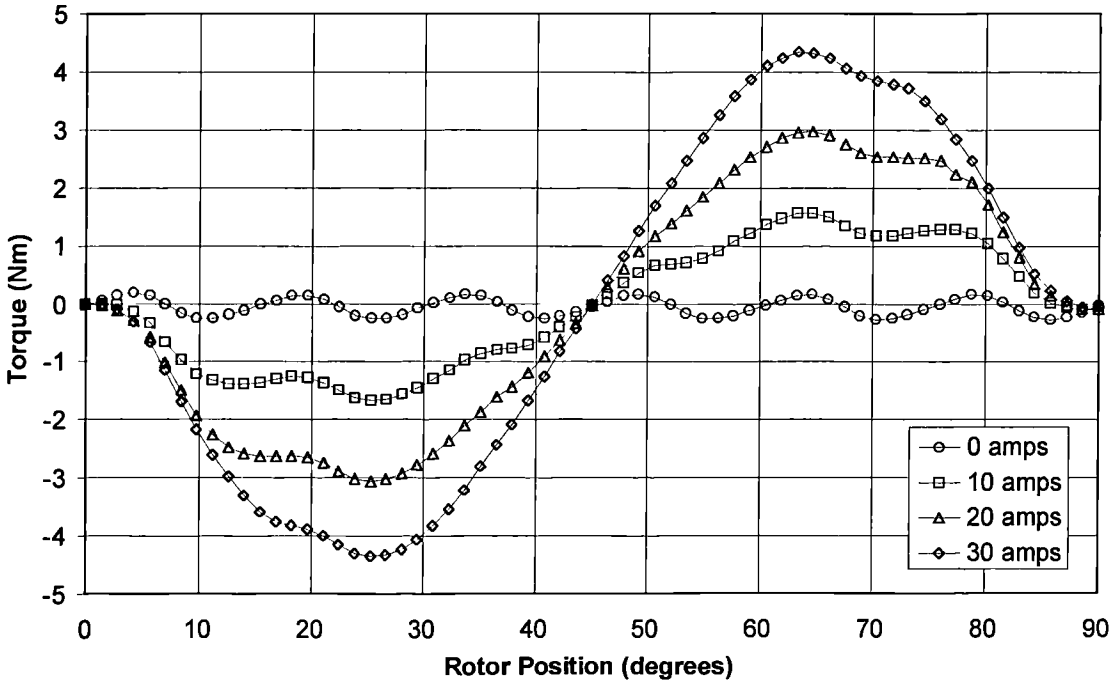

Figure 12 : Static Torque

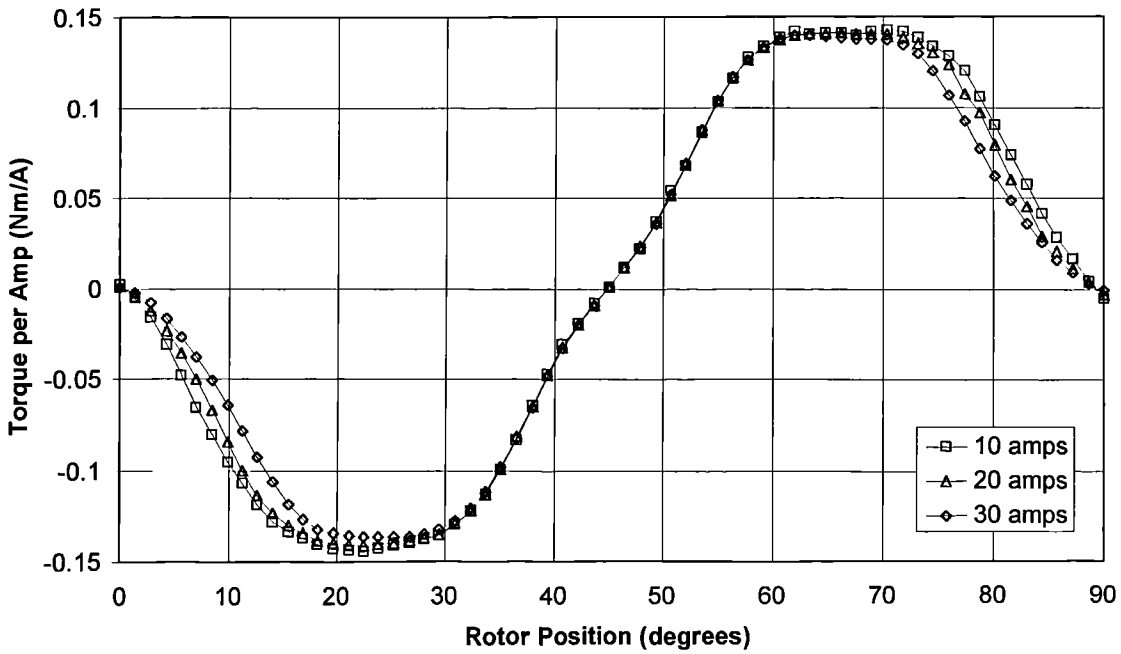

Figure 13 : Static Torque per Amp

for the drive.

Figure 13 shows the torque per unit current vs. angle for a range of phase currents. Note that the cogging torque has been removed from the result in Figure 13.

2.6 Penalties for the Fault Tolerant Design.

This section will examine the penalties incurred by the demonstrator design, in terms of increased parts count and additional material requirement. These two penalties translate to additional mass, volume and cost and reduced MTBF. The penalties incurred fall into two groups, those resulting from the unusual drive design and those associated with redundancy. The following sections will examine the effect of the machine and power converter design on the material requirements. Wherever the increase in parts count and material requirement is presented as a factor, this is relative to a well designed 3 phase PMSM, with negligible per unit reactance, supplied by a standard six switch inverter.

The effect of redundancy is simply to factor the material requirement and parts count by 6/5, because the machine must produce rated power from 5 phases. This factor is not applied until after the penalty for the unusual drive design has been established. Excluding the redundancy factor until the very last stage will make comparisons with other drive types easier.

2.6.1 Additional Material Requirement in the Machine

Although the machine design is unusual, it still makes good use of the iron and copper. The only penalty is that the amount of stator iron must be increased by an amount $\sqrt{2}$ since the high reactance results in an increase in flux in the stator iron compared to a machine with negligible reactance, Figure 14. Whilst the amount of iron in the tooth tips

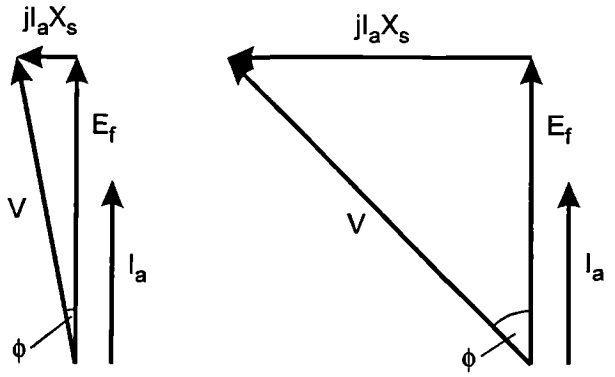

Figure 14 : Operating Points for PMSM Machines with 0.2 p.u .and 1.0 p.u. Inductances

must be increased to provide the high per unit inductance this increase is small compared to the total iron mass (see Appendix A). The increase in stator iron results in an increase in the active mass of approximately 15%, [39].

$$\text{Factor by which the total machine mass is increased} = 1.15 \quad (1)$$

Note that the machine has no significant saliency, so $X_q = X_d$.

2.6.2 Additional Power Switches

The penalties for the power converter are in two forms. Most obviously the number of power switches has been increased from 6 for a three phase star connected system, through 12 for a six phase star connected system, to 24 for a six phase machine with each phase driven from an independent single phase bridge. Thus the demonstrator drive has 4 times the number of power switches of a drive supplied by a six switch, three phase inverter.

This seems like a serious penalty. The gate drive electronics increase by a volume and cost in line with the switch count. However the amount of power silicon increases in line with the volt-ampere, VA, rating of the power converter. Whilst the number of switches has quadrupled, each switch is now smaller in terms of VA rating. The increase in the quantity of power silicon will be calculated in the following section.

2.6.3 Additional Material Requirement in the Power Converter

The material requirement in the power converter is often measured in terms of volt-ampere, VA, rating, thus avoiding the details of component selection. When discussing SRMs, Miller [65] defines the peak volt-ampere rating of a power converter as the product of the power converter peak current capability, the DC link voltage and the number of switching devices. This peak VA rating, based on peak current is more appropriate than a VA rating based on rms current, when sizing transistor based converters.

Consider a three phase, star connected PMSM, supplied by a standard six switch inverter. What is the kVA requirement, assuming a sinusoidal air gap flux, sinusoidal excitation and no significant per unit inductance? For operation with the stator current, I_a, in phase with the phase back-EMF (E_f) and no significant resistance, the peak phase current is given by:

$$\hat{I}_a = \sqrt{2} \cdot I_a \tag{2}$$

The DC link voltage requirement is given by:

$$V_{LINK} = \sqrt{2} \cdot (\sqrt{3} \cdot E_f) \tag{3}$$

The air gap power for three phases, P_{gap}, is simply:

$$P_{gap} = 3 \cdot E_f \cdot I_a \tag{4}$$

The peak power converter VA rating, Q_m, is given by:

$$Q_m = 6 \cdot V_{LINK} \cdot \hat{I}_a = 6 \cdot \sqrt{2} \cdot \sqrt{3} \cdot E_f \cdot \sqrt{2} \cdot I_a = 4 \cdot \sqrt{3} \cdot P_{gap} = 6.93 \cdot P_{gap} \tag{5}$$

Now consider how each feature of the demonstrator design affects the VA rating. The amount of power silicon rises by $2/\sqrt{3}$ because each phase has its own single phase bridge (twice the number of switches but each switch is rated for the phase volts rather than the line volts). The power converter VA rating is unaffected by the increase in phase number as the number of switches is doubled but each phase carries only half the current. The VA rating increases by a factor of $\sqrt{2}$ because each inverter carries reactive power associated with the 1 p.u. reactance (i.e. operation at a power factor of $1/\sqrt{2}$).

$$\text{Change in power silicon VA rating (as a factor)} = \frac{2}{\sqrt{3}} \times \sqrt{2} = 1.63 \tag{6}$$

The above analysis assumes that the machine is operated in vector control, with sinusoidal stator current in phase with the no load back-EMF as shown previously in Figure 14. A proportion of the increased power silicon requirement is associated with the poor power factor. The power factor can be improved by phase advancing the current with respect to the no load back-EMF. The angle through which the current is phase advanced with respect to the back-EMF is referred to as the 'current angle' and given the symbol γ. Now it will be shown, in two stages, that phase advancing the current reduces the kVA / kW requirement in the power converter.

First consider the case where the current is phase advanced with respect to the back-EMF. As phase advance is applied to the current, the load angle, δ, increases and the current must be increased to maintain the required power output. The locus followed by

the current vector given a requirement for constant power is shown in Figure 15. The locus of the voltage vector required to achieve the current is also shown in Figure 15. Clearly as the current is initially increased, the voltage requirement decreases at a higher rate thus reducing the total kVA/kW requirement.

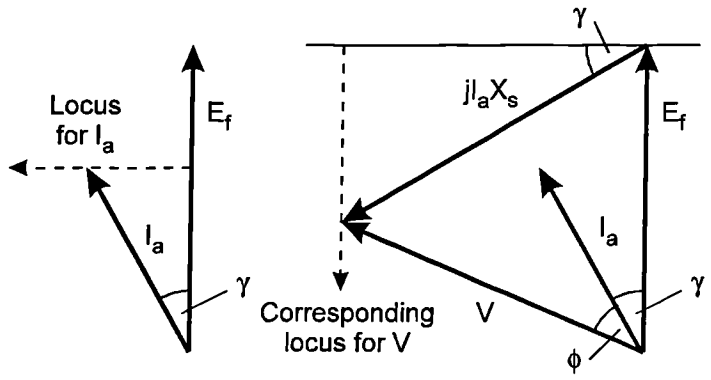

Figure 15 : Increasing Current Angle With Constant Torque, Without Reducing X_s.

With the approach used in Figure 15, the magnitude of I_a must be increased and so must its base p.u. value. Thus the short circuit current will be less than 1.0 p.u. so the machine is excessively reactive. If, whilst increasing the phase advance of I_a, the value of X_s is reduced to keep I_aX_s at 1.0 p.u. then the voltage requirement and thus the kVA / kW requirement will be further reduced. This situation is shown in Figure 16.

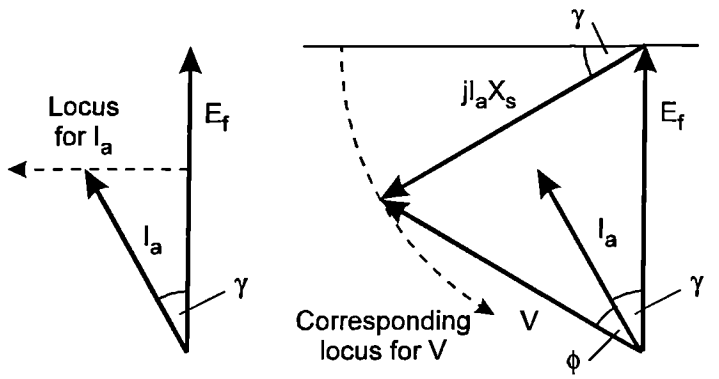

Figure 16 : Increasing Current Angle With Constant Torque, Maintaining I_aX_s at 1.0 p.u.

The variation in current and voltage requirement resulting from applying phase advance to the current, as described in Figure 16, is shown in Figure 17. These graphs show that to obtain a given power output with minimum VA requirement it is best to operate the machine with the current vector phase advanced with respect to the back-EMF vector, i.e. with a current angle greater than zero.

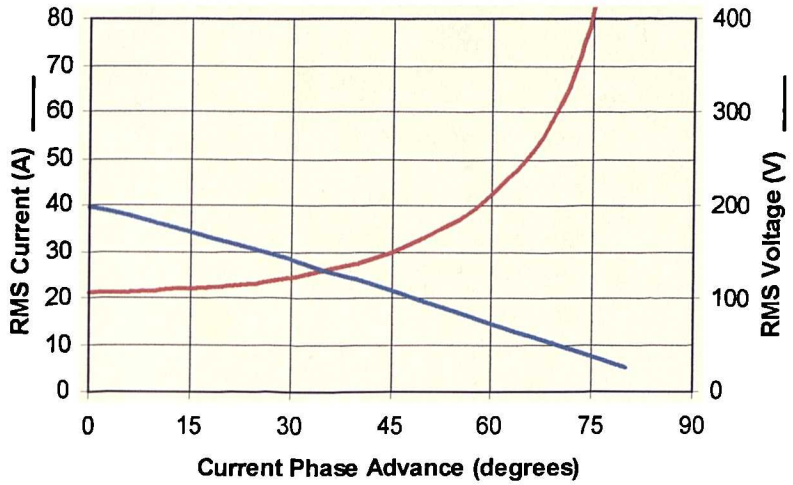

(a) Phase Current and Voltage for a Given Torque

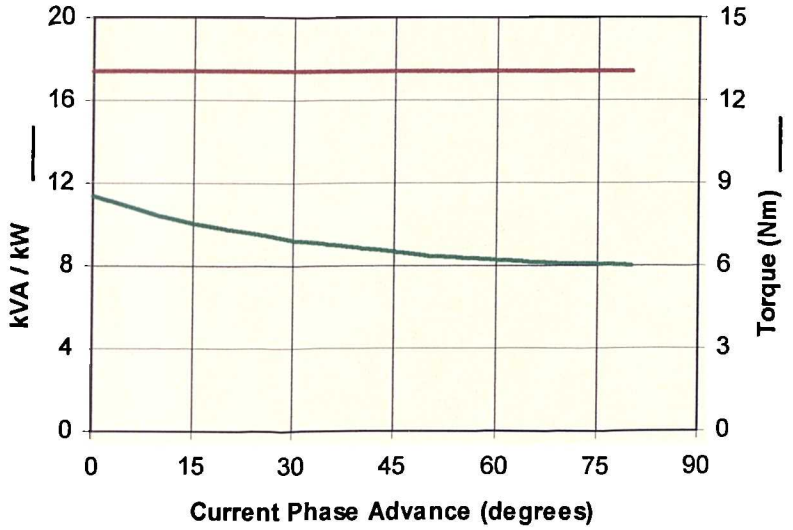

(b) kVA/kW Requirement for a Given Torque

Figure 17 : Best Operating Point for the Demonstrator Machine.

The graphs in Figure 17 show that the power converter kVA requirement falls as more phase advance is applied to the current. This is a simply because the increasing current phase advance results in a reducing power factor. Unfortunately the reduction in power converter size is at the price of an increased machine mass.

The magnetic loading of the machine decreases from $\gamma=0°$ until $\gamma=30°$ at which point the stator flux equals the open circuit flux. Increasing γ beyond 30° will result in the stator flux reducing below the open circuit flux level so no further reduction in magnetic loading is possible beyond $\gamma=30°$. The electric loading increases from $\gamma=0°$ to $\gamma=30°$, but unlike the magnetic loading it continues to increase more and more rapidly beyond $\gamma=30°$. Thus it seems reasonable to design the drive for operation at $\gamma=30°$. At this angle

the current is increased by 15% but the voltage required decreases by 29%. The resulting converter VA rating decreases by 18% to 9.29 kVA/kW.

2.6.4 Summary of Material Penalty for Fault Tolerance

The material penalties incurred in the demonstrator design are summarised in Figure 18. Compared to a PMSM machine without any fault tolerance (first columns on the figure), the machine mass is increased by a factor of 1.15 and the power converter kVA increases by 1.34. This is a significant penalty but can be put into perspective by comparison with the SRM. The machine is approximately half the mass of a similarly rated SRM. The power converter requirement is 9.29 kVA/kW as compared to figures of 9.5 kVA/kW for a typical SRM [66]. Finally the material requirement increase associated with the 5+1 redundancy should be considered. The last columns in Figure 18 show the effect of increasing the total material requirement by 6/5 to provide the 5+1 redundancy.

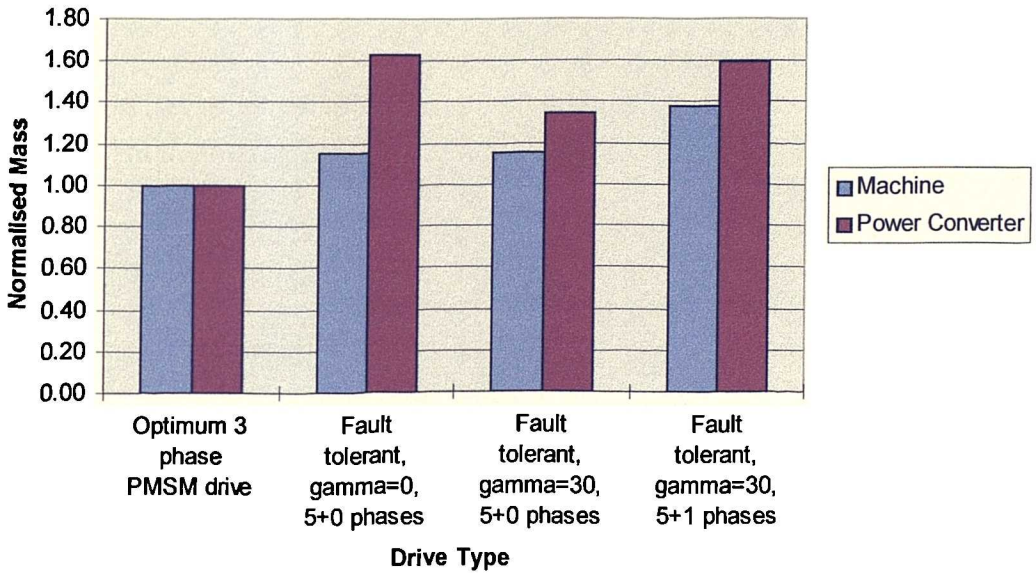

Figure 18 : Material Penalty for Fault Tolerance

2.7 Conclusion

This chapter has proposed the partitioning of a drive into independent phase units. It has been shown that the use of 5+1 redundancy provides an improvement in system failure rate of 4 orders of magnitude, making the difficult MTBF requirements of aerospace applications achievable.

A PM synchronous machine has been chosen over other machine types for its higher power density. Effective partitioning of the drive requires that each phase of the machine

be fed from an independent power converter. Certain key requirements have been identified, to enable continued operation of a faulted drive. These may be summarised as:

- The drive must have effective isolation between phases.
- Steady state fault currents must be limited to rated value.

A novel design of PM synchronous machine drive has been presented which provides the necessary isolation between phases. The resulting machine has a level of inherent fault tolerance, arguably greater than that of a typical SRM. Fault currents are limited by designing the machine with a 1 p.u. inductance, through the choice of an H-bridge power converter topology and through the selection of IGBTs for the power converter switching elements.

A demonstrator drive has been built to the proposed design. A number of important machine characteristics have been measured and these show that the design has been successful in achieving effective phase to phase isolation. The short circuit current for the machine, given in Table 3, is close to rated current demonstrating that a per unit reactance close to 1.0 has been achieved.

Finally the penalties incurred in terms of additional material requirement have been examined. When compared to a standard 3 phase PMSM drive, the demonstrator drive will have 38% more active machine mass and a 59% higher power converter VA rating. The resulting system failure rate is approximately 10^4 times lower, which is clearly an enormous improvement. In spite of the increase in active mass, it is thought that this PMSM still delivers more specific power than its nearest competitor, the SRM, at this power level. Even with the increase in power converter VA rating, the VA rating is still similar to that required to drive a SRM of the same power rating.

Chapter 3

CURRENT CONTROL

3.1 Overview

This chapter introduces some of the methods of current control used in modern drives. The merits of each of these techniques are examined. The PI control method is used initially to control the current in the demonstrator machine. The performance of this simple PI loop is demonstrated. Various improvements to the PI current controller are investigated and the change in performance examined.

A new current control scheme based on a flux linkage model of the machine is introduced. Factors affecting the accuracy of the model based control are given and a method for estimating the impact of these factors on the accuracy of the current control is presented. The performance of the model based control is demonstrated and compared with the performance of the PI current controller. Finally implications of the model based control scheme for fault detection are introduced.

3.2 Established Current Control Techniques

The objective of the current controller is to provide a voltage demand for the machine terminals which will force the phase current to follow the current reference. The actual current should follow fast transients in the reference without excessive overshoot. Any oscillation should be strongly damped and the settling time should be short. Finally, the steady state error should be minimised. If it were possible to design a perfect discrete time controller, its response would be deadbeat. A current controller with deadbeat response will drive the phase current to the demanded level in one sampling interval with no overshoot, oscillation or steady state error.

3.2.1 Hysteresis Control verses PI Control

The methods commonly used for controlling machine current are hysteresis controllers (also referred to as bang-bang controllers) and PI (or PID) controllers.

Hysteresis controllers are unconditionally stable, simple to implement and give a controlled level of current ripple. They are by nature continuous time controllers. Their implementation will be in largely analogue electronics. Unfortunately the switching speed is not controlled. There will always be a trade off between excessive current ripple on one hand and poor noise immunity and high switching frequency, resulting in excessive switching loss on the other.

PID controllers have been adopted to control a wide range of systems because they are robust and easy to tune. When applied to current control the output of the PID controller is pulse width modulated to provide a fixed frequency switching signal to drive the power converter. The pulse width modulation (PWM) is at a fixed frequency. The current ripple is limited by appropriate selection of this PWM frequency. These controllers may be implemented in continuous time or discrete time. If the controller is to be implemented in analogue components a continuous time controller will be chosen. The use of a digital implementation will force the use of a discrete time PID controller.

In general, controllers implemented in discrete time will be less stable than their continuous time counterparts because of the time delay introduced by the sampling period. However when Elliott et al [67] compare a hysteresis current controller to a continuous time PID current controller, they observe that the analogue current feedback signal used by the PID controller must be filtered to remove higher frequency components. These components are introduced by the switching power converter. Filtering is necessary to prevent the PID controller from responding to the components of current at the switching frequency and thus cause additional switching transitions in the PWM output. It is clear that since the filter must remove components at the switching frequency it must introduce an additional delay into the control loop equivalent to sampling once per switching cycle. Thus the stability advantage of the continuous time controller is lost when the controller drives a PWM output.

Elliott et al [67] conclude that if dynamic performance considerations rule then a hysteresis controller must be chosen. Whilst this is arguably true, the increased current ripple introduced by a hysteresis controller is not fully apparent, as the current plots presented have been filtered to remove higher frequencies.

In work examining high performance current control directed towards SRMs, Barrass [68] provides a detailed comparison of hysteresis and PI controllers. It appears that the comparison is a little hard on hysteresis control. At all points in the current waveform the switching frequency for hysteresis control seems slower than for the PI controller, suggesting that the hysteresis band could be narrower. This would reduce the current ripple for hysteresis control. Barrass goes on to demonstrate that the digital hardware in which he implements the PI controller also allows the implementation of a more sophisticated discrete time controller, termed a flux controller. This flux controller gives performance clearly superior to either the hysteresis controller or the PI controller and will be discussed in more detail later.

3.2.2 Torque Control versus Current Control

When examining current control, it is important to be aware of progress in the area of direct torque control. The term 'direct torque control' is applied to PMSM control schemes which seek to eliminate the current loop and/or give instantaneous torque control. This is achieved by measuring the phase currents and measuring or estimating (implicitly or otherwise) the rotor position and thus calculating the air gap torque. The calculated torque is used as a feedback term to close the torque control loop.

A paper recently presented by Rahman et al [69] generated some interest as it appeared that the direct torque control scheme proposed for a PMSM was able to track the changes in rotor position without the aid of any shaft encoder. However this scheme is in fact implicitly estimating the change in rotor position from the phase currents and the history of applied voltage. Instantaneous torque control is clearly important in applications where smooth torque production is a requirement. Outside these applications it seems a little arbitrary as to whether the torque demand is converted into a current demand and phase current controlled or whether the measured current is converted into a torque estimate and the machine torque controlled. Perhaps the most important thing to recognise is that, with the processing power currently available, it is possible to perform the calculations required to transform current and voltage to torque and flux and vice versa in real time.

3.2.3 The Choice of Controller Type for this Application

For this application current rather than torque was chosen as the quantity to be controlled, since it is easier to diagnose faults from the measured current than it is from calculated air gap torque.

The possibility of realising a more sophisticated current controller was the principle factor which led to the decision to implement the current controller in software running on a microprocessor. Later it will be seen that the use of a microprocessor to close the current loop is a significant advantage in providing fault tolerance, allowing fault detection schemes and post fault control strategies to be easily integrated into the current controller. The resulting discrete time controller could have a hysteresis output but in practice this is not common as a much higher sampling frequency (and faster current loop) is required compared to a controller with a PWM output. The use of a PWM output also has the advantage of providing a controlled switching frequency.

Given the use of a discrete time controller with a PWM output, a PI control law seemed to be the most appropriate starting point for the comparative study presented in the following sections.

3.3 Constraints of the Hardware

Before comparing various methods of current control it is important to understand how the current controller fits into the system. Examination of the interface between the controller and the system will reveal any constraints placed on the current controller by the system.

The fault tolerant requirement results in each machine phase being driven by its own single phase bridge. If the independence between phases is to be maintained then this naturally results in a separate current controller for each phase. One phase of the controller is shown in Figure 19.

A measurement of the current delivered to the phase is fed back to the current controller. The current controller also receives a measurement of the rotor angle, the DC link volts and a peak current demand, but these signals are common to each of the six single phase controllers. Each of the inputs and outputs of the current controller is sampled and this sampling is triggered by a single sampling signal shared by all phases. The sampling

signal is generated by the PWM unit at the start of each PWM cycle. Thus each loop runs at the PWM frequency, 10 kHz. All six of the phase current controllers run on a single Texas Instruments TMS320C31 floating point DSP and the PWM unit is realised in digital hardware using a field programmable gate array (FPGA).

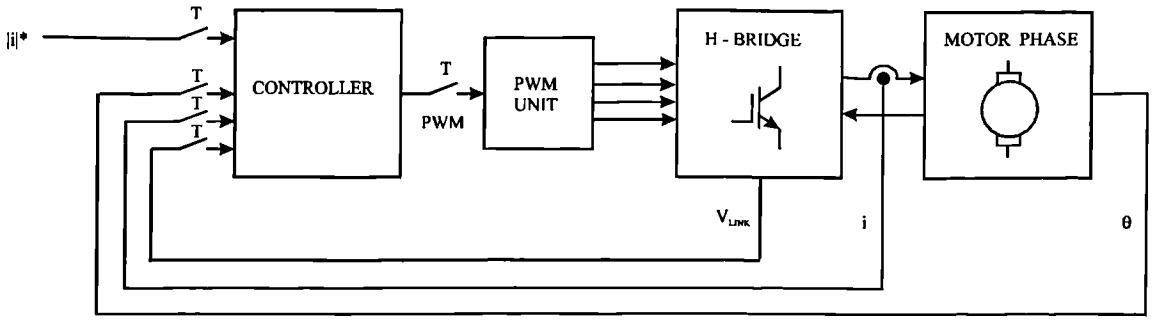

Figure 19 : Current Controller Structure

3.3.1 Current Controller Timing

The delays introduced by a digital, discrete time controller are a serious limiting factor on the performance and stability of the resulting closed loop system. Many texts discussing discrete time controllers assume that the time taken to sample the feedback variables, to perform the control calculations and to apply the resulting drive to the system is negligible when compared to the system sampling time. This is not so here. The time taken to perform analogue to digital conversion of the phase currents and to perform the calculations in the current control loop amount to the majority of a sampling interval.

The decision was made to synchronise both the sampling of the currents and the transfer of new PWM demands to the PWM unit to the start of the PWM cycle as shown in Figure 19. This results in a delay of exactly one sample interval between sampling the currents and the point when the values sampled start to affect the PWM output. A further delay of one sampling interval is experienced before the PWM applied is next updated. Thus the total delay through the controller and plant is 2 sampling periods. This situation is illustrated in Figure 20, which shows the progress of three sets of sampled data, each of which is represented by a different colour. Note that each control cycle overlaps with the previous one, so that the analogue to digital conversion for the second cycle shown in the figure, starts at the 100 μs point.

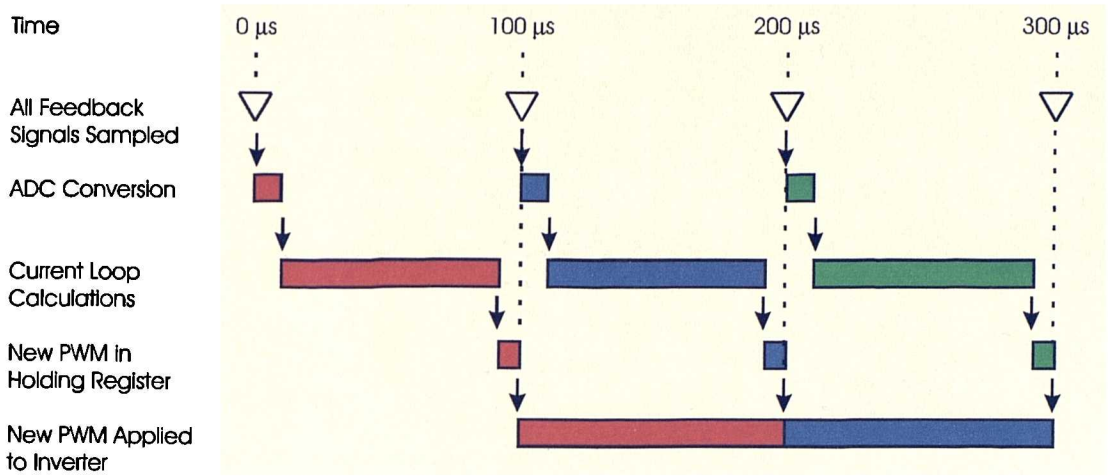

Figure 20 : Current Loop Timing

Monmasson et al [70] define the equivalent small time constant, τ_S, of a discrete time controller as:

$$\tau_S = \tau_O + Te/2 + Tc \quad (7)$$

Where:

- τ_0 is the sum of the small time constants due to the process itself, the feedback transducers and the delay of the analogue to digital conversion.
- Te is the sampling period.
- Tc is the computation time of the controller.

Increasing any of the terms in the expression will increase τ_S, thus reducing system stability. However, second for second, an increase in τ_0 or Tc will be twice as detrimental to system stability as an increase in Te.

The time τ_0 + Tc is the result of latency in the system, and may be referred to as the latency period. The difference between the latency period and the sampling period is as follows:

- The sampling period is the time in which the system state may change between successive measurement points and represents the potential for future errors.
- The latency period is the period before each sampling instant in which the system state may have changed without being reflected in the measured state. Any

latency degrades the extent to which the present state of the system can be measured, which is why it has such a serious effect on system stability.

From Figure 20 it can be seen that in this system τ_0 + Tc = 100 μs and Te = 100 μs. Later it will be shown that latency period can be effectively converted into sampling time, thus improving controller stability.

3.4 Current Control Using a PI Current Loop

The PI current controller which has been implemented is best described in terms of a block diagram, see Figure 21. Central to the diagram is the PI controller block, comprising an integrator with gain, Ki, plus a proportional term with gain, Kp. The output of the PI block is in volts which are converted into PWM demand by dividing by the nominal DC link voltage of 315.0 V. Following the forward path, the PWM is limited to ±1.0 and then a unit delay is applied to represent the delay of 1 sample time through the controller. The other blocks in the figure are related to two improvements made to the standard PI controller.

The first improvement to the current controller is to apply phase advance to the rotor position used to calculate the current reference. First consider why this is necessary. In general, the current reference is a function of the torque required and the rotor angle. In this current controller, the variation in current demand with rotor angle is described by a current profile lookup table. The torque output is controlled by scaling the current profile. Thus the scaling factor, referred to as the current magnitude demand, $|i|_k^*$,

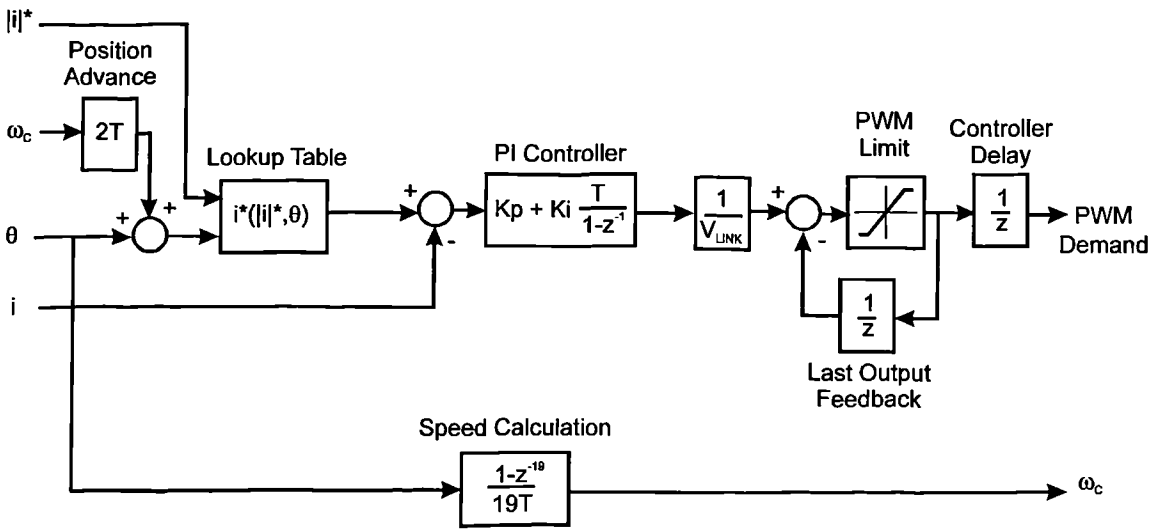

Figure 21 : PI Current Controller Structure

replaces the torque demand. At time t_k, the current reference, i_k^*, is a function of the current magnitude demand, $|i|_k^*$, and the rotor position, θ_k, (8).

$$i_k^* = f(|i|_k^*, \theta_k) \tag{8}$$

The current controller must regulate current in an 8 pole machine operating at 13000 rpm resulting in an electrical cycle period of 1154 μs. There is a delay through the controller of 2 sample intervals, i.e. 200 μs. Thus the delay through the controller and PWM unit amounts to 17% of an electrical cycle or 62 electrical degrees. The phase relationship between the phase current and the rotor angle is crucial to torque production in the machine and a lag of 90 degrees between the current reference and the phase current will result in zero net torque. Thus a phase lag of 62 degrees is very serious.

Provided that the rotor speed and current magnitude demand remain constant, the future current reference can be calculated. This requires knowledge of the rotor speed. The speed is calculated from the change in the rotor angle over the last 19 sampling intervals as shown in the speed calculation block of Figure 21. The time delay through the controller means that the phase current and rotor position measurements at time t_k will affect the current at time $t_{k+2} = t_k+2T$. The phase lag introduced by the controller may be cancelled out by predicting the rotor position at time t_{k+2}. This predicted rotor position is used to look up the current reference, i_{k+2}^*, used in the current controller, (9), and this is achieved by the position advance block in Figure 21.

$$i_{k+2}^* = f(|i|_k^*, \theta_k + \omega \cdot 2T) \tag{9}$$

Note that it is only possible to compensate for the change in current demand with rotor angle in this way because the change in rotor angle is predictable. In general it is not possible to predict changes in the torque demand, so changes in the phase current resulting from changes in the current magnitude reference will always lag the reference by the delay introduced by the controller.

The second improvement to the current controller is to apply 'last output feedback'. Controllers normally operate without last output feedback. The resulting controller measures the current and calculates the PWM demand to drive the actual current towards the current demand. However in a sampled system there is a delay between

sampling the current and the application of the resulting PWM signal to the inverter, τ_0 + Tc. During this period the current in the machine may have changed substantially, 55% of peak rated current for sinusoidal excitation in this drive. Thus the delay between sampling the current and applying the resulting PWM signal to the inverter has a seriously detrimental effect on the stability of the system. This is predicted by equation (7). Clearly if the output of the controller during the period τ_0 + Tc is known and its effect on the system can be accounted for, then the system stability may be improved.

Here the change in the current during the time τ_0 + Tc results from the application of the PWM demand calculated by the previous iteration of the current loop. Thus the effect of the PWM applied over the present sampling interval is accounted for by subtracting it from the PWM demand calculated by the controller, to give the PWM required over the next interval. This is referred to here as last output feedback.

The effect of last output feedback here is to convert latency period, τ_0 + Tc, into an extended sampling period, Te. Thus τ_0 + Tc is reduced from 100 μs to 0 μs and Te is increased from 100 μs to 200 μs. Thus the equivalent small time constant of the controller, τ_S, is reduced from 150 μs to 100 μs with a corresponding improvement in system stability.

Lee and Kang [71] introduce a technique for controller delay compensation very similar to that used here. However Lee and Kang subtract the calculated effect of the last output from the current error, i.e. from the controller input, rather than simply subtracting the last controller output from the present controller output. It is possible that the presence of an integrator in the Lee and Kang controller necessitates this more complex approach.

3.4.1 Current Profiles Used to Evaluate Current Loop Performance

During the performance comparisons which follow two current profiles have been used. The first is a square wave with a 120 degree conduction period. This is the waveform most commonly used to drive 3 phase brushless DC machines. It is used here to demonstrate the current controller's step response. In each step response the machine is rotating very slowly so the back-EMF is negligible. During this low speed step response, the slow rotation of the machine means that only a small section of one electrical cycle is shown. The part of the cycle shown is chosen to include a positive current step.

The second current profile used is identical in shape to the phase back-EMF, Figure 11. This may be seen as the ideal waveform with which to drive the machine, as this waveform gives the maximum torque for a given rms current and thus minimises copper loss [72]. The back-EMF is nearly sinusoidal and this current profile is used to assess the performance of the current control with the machine running at full speed, when the back-EMF and electrical frequency are both at a maximum.

The performance of the different current controllers covered in this chapter is compared on the basis of oscillograms showing results for either a low speed step test or for high speed operation with a back-EMF shaped current reference. The 3 waveforms in each of these oscillograms represent the same quantities in every case. An example of an oscillogram from each test can be seen in Figure 22. In each oscillogram the top trace is the actual current measured by a current probe. Below this the current sampled by the controller, i_k, is super imposed on the sampled current demand, i_k^*. These values are output from the controller via DACs. There is a delay of approximately 90 μs between the instant when the phase current and current demand are sampled and the instant when these sampled values are written to the DACs. Therefore the values output by the controller lag the actual current by a short period.

3.4.2 Performance of a PI Current Loop

There are many techniques used to tune PI controllers. The object of all these methods is the same, to select values for the proportional and integral gains to optimise the controller performance. Initial values for Kp and Ki were derived by applying the 2nd method proposed by Ziegler-Nichols [73], also referred to as the 'ultimate sensitivity method'. The gains this yielded gave excessive overshoot on the step response. The step response was improved by halving the integral gain, giving values Kp = 11.1 VA^{-1} and Ki = 14800 $VA^{-1}s^{-1}$. Figure 22(a) shows the response of the phase current to a 15 amp step in current demand as the rotor passes a commutation point, 150 degrees before the aligned position. The current controller is clearly well tuned as the response to the transient is fast without the over-shoot becoming excessive. Note that the lag normally present between the current demand and the sampled current measurement in a discrete time controller is completely eliminated here by using a phase current reference which has been phase advanced by the 2 sample delay through the controller.

ORIGINAL IN COLOUR

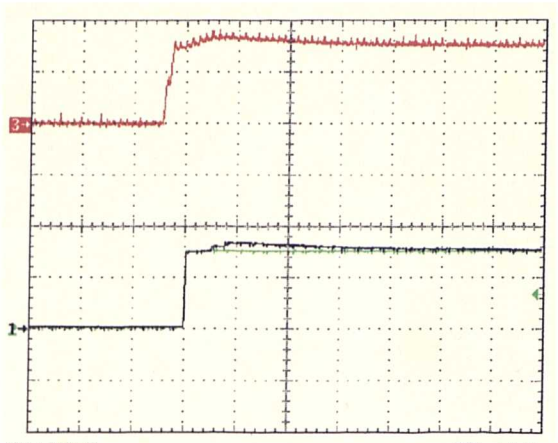

(a) Low Speed Step Response

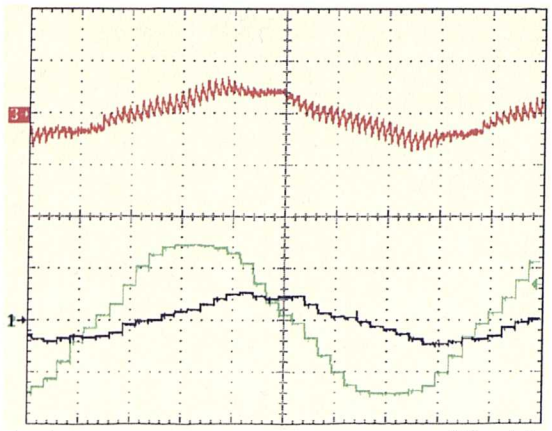

(b) High Speed Test Response

Figure 22 : Performance of PI Current Control

Figure 22(b) shows the response of the phase current to a back-EMF shaped current reference at higher speed. At this speed the gain in the controller is no longer sufficient to overcome the back-EMF. As a result the actual current is lagging the current reference by 90 degrees. Whilst the amplitude of the actual current is also less than required it is the phase lag which is most serious, as a 90 degree phase lag will result in zero net torque from the phase. Note that the oscillogram in Figure 22(b) shows operation at approximately 30% of full speed. The phase lag in the current prevented any higher speed being reached.

The drive has been designed so that there will always be sufficient DC link volts available to overcome the back-EMF and to drive rated current with a sinusoidal profile. The difficulty is that the gain required in the controller to overcome the back-EMF and force the current to follow a sinusoidal reference results in instability. The difficulties introduced by the back-EMF may be overcome to some extent by applying a speed dependent phase advance to the current reference. This is not an elegant solution as it accepts that the current control loop is not stiff enough.

The back-EMF may be viewed as a disturbance in the system. However, it is an unusual disturbance in the sense that it is predictable. Whilst the machine is operating outside saturation, the back-EMF at a given rotor angle is proportional to speed. The profile of back-EMF with rotor position may be held in the controller as a lookup table and used to estimate the back-EMF over the control period. This estimate may be added onto the

ORIGINAL IN COLOUR

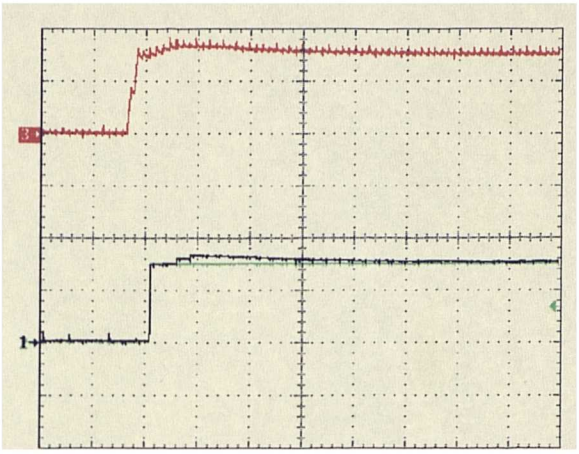

(a) Low Speed Step Response

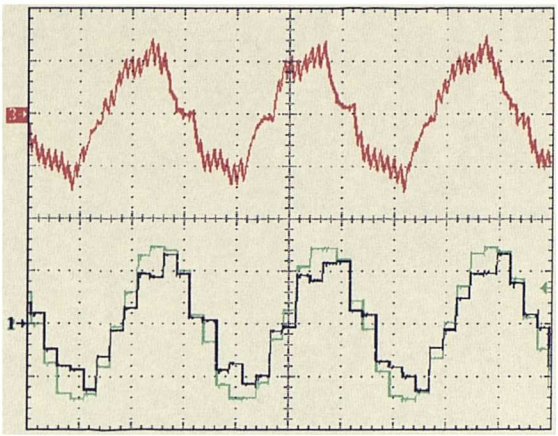

(b) High Speed Test Response

Figure 23 : Performance of PI Current Control with Back-EMF Feed Forward

output of the PI controller block in the controller diagram in Figure 21. The benefits of using the estimated back-EMF as a feed forward term in the current controller have been demonstrated by several authors, including Champenois et al [74].

The performance improvement gained by adding back-EMF feed forward into the PI controller is demonstrated by the results in Figure 23. The step response is unchanged as the controller gains have not been altered. The high speed test shows that back-EMF feed forward has resulted in much faster tracking of the current reference and the drive is now capable of delivering useful torque, at full speed.

3.5 Current Control Using a Flux Linkage Model

In the last section, the improvement in the current control resulting from including knowledge of the back-EMF in the controller was demonstrated. If instead of just knowing the back-EMF it were possible to completely characterise the plant then a further improvement in the quality of current control should be possible. In principle all that is required is a model of the machine which relates phase current to applied voltage. If a valid relationship could be established then it would be possible to calculate the exact voltage necessary to bring about the desired change in current. This should lead to dead beat control at all operating points.

For this machine there is little mutual coupling between phases so the voltage appearing at the terminals of a phase may be related to the current in that phase by:-

$$v = i \cdot R + L \cdot \frac{di}{dt} + e \tag{10}$$

In this model the phase resistance and phase inductance are considered to be constants. The back-EMF is a function of rotor speed and angle. This model will break down as the phase is driven into saturation. However by using the more general form of this relationship the effects of saturation are embodied in the model:-

$$v = i \cdot R + \frac{d\psi}{dt} \tag{11}$$

In general the flux linkage, ψ, will be a function of the rotor position and the current in each phase [75]. However in this machine each phase is magnetically isolated from all other phases so the flux linking a phase is a function of rotor position and the current in that phase alone. Thus expression (11), above, relates the terminal volts for a phase to the current in that phase and the rotor position. If this expression is to be used to model a machine phase then it will be necessary to characterise the relationship between phase current, rotor angle and flux linkage. This is the subject of the next section.

3.5.1 Acquisition of Flux Linkage Data for a Phase.

The flux linkage in a permanent magnet machine has 2 components. These are the magnet flux linking the phase and the flux resulting from current in the phase which links that phase. It is simplest to measure the flux linkage in 2 stages.

Measurement of Magnet Flux Linking a Phase

The change in magnet flux linking the phase gives the back-EMF for the phase. It is possible to work backwards, from the back-EMF profile for a phase, to the magnet flux linking that phase. This is achieved by integrating the back-EMF, so the magnet flux linking the phase, $\psi_f(\theta)$, is given by:

$$\psi_f(\theta) = \int_0^{\theta} e(\theta)\, d\theta + \psi_f(0) \tag{12}$$

The value of magnet flux at the starting position, $\psi_f(0)$, is not known. However it is known that the magnet flux linking a phase averages to zero over an electrical cycle:

$$\int_0^{2\pi} \psi_f(\theta) d\theta = 0 \tag{13}$$

Using equations (12) and (13), it is a simple matter to obtain $\psi_f(\theta)$ from $E(\theta)$.

Measurement of Flux Linking a Phase due to Current in that Phase

The flux linkage associated with a given phase current can be calculated from the time integral of the voltage required to drive the current from zero to the level of interest, with the rotor stationary. The total flux linking the phase at this point is the sum of the magnet flux linkage and the flux linkage associated with the current:

$$\psi(i,\theta) = \psi_f(\theta) + \int v - i \cdot R \, dt \tag{14}$$

The total flux linkage characteristic, measured on the demonstrator machine, is shown in Figure 24. The effect of saturation is clearly visible at the point on each curve where the flux linkage per unit current begins to increase more slowly. The process of gathering flux linkage data for the demonstrator machine is described in more detail in Appendix E.

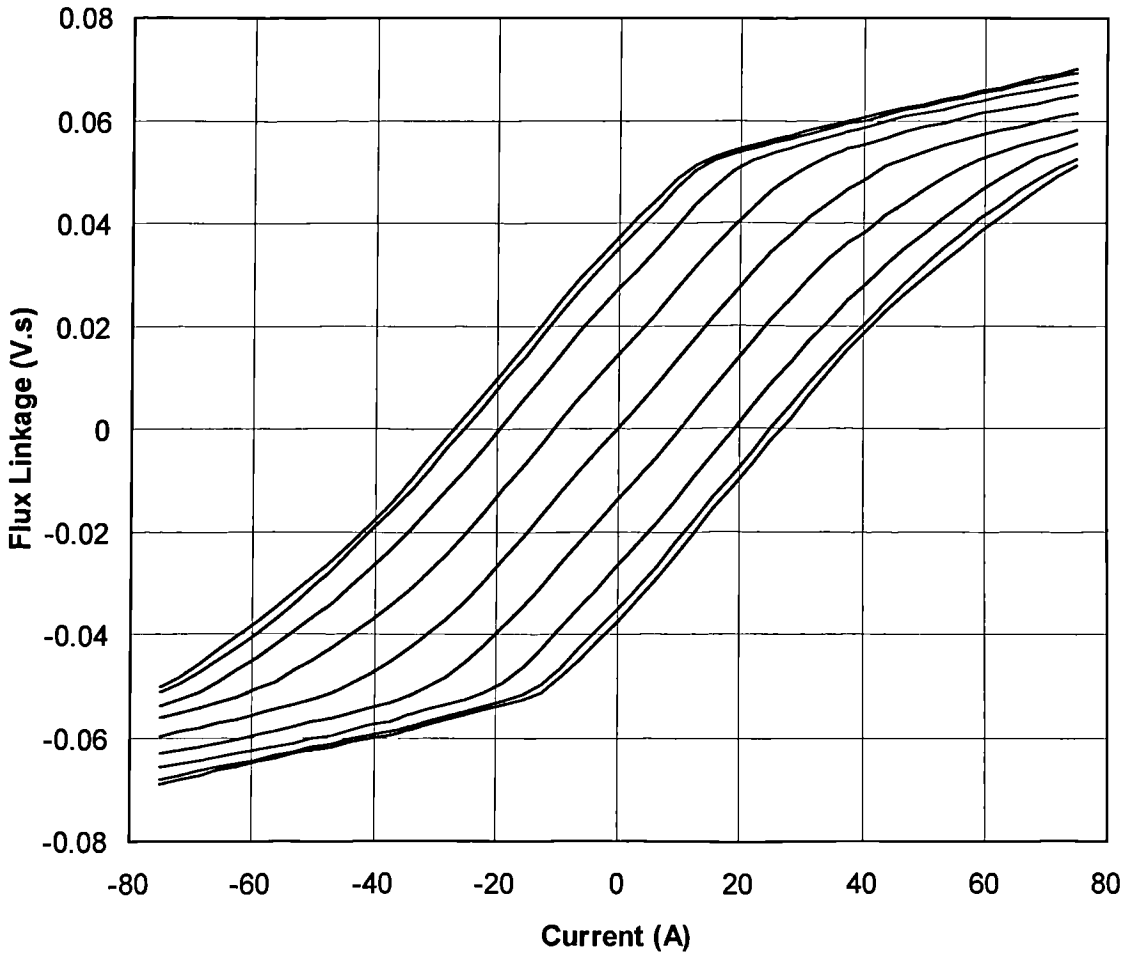

Each curve represents one rotor position from 0 degrees (top curve) to 180 degrees (bottom curve) in steps of 22.5 degrees.

Figure 24 : Flux Linkage Data for Demonstrator PM Machine

3.5.2 Controller Realisation

Barrass [68], [76], shows how flux linkage data may be used to implement a high performance current control on an SRM. Barrass calls the control scheme a flux controller. However, as the controller detailed in the following sections is used to control current it is more correct to refer to it as 'a current controller based on flux linkage data', or more simply as 'flux linkage model based current control'.

The relationship between applied terminal volts and flux linkage in the machine is given by (11). The iR component of the expression is the portion of the applied volts which are dropped across the winding resistance. For machine sizes above 1 kW the resistive drop is typically less than 0.05 p.u. If the resistive drop is neglected the expression may be rewritten as:

$$\psi(t) = \int_0^t v(t)dt + \psi(0) \tag{15}$$

This expression and the flux linkage data form the basis of the current controller as follows. The flux linkage characteristic for the machine shown in Figure 24 is held within the controller as a lookup table. In each control cycle the flux linkage corresponding to the present current and rotor angle is derived from the table. Then the flux linkage corresponding to the rotor angle and current demand predicted at time t is derived from the table. The difference between the present flux linkage and the flux linkage at time t is equal to the volt seconds required to bring about the desired change in current. Finally the volt second demand is converted to a PWM demand. The resulting flux controller is

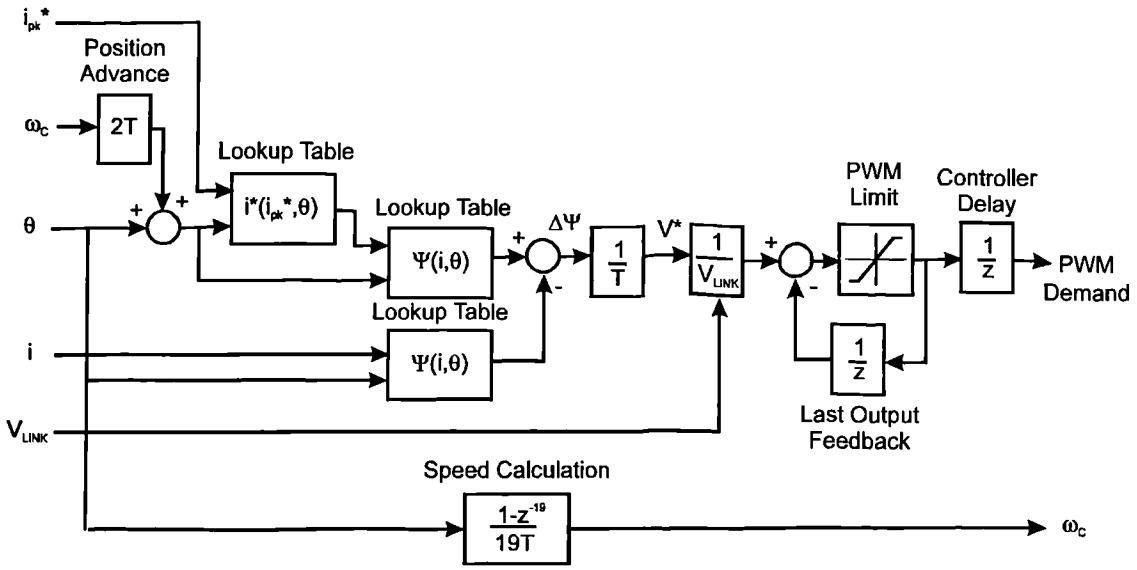

Figure 25 : Current Controller Structure

shown in block diagram form in Figure 25.

Note that the control structure embodies the speed calculation, current reference phase advance and last output feedback elements discussed in the section on PI control. The positive input to the summing junction, associated with the last output feedback block, is the PWM required to achieve the demanded change in current from time t_k to t_{k+2}. The last output feedback subtracts the PWM applied in the period t_k to t_{k+1} to give the PWM requirement for the period t_{k+1} to t_{k+2}, i.e. the required PWM demand for the following PWM cycle. Note that the time t_k to t_{k+1} is the latency period, τ_0+Tc, and the time t_k to t_{k+2} corresponds to the period τ_0+Tc+Te.

3.5.3 Implications of Volt Seconds Error on Current Error

Before presenting results demonstrating the performance of the model based current controller it may be helpful to understand the factors contributing to the current error in this type of control. Control of the phase current based on an accurate model of the phase should result in a phase current equal to the current demand at the end of each control period. i.e. at each sampling interval the current measured should be equal to the current demanded for that instant. In a healthy drive, any difference may arise for one of 3 reasons:

1. The PWM demand is saturated. This happens when the number of volt seconds required to make the actual current follow the current reference exceeds the number of volt seconds available.

2. The model of the drive held within the controller is inaccurate in some way. e.g. the ψ-i-θ data for the machine is not accurate or changes with speed, the inverter dead time has not been compensated for, the resistance of the phase is significant or the inverter losses have not been accounted for.

3. The actual current, DC link voltage or rotor position measurements are deficient.

It is possible to make an estimate of the current error which will result in a phase for a given error in volts applied as follows:

$$i_{error} = \frac{v_{error} \cdot \Delta t}{L} \tag{16}$$

Similarly an estimate may be made of the phase current error resulting from a given error in switching time as follows:

$$i_{error} = \frac{V_{LINK} \cdot t_{error}}{L} \tag{17}$$

Where Δt is the time period between the current measurement and the completion of the subsequent actions taken using that measurement. In the control scheme here it can be seen that Δt is 2 sampling periods. In this expression, L is the incremental inductance at the operating point. The significance of each of the inaccuracies in the flux linkage model of the machine phase may be compared on the bases of the unsaturated inductance, 1.275 mH, a DC link voltage of 300 V and a sampling interval of 100 μs. This comparison is shown in Table 5.

Error resulting from:	Error in volts or seconds	Resulting current error
Dead time in PWM output	4 us	0.94 A
Device drop in inverter	3.2 V at 30 A	0.50 A at 30 A
iR drop in phase winding	4.7 V at 30 A	0.74 A at 30 A

Table 5 : Current Error Resulting from Features not Included in Flux Linkage Model

The sum of these errors is 2.18 A. If the phase is driven into saturation the incremental inductance will fall, causing an increase in the current error.

3.5.4 Performance of Model Based Control

The performance of the model based current controller was tested in the same way as the PI current controller so the oscillograms in Figure 26 may be directly compared with those obtained using PI control. On the step response there is a steady state error of approximately 2 A which is not present in the PI controller. This steady state error results from the accumulated voltage drop across the power devices and winding resistance, and from the volt-seconds lost in dead time. There is some overshoot in the step response. This is the result of the reduction in the current error due to lower average current and therefore lower resistive losses in the sampling period preceding the overshoot. There may also be an element of error introduced by AC effects in the magnetic circuit. The AC effects include additional copper losses resulting from skin

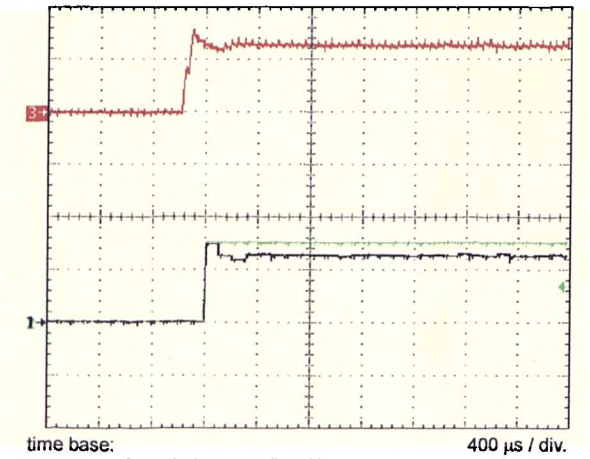

(a) Low Speed Step Response

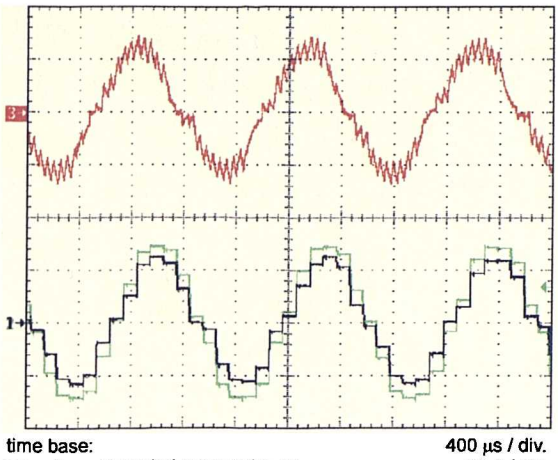

(b) High Speed Test Response

Figure 26 : Performance of Model Based Control

effect, eddy current losses and iron losses, none of which is modelled by the flux linkage data.

Looking at the current waveforms for the full speed test, the steady state error results in a small amplitude error in the sampled current. However the sampled current is perfectly in phase with the current demand. From comparison with the waveforms for PI current control it can be seen that model based current control gives poorer steady state performance but better transient performance.

3.5.5 Including Sources of Error in the Plant Model

It has been shown that a current control scheme based on flux linkage data can provide exceptional transient performance but that there will be some steady state error. It will now be shown that the principle sources of error can also be characterised and that this leads to a further improvement in performance.

iR Drop

The voltage dropped across the winding resistance at any instant is easily calculated. The average voltage dropped across the winding resistance in the two PWM periods from the sample time is given by:

$$v = \left(\frac{i_k + 2 \cdot i_{k+1} + i_{k+2}}{4} \right) \cdot R \qquad (18)$$

Where i_k is the measured current and i_{k+1} and i_{k+2} are the current demands 1 and 2 PWM cycles ahead of the measured current.

Device Volt Drop

The voltage drop across a power device is dependent on the magnitude of the current in that device. This relationship is easy to measure and may be supplied to the controller in the form of a lookup table. At different times in the PWM cycle the current will switch from IGBTs to diodes. It is difficult to calculate the average voltage from the mixture of IGBT and diode conduction. Fortunately the shape of the diode forward voltage drop characteristic and the IGBT on-state voltage drop characteristic is similar. So a single voltage drop characteristic can be defined providing a compromise between IGBT and diode data. The total voltage drop in the inverter can then be calculated from:

$$v = 2 \cdot \left(\frac{Vf_k + 2 \cdot Vf_{k+1} + Vf_{k+2}}{4} \right) \tag{19}$$

Where Vf_k is the forward voltage drop corresponding to the current i_k. The leading '2' in the expression results because the phase current always flows through 2 power devices.

Dead Time

The voltage imposed on the machine terminals during the dead time depends whether the phase current is positive, negative or zero. Further complications are introduced if the current is discontinuous. It may seem most natural to calculate the dead time using a series of conditional statements based on the sign of the current. However it is faster in a DSP to characterise the dead time as a lookup table and calculate the dead time from:

$$v = \frac{V_{LINK}}{2 \cdot T} \cdot \frac{4}{2} \cdot \left(\frac{Td_k + 2 \cdot Td_{k+1} + Td_{k+2}}{4} \right) \tag{20}$$

Where Td_k is the dead time scaled and with polarity assigned according to the current i_k. The '4/2' appears because there are 4 dead time periods in the two PWM cycles but only 1/2 of each dead time period results in a different conduction pattern to that which would be present with zero dead time. The '$V_{LINK}/2T$' factor is required to convert from effective dead time per PWM cycle to the resulting volt seconds error accumulated over 2 PWM cycles.

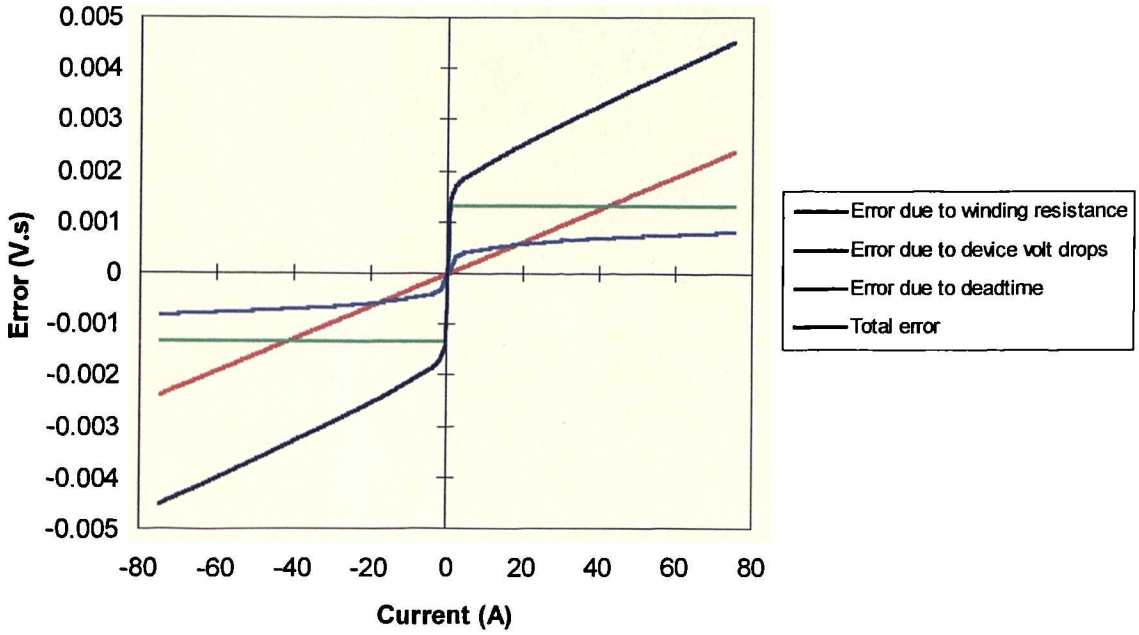

Figure 27 : Combined Volt-Second Error per PWM Cycle

If a constant DC link voltage is assumed for the purposes of calculating dead time then the additional volt seconds required to overcome the dead time and device and iR volt drops may be expressed as a single lookup table, whose form is shown in Figure 27.

3.5.6 Performance of Model Based Current Control with Improved Plant Model

The models described for device voltage drop, iR voltage drop and dead time were incorporated into the flux controller. Figure 28 shows the low speed step response and current waveforms at full speed. Both responses are near perfect.

Thus far PI current control and model based current control have been compared on the basis of peak current demands of 15A. If the peak current demand is increased to 30 A then the phase will just be entering saturation as the rotor approaches the aligned position. The behaviour of the PI controller with back-EMF feed forward and the model based current controller with device drop, iR drop and dead time modelled, may be compared by looking at Figure 29. The poor current waveform given by the PI controller is the result of several effects. The huge overshoots in current occur as the rotor is approaching the aligned position. At this position the effects of saturation reduce the back-EMF. The integrator which helped to achieve the peak current is now resisting the reduction in current which is now required. These two effects result in the controller

ORIGINAL IN COLOUR

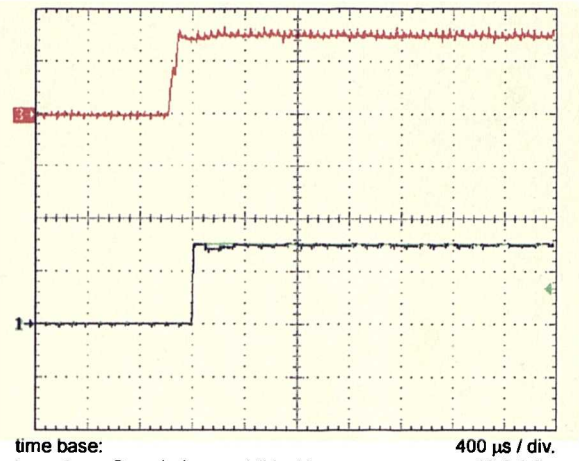

(a) Low Speed Step Response

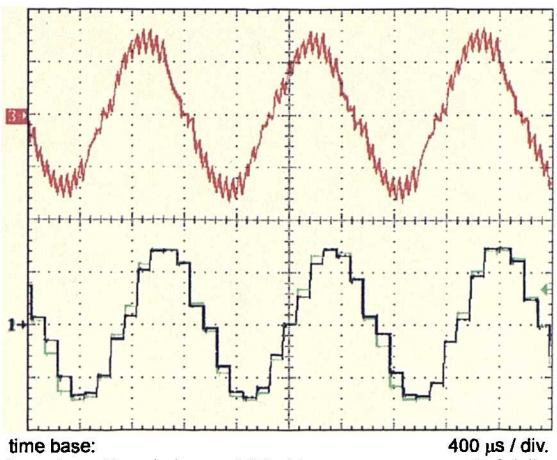

(b) High Speed Test Response

Figure 28 : Performance of Enhanced Model Based Current Control

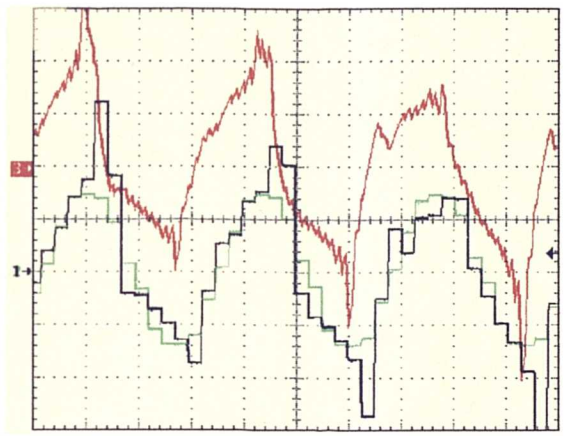

(a) PI Control with Back-EMF Feed Forward

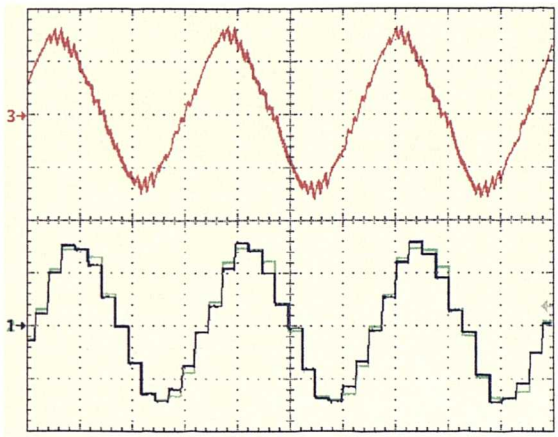

(b) Enhanced Model Based Current Control

Figure 29 : Comparison of Current Controllers as Machine is Driven into Saturation

applying more than the required voltage to the plant. This excess drive combines with the reduced incremental inductance as the phase saturates to give a very large current error.

In contrast the flux controller continues to give a near perfect current waveform. It should be noted that the current waveform is on the limit of what is possible with the DC link voltage used. The absence of current ripple during the fast current transitions indicates that 100% PWM is applied during these periods.

3.6 Comparison of Current Controllers Using RMS Current Error

So far different current controllers have been compared on the basis of graphical evidence. The interpretation of the oscillograms could be a little subjective. Is it possible to make a more quantitative comparison?

Several numerical methods attempt to minimise mean squared error in order to drive solutions towards a best fit [77]. The mean squared current error would be a suitable measure of the quality of current control. However the root mean squared current error defined by (21) has been used, since an error measured in A rms may be visualised more easily than an error measured as A^2. The quality of the various current controllers is compared in Table 6. The RMS current errors were measured for operation with a back-EMF shaped current demand with a peak value of 15 A. These figures correspond directly to the high frequency sine response test results shown in the preceding oscillograms. The figures in Table 6 support the earlier graphical evidence, showing that the simplest flux linkage model based controller gives significantly better control than the best PI based controller.

$$i_{error\_RMS} = \sqrt{\int (i(t) - i^*(t))^2 dt} \tag{21}$$

Control Scheme	RMS Error (A rms)
PI current controller.	9.56
PI current controller with back-EMF feed forward.	3.36
Model based current controller.	2.48
Model based current controller with terms for power device voltage drops, iR drop and dead time.	0.92

Table 6 : Comparison of Current Control Schemes using RMS Error

3.7 Conclusion

PI current control and flux linkage model based current control have been compared on the basis of step and frequency response tests. With the sampling rate used, it was not possible to achieve the gain required in the PI controller to drive the current waveform in phase with the back-EMF. Modifying the PI controller to add back-EMF feed forward gave a large improvement in the current control performance at high speed. However, even with this improvement the PI controller gives very poor current reference tracking as the machine phase approaches saturation.

The idea of a current controller based on the flux linkage characteristic of the phase has been introduced. This model based controller replaces the PI controller with a simple model of the plant. Thus the model based controller requires no tuning and there is no trade off between stability and transient response.

It has been demonstrated that the transient response of the model based controller is exceptional. Some steady state error may be present due to voltage drops in the power devices and phase winding and due to the dead time introduced by the PWM unit. The size of the steady state error may be estimated easily. It has been shown that by modelling the sources of this error in the controller, the steady state error can be reduced towards zero. The resulting controller will always drive the current to the level demanded, subject to there being sufficient volts available to do this. This is true even as the phase is driven into saturation.

Even in the situation where there are insufficient volts available to achieve the demanded level of current, the actual current which will be achieved may be accurately predicted. This has important implications for fault detection. Any substantial difference between the current measured and the current which the controller had expected to achieve at that time can only indicate that there is some fault in the phase. In later chapters it will be shown that the way in which the measured current deviates from the expected current may be used to diagnose certain faults.

Chapter 4

DRIVE PERFORMANCE

4.1 Objectives

In Chapter 2 the design of a fault tolerant demonstrator drive was presented along with measured parameters for the resulting machine. In Chapter 3 a current controller was developed with the very fast dynamic response required to supply individual phase currents to the machine. This chapter will bring together the elements developed in the preceding two chapters and investigate the performance of the drive in its normal, fault-free, state. The tests presented here will look briefly at the electromagnetic and thermal performance of the drive and establish the limits of the viable torque-speed envelope.

Chapter 2 acknowledged that the high per unit inductance of this machine results in a penalty in terms of increased rating in the power converter. It was shown that applying phase advance to the current should reduce the power converter volt-ampere, VA, requirement. The principle aim of this chapter is to investigate the effect of current profile on drive performance. The objective will be to determine the best operating strategy for this unusual machine design.

The study of different current profiles will use the results measured whilst operating with a current profile shaped to match the back-EMF as a starting point. This profile should give maximum torque for a given RMS current. Next the back-EMF shaped current profile will be replaced by a sinusoidal current profile and the resulting effect on power converter VA requirement established. Phase advance will be applied to the current profile and it will be shown that the reduction in converter VA requirement predicted in Chapter 2 really can be achieved. This investigation will provide evidence to back up the analysis of how the drive design might be further improved by considering trade-offs in both the machine and the power converter, presented in Chapter 2. Finally a technique

will be proposed for designing an optimum current profile to further reduce the power converter VA requirement whilst improving the specific power capability of the drive.

4.2 Calculation of Instantaneous Torque

4.2.1 Calculation of Instantaneous Torque from Static Torque Data

At several points it will be useful to derive torque waveforms. The shaft mounted torque meter on the test rig is capable of providing average shaft torque. However this instrument does not have a sufficiently high bandwidth to measure the torque ripple at the operating frequency of 800 Hz. To provide high bandwidth torque information, a technique for calculating electromagnetic torque contributed by each phase is used. The total instantaneous electromagnetic torque for the machine is then the sum of the torques for each phase.

The instantaneous electromagnetic torque for each phase is calculated from the current in the phase and the rotor angle using static torque data. This technique will yield the torque produced by the machine before the stray loss and other rotational losses have been subtracted. The result is that the torque calculated by this method will be higher than the shaft torque actually delivered by the machine. Because the rotational losses increase with speed, the discrepancy will be largest at higher speeds. Whenever a torque derived by this method is presented it will be referred to as 'calculated torque' or 'electromagnetic torque' to differentiate it from measured shaft torque. Whilst it is important to be aware that the calculated torque will always be a little larger than the shaft torque, the usefulness of the technique in deriving the torque waveforms outweighs all shortcomings. This technique is particularly helpful when discussing faulted operation in later chapters, since it allows the torque contribution of each phase to be examined in isolation.

The current and rotor position data used to calculate the electromagnetic torque are output from the controller, so the limitation in the bandwidth of this method is set by the controller sampling rate. Since the controller samples at 10 kHz the calculated torque will have a Nyquist frequency of 5 kHz.

4.2.2 Calculation of Static Torque from ψ-i-θ Data

The static torque data can be obtained from a static torque test such as that presented in Chapter 2, or directly from the ψ-i-θ data presented in Chapter 3. The technique for deriving static torque from the Flux-MMF diagram is described by Staton et al [75]. The instantaneous electro-magnetic torque, T_e, is calculated from the rate of change of co-energy, W', with respect to angle as shown in (22).

$$T_e(\theta_n, i_n) = \left.\frac{\partial W'}{\partial \theta}\right|_{i=\text{constant}} \tag{22}$$

Details of how this technique was applied to the ψ-i-θ data to produce a new lookup table containing T_e-i-θ data are given in Appendix F. Plotting data from the look up table yields the set of curves presented in Figure 30.

Calculation of instantaneous torque from the ψ-i-θ data assumes that there is no significant electromagnetic coupling between phases. This is a reasonable assumption since the measured parameters of the machine show that the design aim of effective isolation between machine phases has been largely achieved, (Table 4 in Chapter 2).

Note that the torque calculated using this technique only accounts for torque production resulting from magnet flux which links the stator winding. The phenomenon of cogging in a PM machine results from interaction between the magnets and stator slotting and

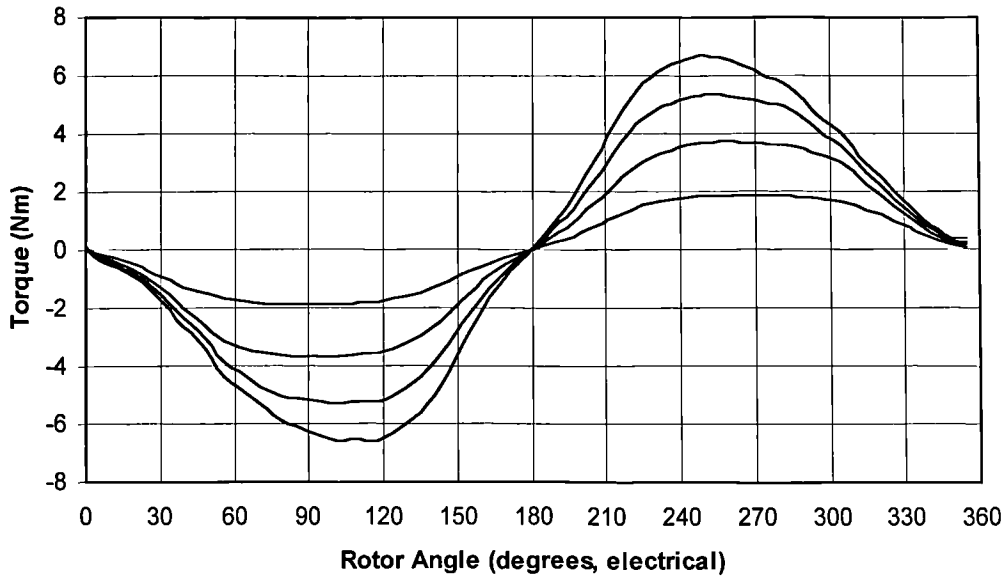

The curves show torque against angle for currents of 12.5 A, 25.0 A, 37.5 A and 50.0 A

Figure 30 : Static Torque Data Derived from ψ-i-θ Data

does not cause any change in the magnet flux linking the stator winding. Thus the electromagnetic torque calculated from ψ-i-θ data does not take account of cogging torque, Deodhar et al [78]. This is not a serious concern here, since the cogging torque over a machine cycle averages to zero and, as shown in Chapter 2, the peak value of the ripple resulting from cogging is only 2.3% of rated torque.

Disregarding the cogging torque, the static torque curves calculated from ψ-i-θ data are in good agreement with those obtained from direct measurement of static torque in Chapter 2. The discrepancy between calculated and measured static torque is always less than 5%. The static torque calculated from ψ-i-θ data, is available in a lookup table and covers a wider range of currents than the directly measured static torque data, so this calculated static torque data is used whenever electromagnetic torque is calculated.

4.3 Operation with Current Profile Matched to Back-EMF Shape

In this section the performance of the drive will be evaluated whilst operating with a phase current profile matched to the back-EMF shape, i.e. with a current profile shaped as shown in Figure 11 of Chapter 2. The current control was arranged to keep the phase current in phase with the machine back-EMF. This should deliver the maximum torque per unit rms current, [72].

The purpose of this performance evaluation is twofold:

- To enable performance of the drive to be compared to the target specification in Chapter 2.
- To provide a reference for comparison of operation with different phase current profiles.

4.3.1 Electrical Performance

The machine was operated at a range of loads and speeds and extensive data gathered. The machine has a thermal time constant of approximately 2 minutes so operation at each load point was maintained for approximately 20 minutes to allow the machine temperature to stabilise. The most important results for each operating point are presented in Table 7 and Table 8.

Test identifier	Speed / rpm	DC link power / kW	Phase current / A RMS	Calculated torque * / Nm	Air gap power * / kW	Shaft torque / Nm	Shaft Power / kW
1	4735	0.59	1.1	0.69	0.34	0.49	0.24
4	4144	1.90	5.3	3.40	1.48	3.25	1.41
5	4165	4.25	12.0	7.60	3.32	7.49	3.27
6	4123	6.46	18.2	11.30	4.88	11.05	4.77
2	8288	1.23	1.5	0.93	0.81	0.58	0.50
7	8350	3.63	5.4	3.45	3.02	3.05	2.66
8	8329	7.94	12.3	7.79	6.80	7.23	6.31
9	8288	11.73	18.7	11.63	10.10	10.45	9.07
3	12359	2.08	2.0	1.21	1.57	0.67	0.87
10	12453	5.27	5.6	3.57	4.66	2.92	3.81
11	12453	11.46	12.9	8.18	10.67	6.95	9.06
12	12391	14.64	17.4	10.79	14.00	8.79	11.41

* *Air gap power is derived from calculated torque and the shaft speed. The calculated torque is derived from rotor position, phase current and flux linkage data.*

Table 7 : Power Flow

Test identifier	Speed / rpm	Shaft torque / Nm	DC link power / kW	Inverter loss ◊ / kW	Copper loss * / kW	Other loss / kW	Shaft Power / kW
1	4735	0.49	0.59	0.05	0.00	0.20	0.24
4	4144	3.25	1.90	0.25	0.03	0.21	1.41
5	4165	7.49	4.25	0.58	0.17	0.23	3.27
6	4123	11.05	6.46	0.87	0.46	0.46	4.77
2	8288	0.58	1.23	0.07	0.00	0.64	0.50
7	8350	3.05	3.63	0.26	0.03	0.68	2.66
8	8329	7.23	7.94	0.59	0.19	0.85	6.31
9	8288	10.45	11.73	0.90	0.53	1.23	9.07
3	12359	0.67	2.08	0.10	0.00	1.11	0.87
10	12453	2.92	5.27	0.27	0.04	1.15	3.81
11	12453	6.95	11.46	0.62	0.23	1.55	9.06
12	12391	8.79	14.64	0.84	0.48	1.91	11.41

◊ *Estimated value of loss.*

* *DC copper loss derived from the RMS phase current and the winding resistance corrected for the operating temperature.*

Table 8 : Losses

Table 7 shows that the operating points cover three speeds at a range of loads. The torque speed envelope covered by tests is illustrated in Figure 32. The last two columns give the shaft torque measured on the torque meter and the shaft power. The shaft power is the product of the torque meter reading and the shaft speed. The third column gives the DC link power, i.e. the power into the inverter. The air gap torque figures have been derived from the product of the electromagnetic torque and the shaft speed. The air gap power is used as an indication of the mechanical power output of the drive before AC losses and other rotational losses have been deducted.

The figures in Table 8 apportion the drive losses. The losses have been derived from the measurements in Table 7. The accuracy of the loss figures is very sensitive to any inaccuracies in the power measurements. The accuracy is further compromised by the absence of any measurement at the interface between the inverter and the machine, thus the calculation of machine losses is dependent on a crude estimate of inverter losses. However, it is clear that the losses in the machine are higher than was expected. The loss figures are further supported by measurements of coolant flow and the temperature difference between coolant inlet and outlet.

Some of the data from Table 7 and Table 8 is shown graphically in Figure 31. The figure shows the inverter losses, copper losses and other losses subtracted from the power supplied by the DC link to give the resulting shaft power. Each component of loss is

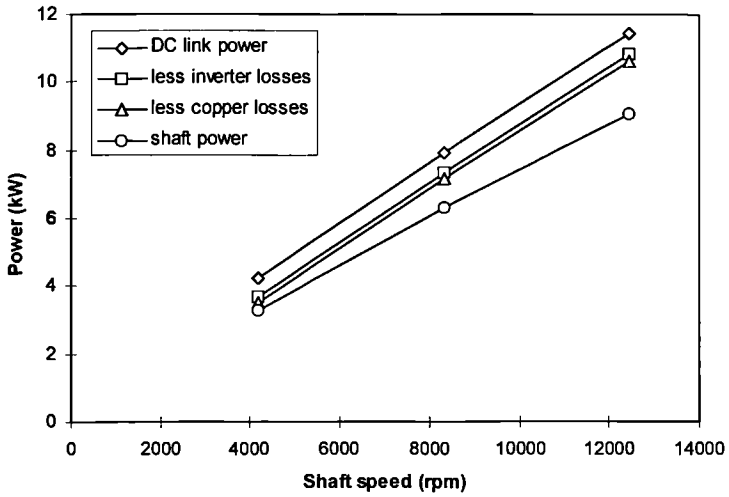

(a) Increasing Speed, Full Load

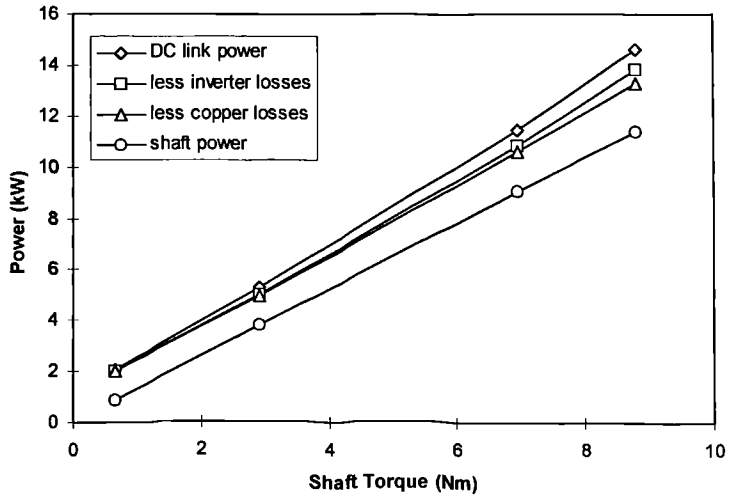

(b) Increasing Load, Full Speed (≈12500 rpm)

Figure 31 : Variation of Losses with Shaft Speed and Torque

represented by the separation of two adjacent lines. Figure 31(a) shows how the losses vary with shaft speed and Figure 31(b) shows how the losses vary with shaft torque. The component of loss referred to, in Table 8, as 'other losses' includes iron losses and other rotational losses. Figure 31 shows that this component is the largest and is most significant at high speeds.

4.3.2 Thermal Performance

The steady state thermal situation resulting from motoring operation is given in Table 9. Note that the operating points used in Table 9 are the same as those used in Table 7 and Table 8.

Test identifier	Speed / rpm	Shaft torque / Nm	Shaft Power / kW	Machine loss * / kW	Temperature rise◊ / °C
1	4735	0.49	0.24	0.20	22.3
4	4144	3.25	1.41	0.24	26.8
5	4165	7.49	3.27	0.40	61.5
6	4123	11.05	4.77	0.92	124.0
2	8288	0.58	0.50	0.64	44.0
7	8350	3.05	2.66	0.71	50.8
8	8329	7.23	6.31	0.85	95.0
9	8288	10.45	9.07	1.76	160.5
3	12359	0.67	0.87	1.11	64.5
10	12453	2.92	3.81	1.19	64.8
11	12453	6.95	9.06	1.78	116.3
12	12391	8.79	11.41	2.39	172.0

* *Machine losses are calculated from the total loss minus the estimated inverter loss.*
◊ *Average copper temperature minus average coolant temperature.*

Table 9 : Thermal Results, Motoring under PWM Current Control

Whilst the electromagnetic performance of the machine is in line with the requirement in Chapter 2, the thermal performance is poor. The resulting temperatures in the machine prevent sustained operation over the full torque speed envelope. Figure 32 shows the limit of steady state testing on the machine contrasted to the required torque speed envelope for the drive. The results in Table 9 have been used to predict the torque speed envelope projected to give the 80 °C temperature rise specified in the requirement (Table 1 of Chapter 1). This torque speed envelope is also shown in the figure.

ORIGINAL IN COLOUR

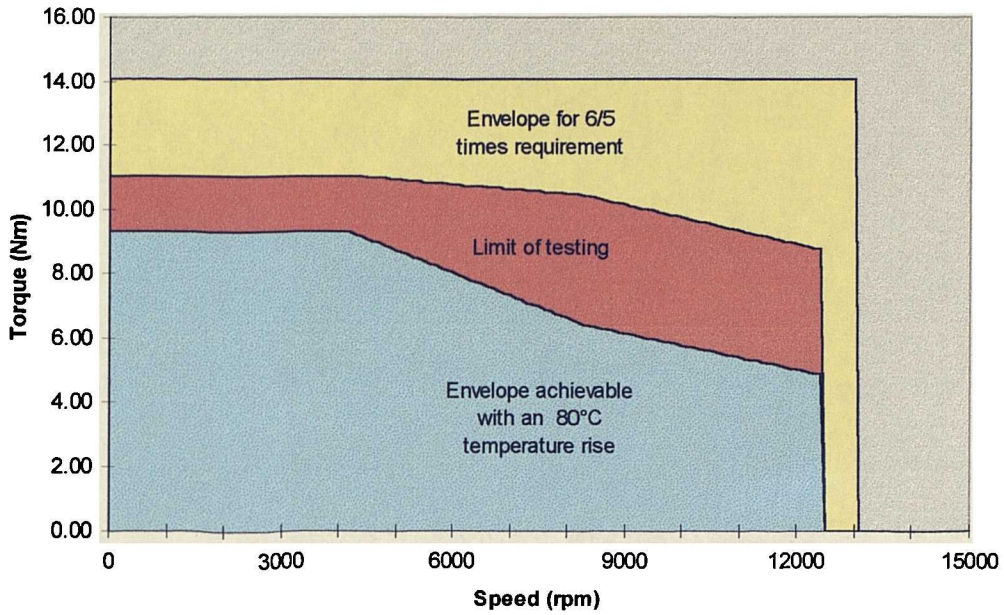

Figure 32 : Desired and Measured Torque Speed Envelopes

The poor thermal performance of the machine is due to two factors:

- Higher than predicted thermal resistance between the copper and the coolant. (1.21 °C/W measured in contrast to 0.71°C/W required).
- High iron and stray losses (including sleeve loss) associated with high speed operation.

Both these factors are associated with efforts to obtain a very high specific power. The thermal resistance could be substantially reduced by operating the machine with a flooded rotor. Thermocouples attached to the stator indicate that much of the iron loss is occurring in the tooth tips. This may be reduced by using thinner stator laminations. Since these problems are not thought to be compounded significantly by the fault tolerant features of the design they are not discussed further here.

It is possible that the iron loss could be reduced by employing an operating regime which results in reduced levels of flux and rate of change of flux in the machine. Later it will be seen that the process of optimising the current profile tends to reduce both the peak flux and the rate of change of flux.

4.3.3 Waveforms from Operation with a Back-EMF Shaped Current Profile

For each operating point, a set of data output from the controller was recorded. This data included the operating speed, the rotor position, the PWM applied to each phase and the current measured in each phase at every sampling interval for a number of revolutions of the machine. This data was processed to generate Figure 33(a) and Figure 34(a) which show phase current in one phase of the machine and the corresponding phase voltage demanded by the controller during tests 6 and 12, i.e. operation at maximum load, one third speed and operation at maximum load, full speed respectively. Note that the voltage waveforms include the volts required to overcome the drops across the power devices and winding resistance, thus the peak voltage applied is a true indication of the DC link voltage requirement. The data has been further processed using the ψ-i-θ data to calculate the ψ-i loci shown in Figure 33(b) and Figure 34(b). The data

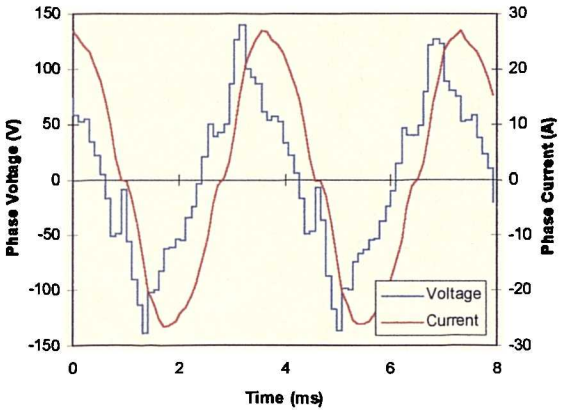

(a) Sampled Phase Current and Voltage Demand

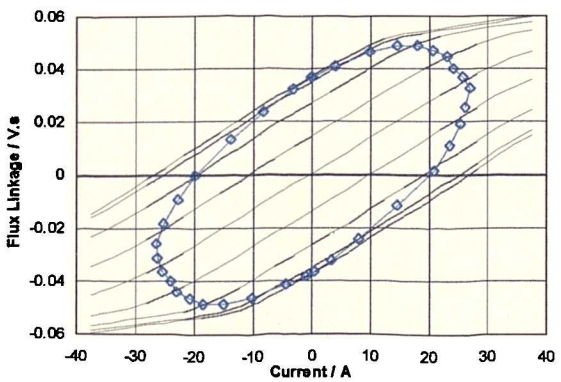

(b) Calculated ψ-i Locus (from measured data)

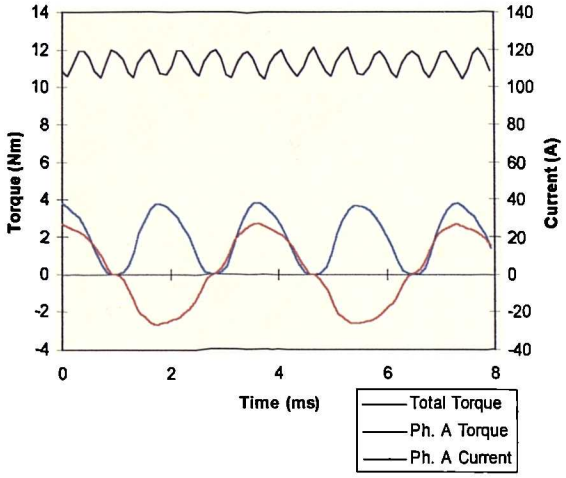

(c) Calculated Torque Output

Figure 33 : Operation at 4000 rpm with Back-EMF Shaped Current Profile

was also processed using a lookup table containing static torque data to obtain Figure 33(c) and Figure 34(c). These show the calculated electromagnetic torque on a cycle by cycle basis.

The figures show that the current control is capable of handling the high frequency current waveforms so that the ψ-i locus is still well formed even at the maximum speed. At this speed the electrical frequency is 800 Hz and the ripple in the total torque is at 6 times this frequency, 4.8 kHz. The frequency of the torque ripple is approaching half the data sampling frequency, 10 kHz, hence some aliasing can be seen in Figure 34(b). Note that the torque ripple shown results from interaction between stator current and rotor flux. This torque is in addition to the small amount of cogging torque which, whilst present, is not included in the electromagnetic torque plots.

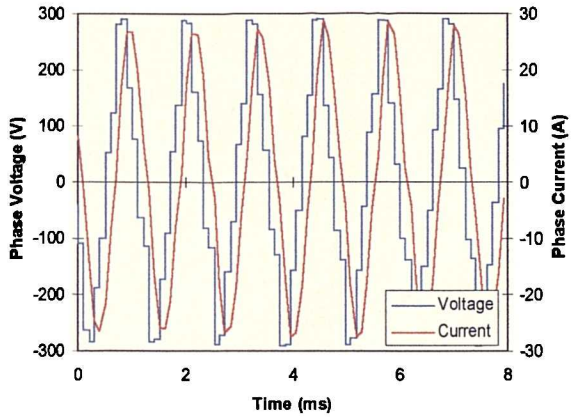

(a) Sampled Phase Current and Voltage Demand

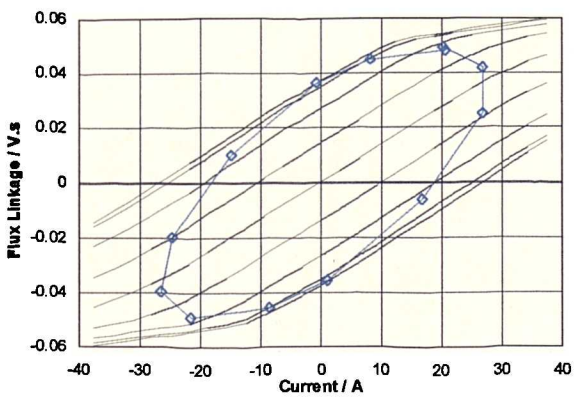

(b) Calculated ψ-i Locus (from measured data)

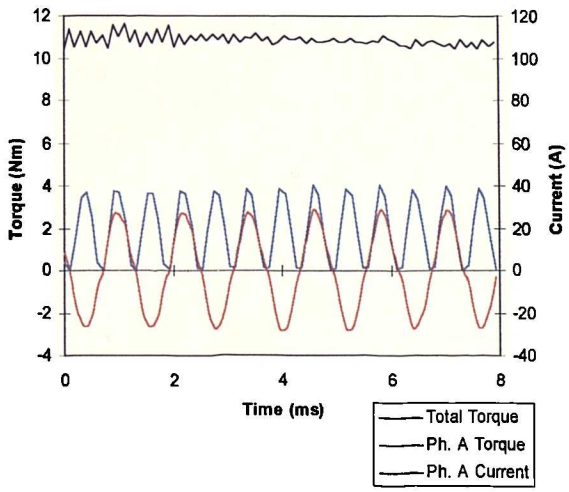

(c) Calculated Torque Output

Figure 34 : Operation at 12000 rpm with Back-EMF Shaped Current Profile

Since operation with a back-EMF shaped current profile, controlled to be in phase with the back-EMF, should yield the maximum possible torque per unit current, the ψ-i locus should encircle the maximum possible area for the rms current used. The penalty is that the peak voltage required to supply the current is high. This is clear from Figure 33(a), where it can be seen that large peak voltages are required to drive the rapid increase in current. The presence of a significant third harmonic in the current demand profile results in a steeper rise in current, compounding the problem. The same situation is also apparent from the ψ-i diagram, which at some current steps requires large changes in flux linkage and hence large applied voltages.

4.4 Operation with Sinusoidal Phase Current and Varying Current Angle

4.4.1 Choice of Criteria Used to Select Optimal Current Profile

The discussion so far has used a current demand profile with a shape identical to the machine back-EMF. This profile was chosen as it gives the best torque per amp. However this is not the only performance measure against which the drive may be judged. It may be desirable to have a drive design and operating regime to:

- Minimise inverter VA rating
- Minimise iron requirement in the machine
- Minimise copper requirement in the machine
- Minimise copper loss (i.e. maximise torque per amp)
- Minimise iron loss
- Minimise total loss
- Minimise winding temperature rise

These requirements may be achieved by altering the machine design or by adjusting the way in which it is driven. The optimisation of some quantities is at odds with others, so in practice all designs are a compromise between these factors.

This machine has been designed with a back-EMF which is only just low enough to permit the current to be driven into the machine by the DC link voltage of 270V. So it

appears that it is not possible to reduce the DC link below it's present value. In Chapter 2 it was suggested that the VA rating of the power converter could be reduced by phase advancing the current with respect to the back-EMF. The angle between the current and the back-EMF is referred to as the current angle and given the symbol γ. This section will show that this works in practice. Phase advancing the current in this type of machine is effectively field weakening, or as Jahns [79] points out, flux weakening. The result is that for a given DC link voltage it is possible to drive more current into the machine. Although the torque per amp reduces, the possible torque actually increases for a given DC link voltage as observed by Jahns in his discussion of an interior PM machine.

The effect of phase advancing the current whilst keeping the torque output constant was examined in Chapter 2. This necessitated increasing the current and required that the volume of copper in the machine be increased, to keep the electric loading unchanged. At the same time the DC link voltage could be reduced. As the flux in the machine was reduced the volume of iron in the machine could also be reduced in line with keeping the magnetic loading constant.

Once the machine has been constructed, the amounts of copper and iron in the machine are fixed, as is the phase reactance, so the effects of phase advance will have to be examined on a different basis to that used in Chapter 2. From this point there are two obvious approaches to examining the effect of current phase advance at a given speed.

1. ***Operation at fixed torque, Figure 35.*** As the current is phase advanced with respect to the back-EMF, the current amplitude is increased to maintain the torque. This makes it easy to directly compare the current and VA required at the particular operating point chosen in the torque speed envelope. The drawbacks are that the electric loading increases, leading to higher winding temperatures, and that IaXs increases above 1 p.u. requiring more reactive VA.

2. ***Operation at fixed current magnitude, Figure 36.*** This will ensure comparison with the short circuit current, I_{SC}, at 1.0 p.u.. The volume of copper in the machine will remain correct and IaXs will remain at 1.0 p.u. The drawback is that the machine is no longer being compared on the basis of the same point on the

torque speed curve. Also the magnetic loading will fall as the current is phase advanced. The result of the surplus iron in the machine is likely to be a reduction in iron loss but also less specific power.

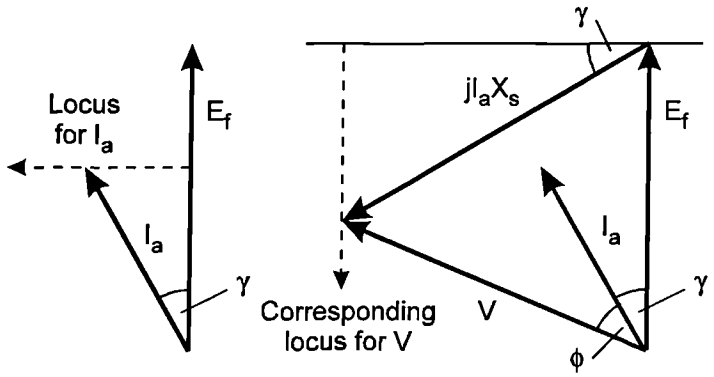

Figure 35 : Flux Weakening with Constant Torque

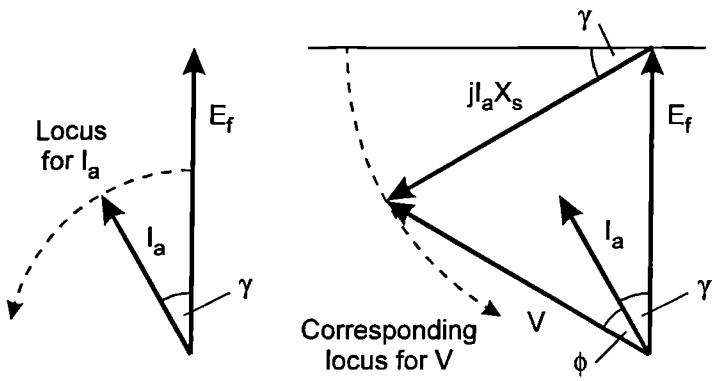

Figure 36 : Flux Weakening with Constant Current Magnitude

The second of these two approaches was adopted, as a comparison where IaXs remains at 1.0 p.u. appears most valuable. Operation with constant current demand but progressively increasing current phase advance was investigated. Note that the vector diagrams assume that the currents and voltages are harmonic free sinusoidal quantities. To avoid unnecessary complications the back-EMF shaped current profile used earlier is replaced by a sinusoidal current profile for the purpose of investigating the effects of phase advance.

4.4.2 Results for Sinusoidal Current with Phase Advance

For all tests the DC link volts are kept at approximately 300 V. The DC voltage requirement for the drive is calculated as the product of the maximum demanded PWM duty ratio and the measured DC link voltage.

The waveforms resulting from operation with a current angle of 0.0° are shown in Figure 37. In spite of the sinusoidal current demand, the current in the machine is not perfectly sinusoidal. From the ψ-i locus it is clear that the distortion of the current waveform is occurring as the phase saturates magnetically. At this point the errors in the model have the greatest effect as the phase inductance is reduced. This compounds the errors which have resulted as the phase resistance increases through heating and as the DC link voltage reduction at high DC link currents is not accounted for in the dead time calculation.

In Chapter 3 the poorer performance of the current control in the saturated region was quantified. The saturation of the phase alone does not seem to account for all of the

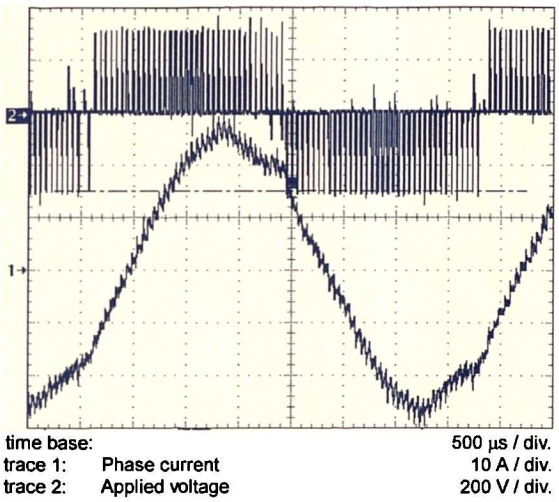

(a) Measured Current and Voltage Waveforms

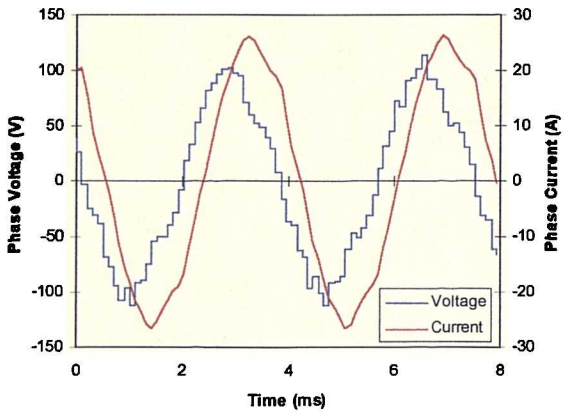

(b) Sampled Phase Current and Voltage Demand

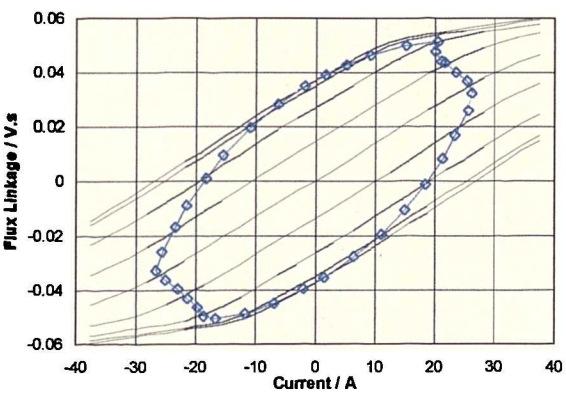

(c) Calculated ψ-i Locus (from measured data)

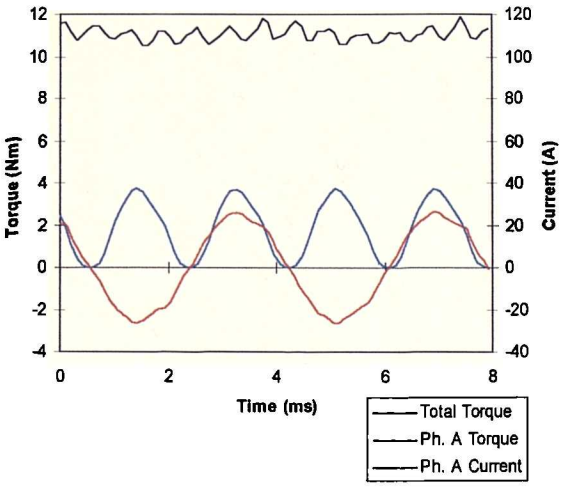

(d) Calculated Torque Output

Figure 37 : Operation at 4000 rpm with Sinusoidal Current Profile and $\gamma = 0°$

current error and it appears that a small amount of cross saturation between phases is also present. By comparing the voltage demand waveform to that in Figure 33, the principle benefit of moving from a back-EMF shaped current waveform to a sinusoidal current profile can be seen. Operation with the sinusoidal current profile here results in a less peaky voltage demand than operation with a back-EMF shaped current profile. This results in a lower DC link voltage requirement and hence a lower power converter VA rating.

The waveforms resulting from operation with a current (advance) angle of 30 degrees are shown in Figure 38. It can be seen in the ψ-i locus that applying phase advance to the current has reduced the peak flux linkage in the machine. Thus operation with current

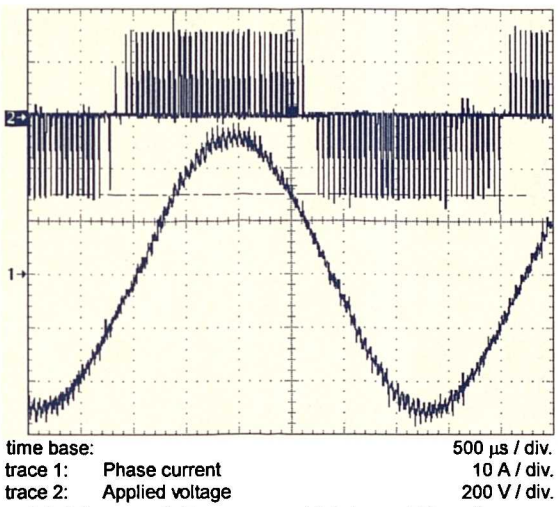

(a) Measured Current and Voltage Waveforms

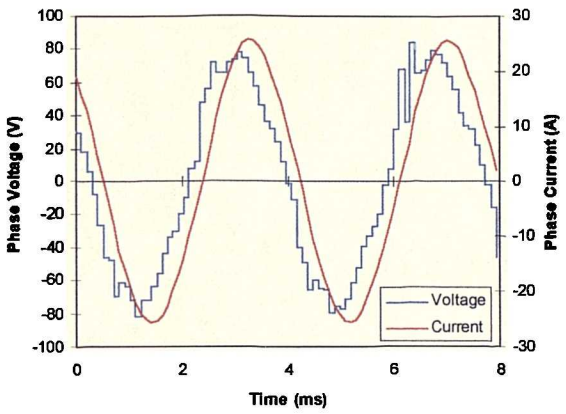

(b) Sampled Phase Current and Voltage Demand

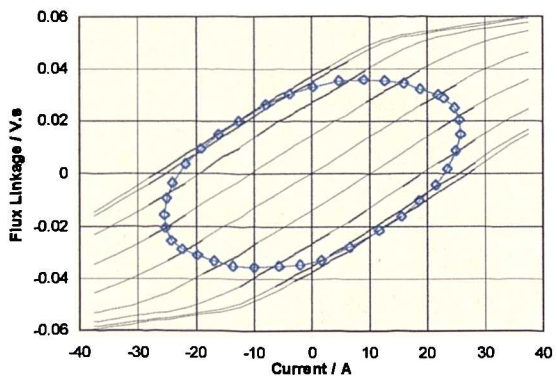

(c) Calculated ψ-i Locus (from measured data)

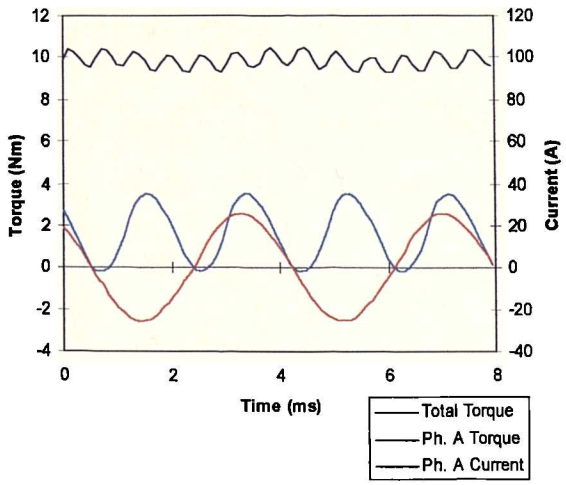

(d) Calculated Torque Output

Figure 38 : Operation at 4000 rpm with Sinusoidal Current Profile and $\gamma = 30°$

phase advance is effective in flux weakening. The reduction in flux linkage also ensures that the machine is no longer driven into saturation. Thus the phase current closely follows the demand, resulting in highly sinusoidal phase currents. Another effect of phase advancing the current is that the peak voltage demand is clearly reduced for a given current demand - as predicted. Finally the average torque contributed by each phase is reduced and at some rotor positions the torque from individual phases is negative.

Operation with constant current demand but progressively increasing current phase advance provided the results summarised in Table 10.

Test identifier	Current Phase Advance /degrees	Speed / rpm	DC link power / kW	Phase current / A rms	Calculated torque * / Nm	Air gap power * / kW	Shaft torque / Nm	Shaft Power / kW
21	0.0	4010	6.26	18.1	11.06	4.64	11.05	4.64
22	10.1	4010	6.26	18.1	11.01	4.62	10.85	4.56
23	20.0	4010	6.06	18.1	10.60	4.45	10.52	4.42
24	25.0	4010	5.91	18.1	10.26	4.31	10.29	4.32
25	30.1	4010	5.75	18.1	9.83	4.13	9.80	4.12
26	35.2	4010	5.51	18.1	9.31	3.91	9.45	3.97
27	40.2	4010	5.24	18.1	8.75	3.67	8.88	3.73
28	50.5	4010	4.63	18.1	7.36	3.09	7.05	2.96
29	60.7	4010	3.86	18.1	5.66	2.38	5.85	2.46
30	70.5	4010	3.07	18.1	3.89	1.63	4.03	1.69
31	80.8	4010	2.16	18.1	1.95	0.82	2.30	0.97

* *Air gap power is calculated from calculated torque. The calculated torque is calculated from rotor position, phase current and flux linkage data.*

Table 10 : Power Flow with Phase Advance

From the calculated torque, shaft speed, rms phase current and peak voltage demand, the peak power converter VA requirement per unit shaft power has been calculated. This is plotted alongside a line showing shaft torque in Figure 39. The graph shows that the power converter VA rating required to supply a given amount of shaft power reduces as phase advance is applied. As more and more phase advance is applied the VA requirement eventually starts to increase again as the current supplied provides less and less torque, making the losses more significant.

It appears that the best operating point is between 35 and 60 degrees of phase advance. Because the curve is constructed from operating points with constant rms current, it does not account for the increased VA requirement of driving the larger currents which are required to maintain torque when phase advance is applied. Thus the kVA/kW

ORIGINAL IN COLOUR

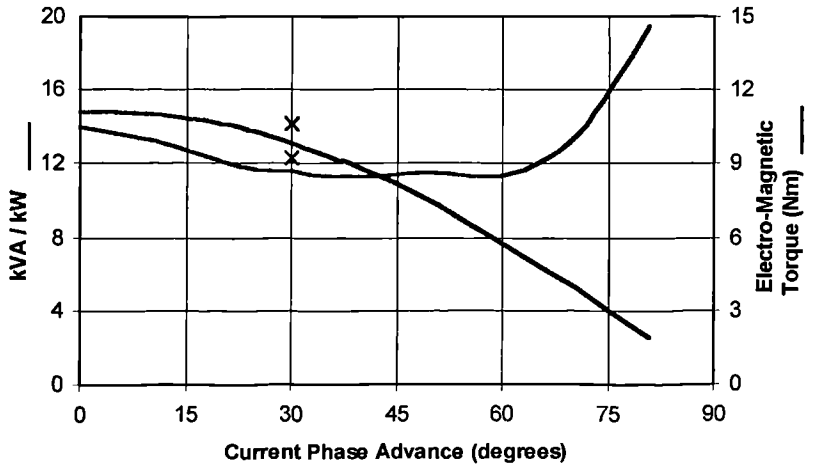

Note. The additional point shown by the crosses is with the current increased to regain the torque delivered with no current phase advance.

Figure 39 : Effect of Phase Advancing Current on VA Requirement

requirement should start to rise again sooner. To illustrate this an additional experimental point has been added at 30 degrees phase advance. The current magnitude at this point has been increased to deliver the same shaft torque as was delivered with zero phase advance. Since the current has been increased IaXs has increased above 1.0 p.u. making the VA requirement unnecessarily high. So the true minimum kVA/kW value achievable will be between this point and the curve.

As discussed in Chapter 2, current angles greater than 30 degrees result in serious penalties for the machine. Thus a current angle of 30 degrees is chosen as the preferred level of phase advance. Figure 39 shows that the use of flux weakening in this machine type results in a reduction in the VA rating of the power converter and in the link voltage required of between 14% and 21% compared to a theoretical figure of 18% from Chapter 2. Note that whilst the kVA/kW reduction achieved in practice is similar to that predicted, the actual kVA/kW required is approximately 20% higher than predicted in Chapter 2. The discrepancy arises because the kVA required in practice must also supply the losses in the drive.

4.5 Operation With an Optimised Current Profile

The discussion of current phase advance has so far assumed sinusoidal current profiles. Is this the best approach?

To make best use of the available DC link volts, motor drives are sometimes operated under voltage control. In voltage control operation the current waveforms are far from sinusoidal. Can the benefits of voltage control be achieved whilst operating under current control, by selecting a current profile which maximises torque output for a given DC link voltage? If this is possible then the drive could remain in current control for all operating conditions.

4.5.1 Designing a Current Profile to Minimise Peak Voltage Requirement

In principle it must be possible to derive a current waveform to minimise the peak voltage requirement - and hence DC link voltage requirement. It will be shown that it is possible to obtain at least a good approximation to this optimum current waveform by graphical means using the ψ-i diagram. A peak current of 26.0 A is chosen as a starting point as this is the peak value of the short circuit current.

To achieve the maximum possible torque with this peak current, the current in the machine should be equal to the positive peak current whilst the back-EMF is positive, and equal to the negative peak current whilst the back-EMF is negative. This would result in a square wave of current. Unfortunately the sharp transitions in current present in the square wave would result in large peak voltage demands. For square wave current profile the current transitions would occur at the aligned positions, with the negative going transition occurring at the aligned position corresponding to positive phase current.

The next step is to modify the square wave by shaping the transitions between the two peak current values. When designing a transition profile to minimise the voltage requirement there are two variables - the timing and the rate of the transition. The objective in shaping the transitions is to give the fastest possible transition subject to a chosen peak voltage requirement.

The transition profile was designed as follows:

- The desired current profile will be defined at 32 points, equally spread across the electrical cycle. Thus each point is separated by 11.25 electrical degrees. The position of the first point is fixed at 0 degrees, i.e. at the aligned position.

- The negative going transition was designed to finish at 45 degrees past the aligned position corresponding to positive phase current. The figure of 45 degrees appears arbitrary but is in fact selected by iteration of this transition profile design process, to maximise the area enclosed by the ψ-i locus.
- The rate of transition was mapped out back to the peak positive current in steps of 11.25 degrees. The change in current over each angle step was chosen to be as large as possible without exceeding a voltage requirement equal to the peak phase back-EMF.
- The positive transition was then arranged to mirror the negative going transition.

The resulting optimised ψ-i locus is shown in Figure 40(a). The optimised current profile and calculated phase torque are plotted against time for constant speed operation in Figure 40(b).

The current profile has been devised to result in a peak voltage requirement equal to the peak back-EMF. Thus, neglecting conduction losses, it will be possible to supply the current from a DC link voltage equal to the peak back-EMF. This same low DC link voltage requirement was achieved with a sinusoidal current profile and 30 degrees of phase advance. However by comparison of the ψ-i diagram here with that in Figure 38 it is clear that there is far less flux weakening occurring here and correspondingly more torque available.

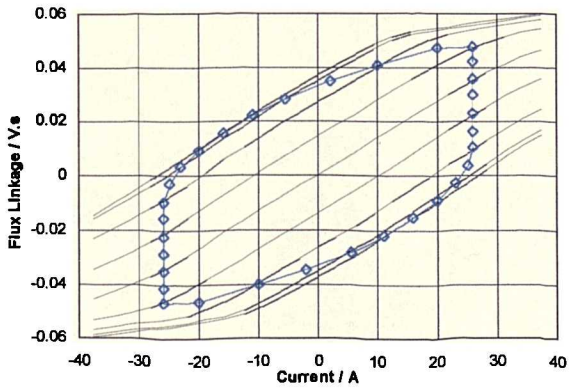

(a) Optimum ψ-i Locus

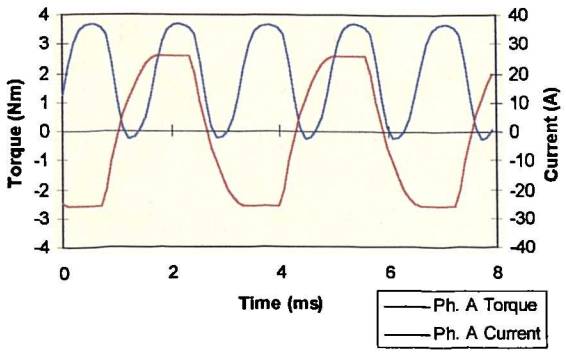

(b) Calculated Phase Current and Torque

Figure 40 : Operation at 4000 rpm with Peak Voltage Minimising Current Profile

4.5.2 Drive Performance Using a Peak Voltage Minimising Current Profile

Figure 41 shows the waveforms obtained using the optimised current profile. The current waveform achieved in practice, Figure 41(a), is close to the desired current waveform in Figure 40(b). Similarly the ψ-i locus calculated from the measured data, Figure 41(c), compares well with the intended ψ-i locus in Figure 40(a). Note from the plot of phase voltage requirement, that the voltage demand is close to a square wave, demonstrating that the available volts are being well utilised, as would be the case under voltage control. However the current is well controlled in contrast to the ragged, peaky current which typically results from voltage control operation, thus the copper losses are likely to be lower for a given torque than under voltage control. These results are best

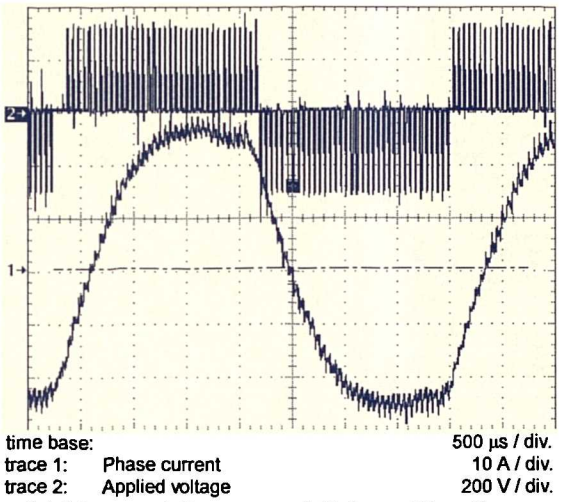

(a) Measured Current and Voltage Waveforms

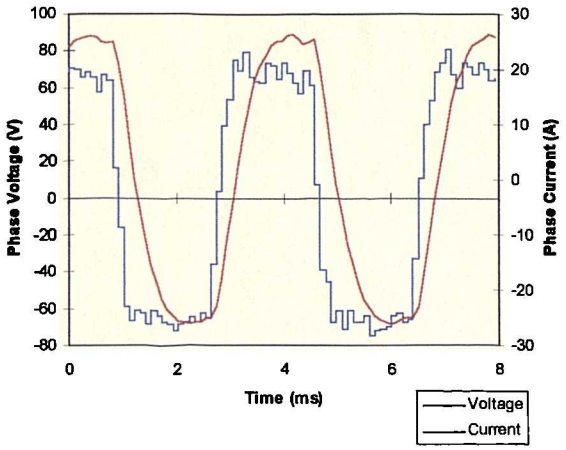

(b) Sampled Phase Current and Voltage Demand

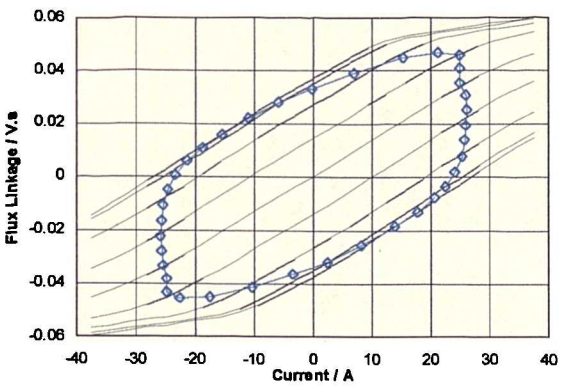

(c) Calculated ψ-i Locus (from measured data)

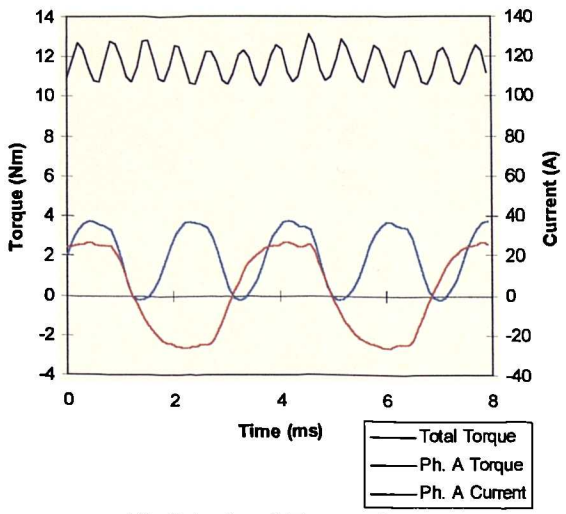

(d) Calculated Torque Output

Figure 41 : Peak Voltage Minimising Current Profile, 4000 rpm, 20.5 A rms

compared with those obtained for a sinusoidal current profile with 30 degrees of phase advance (Figure 38). Here 18% more torque is produced from the same DC link voltage. The drawback is that the rms current has increased by 13% to 20.5 A, resulting in excessive temperature rise if sustained operation is attempted.

Figure 42 shows the effects of reducing the amplitude of the current back to 18.1 A rms. Again comparing these results to those obtained for sinusoidal current profile and 30 degrees of phase advance in Figure 38. Here, 4% more torque is produced from the same rms current and the DC link voltage requirement has actually fallen.

In this section we have shown that a departure from sinusoidal current profiles gives still more power for a given VA requirement and also more power from a given DC link voltage.

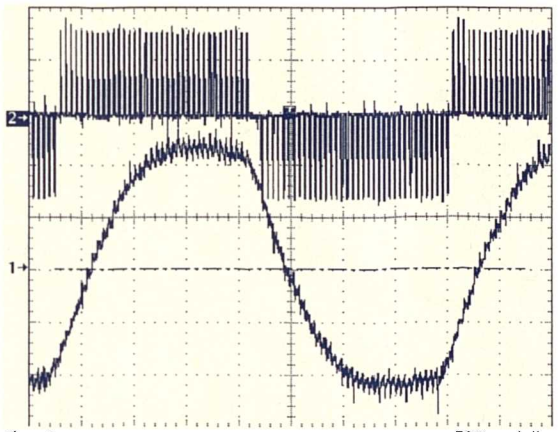

(a) Measured Current and Voltage Waveforms

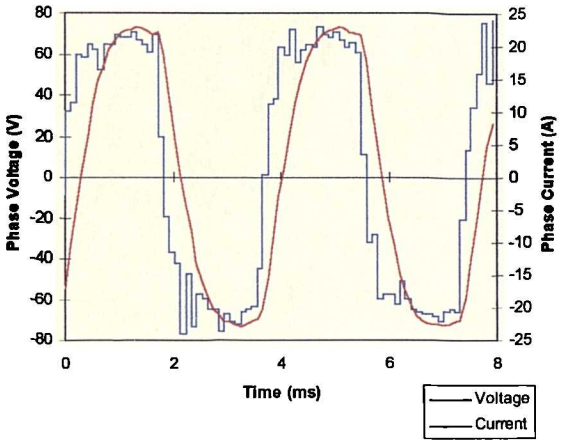

(b) Sampled Phase Current and Voltage Demand

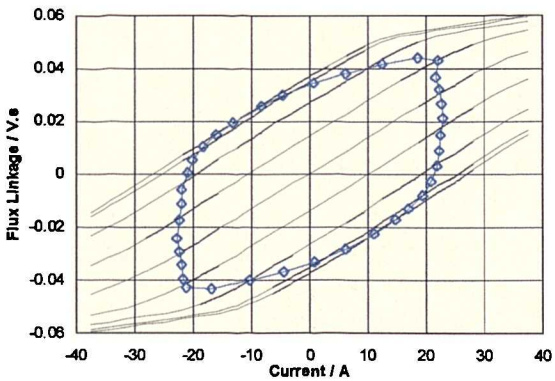

(c) Calculated ψ-i Locus (from measured data)

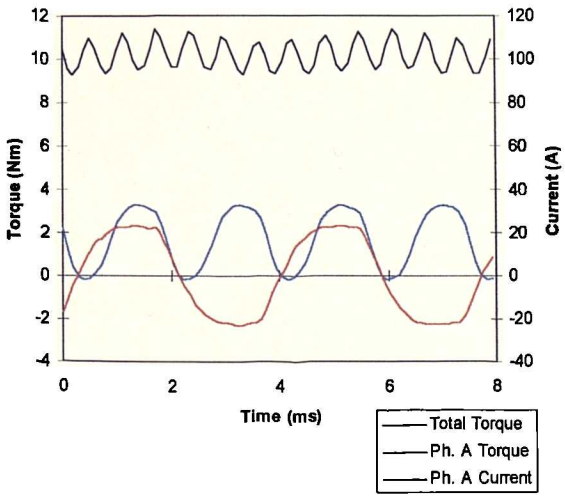

(d) Calculated Torque Output

Figure 42 : Peak Voltage Minimising Current Profile, 4000 rpm, 18.1 A rms

4.6 Comparison of Current Profiles

The key results for operation with each current profile are summarised in Table 11. The results are normalised and shown graphically in Figure 43. Operation with sinusoidal current in phase with the back-EMF should be taken as a starting point for comparisons as the machine was initially designed for operation in this mode. This was the starting point used in Chapter 2 when discussing material penalties for fault tolerance.

Operation with a back-EMF shaped current profile delivers 2% more torque but requires a 14% higher VA requirement in the power converter. This makes operation with this current profile unattractive for a machine with high p.u. reactance.

Phase advancing a sinusoidal current waveform by 30° results in an 11% loss of torque. However this lower torque is achieved with 21% less VA requirement in the power converter, making phase advance an attractive proposition, as predicted in Chapter 2.

Test ID	Comment	Rms current /A	Calculated Torque /Nm	DC Link Voltage Requirement / V	Peak kVA/kW
test 6	back-EMF shaped current profile	18.2	11.30	130	17.3
test 21	0.0 degrees phase advance	18.1	11.06	106	14.0
test 25	30.1 degrees phase advance	18.1	9.83	78	11.6
test 36	current profile for minimum peak voltage requirement	18.1	10.22	72	10.4

Table 11 : Comparison of Operation with Different Current Profiles

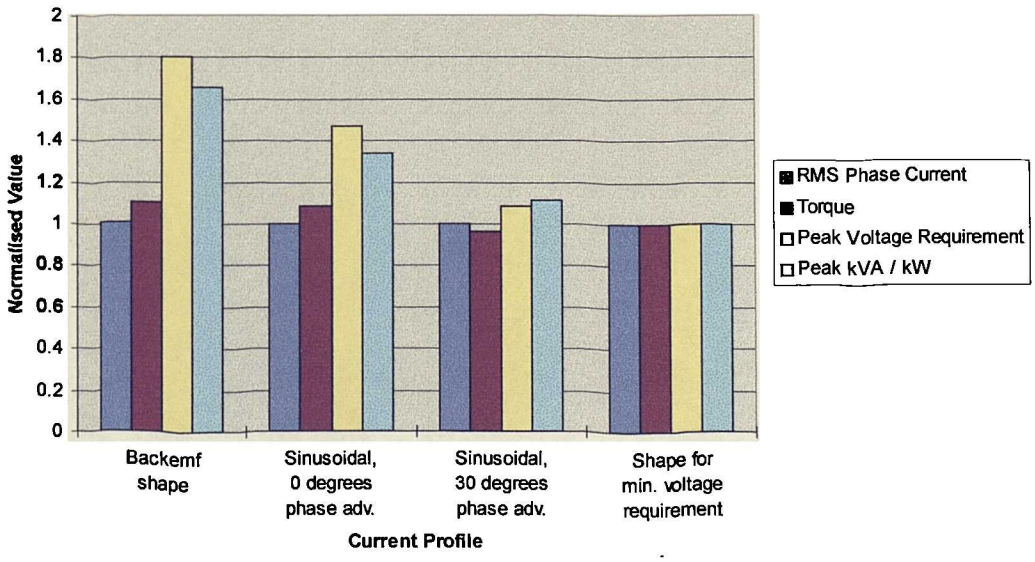

Figure 43 : Comparison of Operation with Different Current Profiles

By using a current profile which has been optimised for minimum peak voltage requirement, some of the torque lost as a result of phase advancing the sinusoidal current has been regained and, at the same time, the kVA requirement has been further reduced. By profiling the current to make best possible use of the available link volts, the torque capability is maximised and the link voltage requirement is minimised, thus optimising on the converter VA rating. The resulting figure of 10.4 kVA/kW is similar to the figure of 10.5 kVA/kW which Miller [65] suggests is achievable in a practical switched reluctance drive.

4.7 Conclusion

The poorer than expected thermal resistance from the winding to the case and the higher than expected iron losses mean that operation with full torque at full speed is not sustainable. Operation at full speed and light load and operation with full torque at 1/3 speed are both sustainable but result in temperature rises above those given in the requirement. Further work is required to reduce the losses in the machine and to improve the thermal performance. These issues are associated with the demands for high specific powers and high speeds, not with design features required provide fault tolerance.

The power converter volt-ampere requirement for a given drive power can be reduced by applying phase advance to the current with respect to the back-EMF. This is in contrast to a PM machine with low phase inductance and no saliency, where the best operating point will be with the current in phase with the back-EMF.

The power output from a given power converter VA rating (given DC link voltage) is traditionally maximised by switching from current control to voltage control. However, the power converter VA requirement can be similarly reduced by choosing a current profile to make the best possible use of the available volts and continuing operation under current control. A method for designing an optimum current profile has been proposed and demonstrated.

In the pump drive discussed here, which naturally requires maximum torque at maximum speed, flux weakening is still useful in reducing the VA requirement of the power converter. Flux weakening in PM machines is often associated with buried magnet rotor designs, which have significant saliency, so it may seem surprising that a surface mounted

magnet design is suitable for flux weakening. The key attribute of this drive which makes flux weakening possible and indeed advantageous is the high per unit inductance.

Chapter 5

FAULT ANALYSIS

Previous chapters have laid the foundations for a fault tolerant drive by partitioning the required functionality into six independent single phase units. The design provides effective isolation between phases to ensure that a fault in one phase will not have a direct impact on the adjacent phases. However, each phase must share a common input, the DC link, and contribute to a common output, shaft torque. The possibility of a single phase fault indirectly degrading the contribution of remaining healthy phases, through its impact on the DC link or shaft torque, must be considered.

Until now, there has been little discussion of the effect of each fault on the DC supply or on the shaft torque. Each fault must be examined in turn to assess its effect on the DC link and on the torque output of the phase. To ensure that no fault indirectly affects adjacent phases it is essential that no fault results in a serious disturbance to the DC link voltage (by resulting in excessive energy flow to or from the link) or contributes net shaft torque which opposes the torque demand.

5.1 Objectives

This chapter will examine the types of faults which may occur in the drive. Some effort will be made to assess the relative likelihood of each fault and a short list of important faults will be presented. The effects of the principle faults will then be investigated in detail with 4 principle aims:

- Verify that the fault does not directly affect the operation of healthy phases.
- Determine the effect of the fault on the DC link supply.
- Determine the effect of the fault on shaft torque.
- Evaluate the time available to detect the fault.

This chapter is only attempting to assess the consequences arising from undetected faults. At this stage no attempt is made to establish what action should be taken in the event of each fault.

5.2 Faults under Consideration

There are many potential faults which can occur in a drive system. Within this work the range of faults under consideration must be restricted. The consideration of fault tolerance in the section on drive design was restricted to the power circuit and the electric machine, so in this section only the faults which may occur in these units are considered. Thus failures in the control electronics, for example, have been disregarded. In general, sensors within the system have also been disregarded as the nature of these sensors could easily change.

One technique used in the analysis of likely faults is a fault tree [80]. Debebe et al [47] use a fault tree to express the likely failures in an induction machine based drive. More recently Calonnec et al [81] have used the fault tree to discuss reliability in an induction machine. Of the 2 pieces of work, that by Debebe et al is more useful to us since it addresses the complete drive. The faults which may occur in the demonstrator drive are presented in a fault tree in Figure 44.

The fault tree can be used to present the factors which cause component failure, sometimes referred to as stressors. This is useful in that it allows the effect of reducing the individual sources of component stress to be examined. Here it is accepted that even in a well designed system there will always be stress inducing factors and that these will result in occasional component failure. The principle stress inducing factors will be heat, vibration and thermal cycling. In the power electronic converter, voltage and current stress will also be factors but it is assumed that these are kept within the limits specified by the manufacturers of each component.

The following assumptions are made about the likelihood of certain faults:

- The physical separation of phases ensures that phase to phase faults can only result from 2 separate phases sustaining individual phase to frame faults.
- Whilst phase to frame faults may occur as a result of the direct breakdown of ground wall insulation, they are most likely to result from heat damage following a turn to turn fault.

These assumptions lead to a short list of important faults:

- Winding open-circuit.
- Winding short-circuit at terminals of a phase.
- Winding short circuit between adjacent turns within a phase.
- Power device open-circuit - caused by an IGBT, power diode or gate drive fault.
- Power device short-circuit - caused by an IGBT, power diode or gate drive fault.
- DC Link capacitor failure - either short circuit or open circuit.

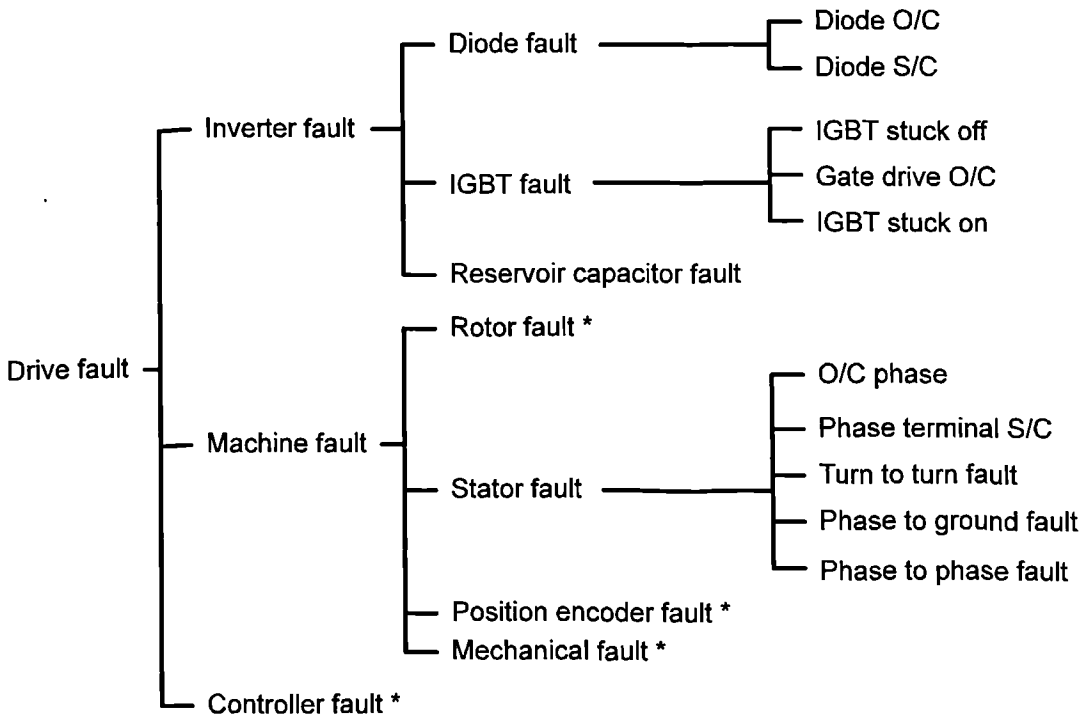

Figure 44 : Fault Tree of Possible Faults in the Demonstrator Drive

The aim of the research is to develop a drive which is capable of operating in the presence of any one of these faults. The following sections will examine the extent to which this aim has been achieved.

5.3 Existing Work in the Area of Drive Fault Analysis

The topic of operation under fault conditions has been the subject of much recent research. Some of the most significant work looking at the consequences of faults in various drive types should be mentioned.

Jahns [14] introduces an induction machine drive with six independent phases, but admits that strong electro-magnetic coupling between phases limits the degree of independence which can be achieved between drive outputs. The paper presents simulated and measured waveforms for phase open circuit and phase terminal short circuit faults. The faults are examined under both current and voltage control. The open circuit phase fault results in 20% reduction in torque and a 15% increase in rms torque ripple. The phase terminal short circuit fault results in several times rated current which is not sustainable.

Kastha and Bose [4] investigate the fault modes of a voltage fed PWM inverter supplying an induction motor. This study includes input line to ground and input rectifier short circuits as well as base drive open circuits (treated as a transistor open circuit fault) and power transistor short circuits. All the work presented is in the analysis and simulation stages and machine faults are not considered. Scezeney et al [43] model faults in an induction motor drive. Ronkowski [42] undertakes a simulation of converter fed synchronous drives in various fault conditions.

Renfrew and Tian [46] list: open circuit rectifier diode; an open circuit capacitor in the DC link capacitor bank; an open circuit power switch or anti-parallel diode; and reduced base drive faults in a single phase rectifier and 3 phase inverter as faults to be considered. Simulated and experimental results for the inadequate gate drive fault are presented. This fault is introduced in the test rig and the simulation by adding a resistor in series with one IGBT. This does not take account of the increased influence of the Miller capacitance with reduced gate drive.

Moo et al [82] simulate the effect of one or more open circuit thyristors in a controlled 3 phase rectifier supplying the armature of a DC machine. This would be of relevance here

if the drive were supplied from a 3 phase supply via a controlled rectifier, rather than directly from a DC power distribution bus.

The possible faults in a switched reluctance drive are briefly considered by Stephens [24] in his paper on fault detection and management. The phase open and short circuits, pole short circuit, phase to ground and phase to phase faults are listed along with the average torque reduction, the effect on machine vibration and the level of fault current. The work of Stephens is added to by Miller [25] who looks at currents, torques and unbalanced magnetic pull in a switched reluctance drive with an open circuit pole winding fault and with a short circuit pole winding (or shorted turn). Miller does not claim to cover all faults, only the ones which are particularly serious in terms of unbalanced forces in the machine. Belfore et al [71] use artificial neural networks and genetic algorithms to predict faulted and fault-free performance in an SRM. Significantly, this work considers open circuit transistor faults and partial armature short circuits.

Lipo et al [45] investigate the possibility of a fault tolerant synchronous reluctance machine. The work concentrates on post fault control which attempts to control the machine flux so that no flux links the faulted phase. This is presented as an alternative to a machine with effective magnetic isolation between phases. The penalties for this approach were high in terms of lost torque with the faulted machine producing only 0.3 p.u. torque.

Wallace and Spee [33] simulate the effects of a variety of faults in a 3 phase brushless DC drive under voltage and current control. The purpose is to understand the 'symptoms of failure' and possible 'remedial strategies'. They consider an open circuit phase, open circuit transistor, an inadequate base drive and an open circuit hall effect rotor position sensor. Significantly they state that the most likely motor winding failure is the short circuited turn due to insulation failure.

Nyamusa and Demerdash [32] look at the effect of a turn to turn short circuit in a 3 phase PM brushless DC machine. It is shown that the fault seriously degrades current control and results in much increased torque ripple. This work observes that the current induced in the shorted turns is many times rated current, so operation with this fault is not likely to be sustainable.

Oliveira et al. [18] present a brushless DC drive with six independent phases designed as a demonstrator for naval propulsion. This machine may have some similarities to the demonstrator machine here but it is difficult to be certain as the level of mutual coupling between phases is not mentioned. A phase terminal short circuit is simulated but results in a near sinusoidal fault current of 120 A peak in a phase with a rated current of 29 A rms so this is certainly not sustainable. The terminal short circuit results in more than 100% torque ripple (i.e. periods of negative torque in the combined output of all 6 phases). Moreover, it is difficult to see how the fuses relied on to interrupt the fault current could be placed to protect the drive from short circuit faults within the machine. The paper shows the torque produced when the drive is operated with only one of the 6 phases functioning, demonstrating the functional independence of each phase.

Jack and Mecrow [35] present simulation results for a PMSM which show that, with the correct parameters, the fault current resulting from a phase terminal short circuit is limited to close to 1.0 p.u.

Bianchi et al [44] investigate the effects of failures in brushless DC motor drives operated in the field weakening regime. Although Chapter 4 investigated the benefits of operating the machine in the flux weakening regime, this is not done in order to increase the speed range of the drive and there was no advantage in flux weakening to a point where the peak back-EMF would be greater than the DC link voltage. Therefore the failures with the degree of flux weakening proposed in Chapter 4 do not have the serious consequences of the faults presented by Bianchi, such as over charging of the DC link.

5.4 Fault Consequences

The first step in examining each fault must be to establish the effects of the fault if it remains undetected. This will help to define the time available for fault detection. In the case of certain faults, a post fault control strategy will be necessary. Detection techniques and post fault control will be covered in subsequent chapters. In the following sections waveforms are presented with various faults imposed on one phase. Except where stated otherwise, the drive is operating at 4000 rpm with the current controller attempting to drive sinusoidal current in phase with the back-EMF. Where torque ripple is discussed, in

the following sections, the peak value is used. Where the total torque output of the drive is mentioned, it refers to the average value.

5.4.1 Machine Phase Open Circuit

This is the most benign fault which could occur, which is why most protection schemes attempt to convert any short circuit fault to the open circuit condition by means of fuses, a circuit breaker or other similar element. In the event of a phase failing open circuit the phase will cease to contribute any torque at all. Similarly the phase will not draw any power at all from the DC link. The fault will not propagate, even if it remains undetected.

The waveforms resulting from operation of the drive with phase A open circuit are shown in Figure 45. Figure 45(a) shows there is no current in the open circuit phase but that the current driven into the adjacent healthy phase remains unaffected. (Compare this figure to the oscillogram in Figure 37 of Chapter 4). Figure 45(b) shows that the

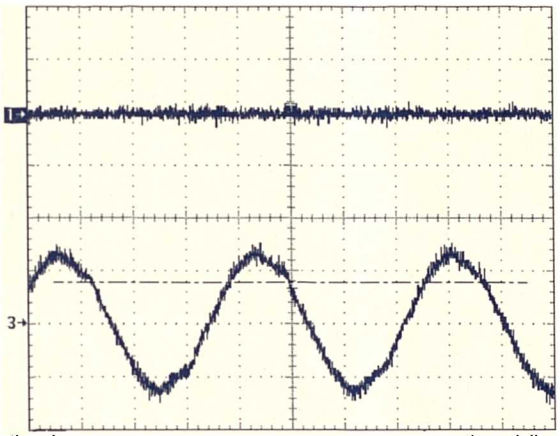

(a) Measured Current Waveforms

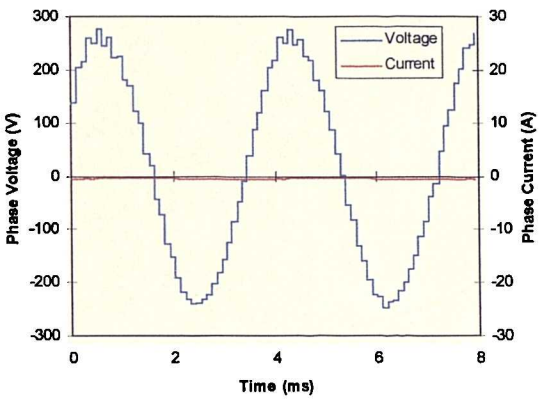

(b) Sampled Phase Current and Voltage Demand

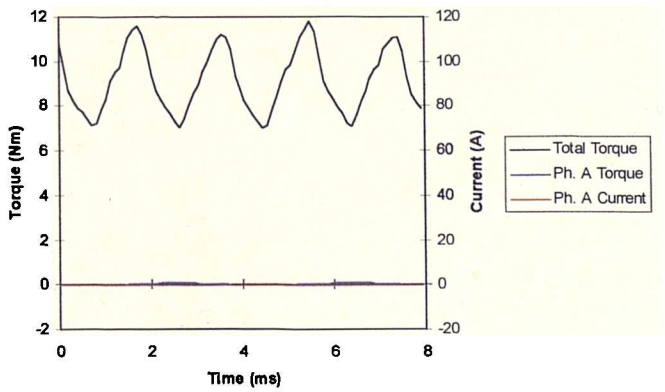

(c) Calculated Torque Output

Figure 45 : The Effect of an Open Circuit Phase Fault at 4000 rpm

open circuit fault results in a large voltage demand for the faulted phase. Figure 45(c) shows that the since there is no current in the faulted phase no torque is delivered by the faulted phase. The loss of the torque from a phase is about 1/6 less net torque and more torque ripple. The peak torque ripple, defined in (23), has increased from 6% to 26%.

$$\text{Peak torque ripple} = \frac{\text{Peak to peak torque ripple} / 2}{\text{Average torque delivered}} \times 100\% \quad (23)$$

5.4.2 Machine Terminal Short Circuit

This could happen at the inverter side or the machine side of the phase current sensor. The machine side is more likely, since the phase current sensor is typically close to the inverter.

With a short circuit applied to the terminals of one phase, the current that is driven by the back-EMF is limited to rated current by the phase reactance. This enables the machine to withstand a phase terminal short circuit indefinitely and continue operation on the healthy phases. Because the current in a shorted phase is reactance limited, it is almost exactly 90 degrees out of phase with the back-EMF. Therefore the average torque contributed by the faulted phase is very small.

Figure 46 shows the drive operating with a terminal short circuit. Note that the short circuit was applied to the phase through the inverter. This is achieved by turning off both upper switches and turning on both lower switches in the H-bridge feeding the phase. Thus a pair of power devices is also incorporated into the short circuit path. This results in additional loss and thus a little more braking torque. The significance of a phase short circuit applied in this way will be made clear later.

Figure 46(a) shows that the current in the faulted phase is limited to rated current - having an RMS value of 18.6 A, compared to 18.2 A for the adjacent healthy phase. The phase current in the faulted phase would lag the phase back-EMF by 270 degrees (lead by 90 degrees) if there were no resistive losses. However, from the oscillogram, it is apparent that the short circuit current in the faulted phase (phase A) lags the current in the healthy phase (phase B) by about 198 degrees and thus must lag the phase A back-EMF by 198 + 60 = 258 degrees. The difference between the expected and actual phase

ORIGINAL IN COLOUR

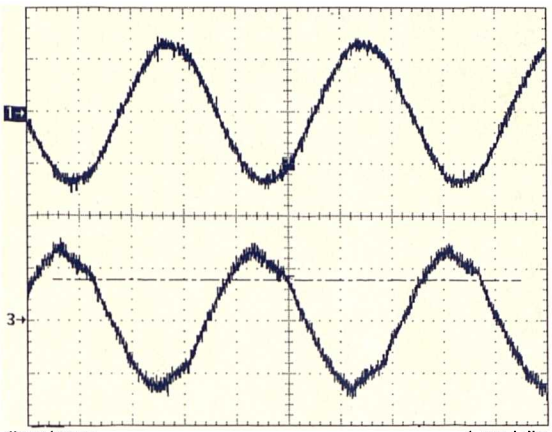

(a) Measured Current Waveforms

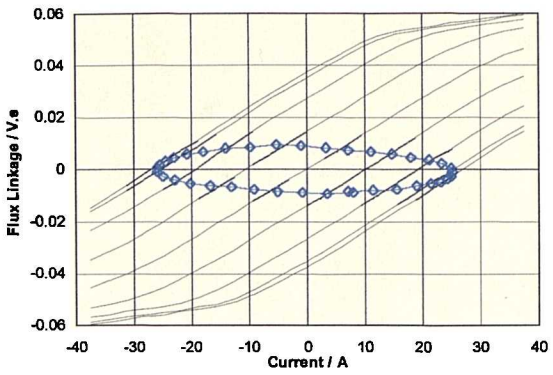

(b) Calculated ψ-i Locus for the Faulted Phase

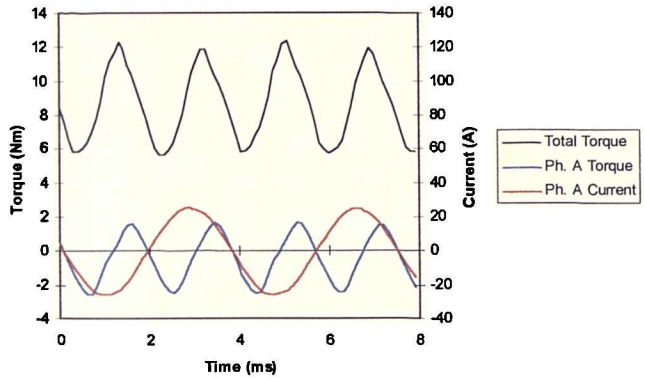

(c) Calculated Torque Output

Figure 46 : Machine Phase Short Circuit Applied Through the Inverter at 4000 rpm

angle of the short circuit current is a result of the resistance of the phase and the volt drop in the inverter. The non-orthogonal relationship of the current and the phase back-EMF results in a relatively small area enclosed on the ψ-i diagram. This area corresponds to a net braking torque but comparison with Figure 37 of Chapter 4 shows that the braking torque is small compared to the torque contribution of a single healthy phase.

Note that the waveforms have been measured with the machine running at approximately 1/3 speed. At full speed the resistance will be a smaller part of the phase impedance and the short circuit will provide less net braking torque. The assertion that a shorted phase does not contribute much braking torque is born out by Figure 46(c). The torque contributed by the faulted phase pulsates at twice the machine electrical frequency. (i.e. 2 × number of pole pairs × speed). The average value of the pulsating torque due to the faulted phase is close to zero. The average machine torque is reduced by a little over 1/6

but the torque ripple is increased from 6%, for a healthy phase, to 40% - much larger than for an open circuit phase.

In the transient case the current may reach twice its rated value and take approximately 20 ms to settle into steady state. Settling time is dependent on the L/R time constant.

The effect of this fault on the supply is more serious. If the machine were line fed the fault would constitute a short circuit applied to one phase of the supply. With the phase fed from a single phase inverter, the fault will result in a short circuit applied to the DC link through the power converter. The supply current into the short circuit will be limited by power device gain. Clearly it is essential that this fault is detected within the short circuit withstand time of the power switches.

5.4.3 Turn to Turn Fault

The adverse conditions of heat and vibration to which electric machines are naturally exposed will occasionally lead to the progressive deterioration and eventual breakdown of the winding insulation in a phase of the machine. Both Wallace and Spee [34] and Penman et al [83] suggest that breakdown of the stator winding insulation is a major cause of motor failure.

In general, insulation failure in the machine may lead to turn to turn faults, phase to phase faults or phase to frame faults. The machine considered here has physical separation between phases, making phase to phase faults extremely unlikely. It is possible to progress from a healthy phase directly to a phase to frame fault. If the machine stator core or the power converter were isolated from ground, drive operation would be unaffected by a single phase to core fault. Moreover it seems likely that most phase to frame faults are the result of the continued operation of a machine with an undetected turn to turn fault, [83]. For these reasons the work here focuses on turn to turn faults

A large body of work exists on the causes and effects of shorted turns in machine windings [32], [84] and [25]. How this type of failure might be detected and strategies for continued operation in the presence of the fault have received little attention.

Analysis of Turn to Turn Faults

The parameters required to examine the effect of a single shorted turn in the machine are given in Table 12. The table contains 2 sets of values, the first being for a small prototype fault tolerant machine and the second being for the 16 kW demonstrator machine. The first half of the table shows that the current in a shorted phase of the demonstrator machine is limited to the rated current for the machine, 20.3 A.

Parameter	Prototype	Demonstrator	Units
Operating speed	12000	13000	rpm
Number of poles	12	8	
Turns per phase	76	50	
Peak phase back-EMF at operating speed	79.9	198.9	V
Phase resistance	0.273	0.156	Ω
Phase inductance	1.10	1.275	mH
Phase reactance	8.29	6.942	Ω
Phase impedance	8.30	6.944	Ω
Current in shorted phase	6.8	20.3	A
Peak back-EMF seen by each turn at operating speed	1.05	3.978	V
Single turn resistance	3.59	3.12	mΩ
Single turn inductance	0.19	0.51	μH
Single turn reactance	1.44	2.78	mΩ
Single turn impedance	3.87	4.18	mΩ
Current in shorted single turn	192	673	A
Estimated mass of single turn	0.51	2.4	g
Rate of rise of temperature with a single shorted turn	674	1545	$°C.s^{-1}$

Table 12 : Machine Parameters Required for Shorted Turn Analysis

If a turn to turn short circuit occurs within a machine winding containing n turns, so as to form a shorted coil of m turns, then the portion of the back-EMF seen by the shorted turns will drive a current of n/m times terminal short circuit current (neglecting resistance). Clearly the most severe case is a single shorted turn. Since this machine contains 50 turns per phase and has a terminal short circuit current of 20.3 A , a single shorted turn will try to carry 1015 A.

In contrast to a phase short circuit, where the reactance dominates, the resistance often plays an important part in limiting the current in a single shorted turn. This difference results because for a phase coil with n turns, $R_{turn} = R_{phase}/n$. but $L_{turn} = L_{phase}/n^2$. The table shows that even at full speed, when the electrical frequency and hence reactance are

at a maximum, resistance makes a significant contribution to the impedance of a shorted turn. Thus the current in the shorted turn is limited to 673A, which is 33.2 times rated current. Finally the table suggests that this large current will cause the temperature of the shorted turn to rise at a rate of 1545 $°Cs^{-1}$. Consideration of the rate of temperature rise is limited to the machine, as the current in the shorted turn does not flow through the power electronics. Clearly the fault must be detected and some action taken in the order of 10 ms.

Some of the work on this fault used a small prototype machine with similar geometry to that of the demonstrator machine. The parameters of the prototype are also given in Table 12. Comparison of the results show that for a smaller machine, with more turns per phase, the fault current is more strongly resistance limited. The result is that the reaction time is longer, but the fault harder to detect in the smaller machine [40].

Imposing a Single Shorted Turn

Phase A of the demonstrator machine was modified to assist investigation of the shorted turn problem. The machine is wound without a star point, so both ends of each phase winding are brought out of the machine as shown in Figure 47. The phase has been modified by attaching an additional lead to the end of the first turn and bringing this point out of the machine. Now a shorted turn may be imposed by connecting the end of the first turn, point Y, to the positive phase terminal, point X. Because the connection is made outside the machine case, it is a simple matter to include a current probe in the circuit to measure current in the shorted turn.

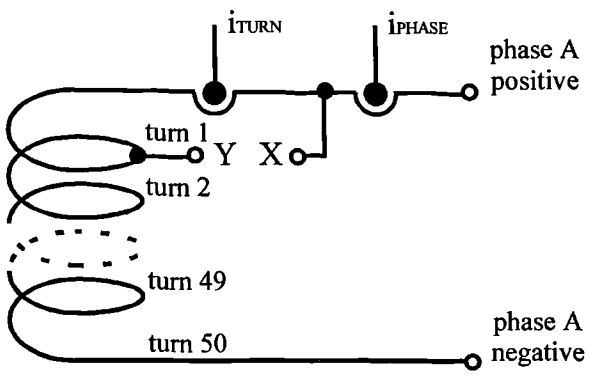

Figure 47 : Modification of Phase A Winding to Allow Shorted Turn Tests

With the machine running, but with phase A open circuit, the current measured in the shorted turn was approximately 140 A rms. The waveform can be seen in Figure 94 of Chapter 9. The current measured in the shorted turn is somewhat lower than predicted because the machine is running below rated speed and because the length of the shorted turn has been considerably increased to allow the inclusion of the current probe.

The ideal shorted turn, i.e. a shorted turn perfectly coupled to the rest of the phase winding and having zero resistance, would have the same torque characteristic as a terminal short circuit.

Effect of a Shorted Turn in an Inverter Fed Machine

Previous sections have looked at the current driven in a shorted turn by the rotating magnet flux. If the machine is to operate as a motor then some voltage must be applied to the terminals of the machine to drive the phase currents required. This applied voltage will induce a voltage in the shorted turn, also resulting in current in the turn.

If the drive is operating in current control then there will be some interaction between the current controller and the faulted phase.

Many of the effects of a shorted turn on an inverter fed drive can be seen in Figure 48. The drive has been operated for a short period with a shorted turn in one phase. The red trace is the current in the adjacent healthy phase and this is unaffected by the fault.

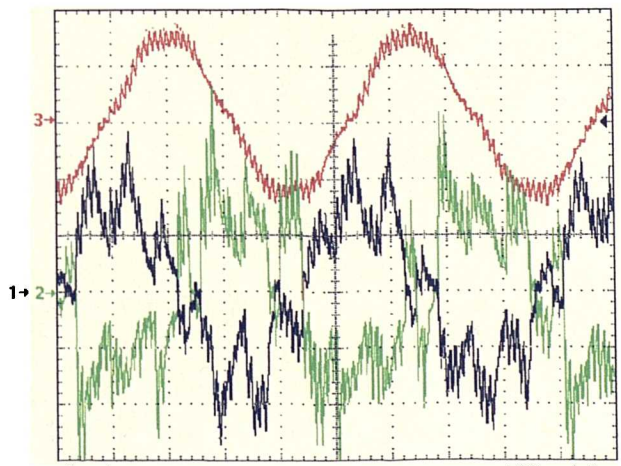

Figure 48 : The Effect of a Shorted Turn in an Inverter Fed Machine at 7000 rpm

ORIGINAL IN COLOUR

The black trace shows the current in the faulted phase. The current driven into the faulted phase clearly includes increased level of ripple at the PWM frequency. The green trace shows the current driven in the shorted turn as a result of voltages induced by the passing magnet flux and by the changing flux resulting from volts applied at the terminals of the faulted phase. The current in the shorted turn has a peak value in excess of 300 A. Whilst the disturbance of the phase current in the faulted phase may be tolerable, the levels of current induced in the shorted turn clearly are not.

The oscillogram in Figure 49 shows quantities sampled by the controller. The lower pair of traces compare the sampled phase current and phase current demand for a phase with a shorted turn. The upper pair of traces compare the sampled phase current and phase current demand for one of the healthy phases adjacent to the faulted phase. Note that a back-EMF shaped current profile has been used for the current demand.

In can be seen that the current in the healthy phase is well controlled. In contrast the current in the faulted phase oscillates about the demand value. This oscillation results from the interaction of the current controller with the reduced reactance seen at the terminals of the faulted phase. It is worth noting that the adjacent healthy phase only remains unaffected by the fault because there is very little mutual coupling between phases in this machine type. This is not the case in most synchronous PM machines, [32].

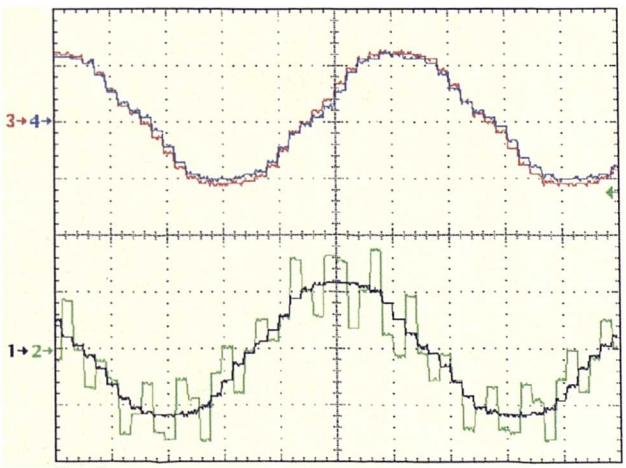

Figure 49 : The Effect of a Single Shorted Turn in Phase A on Current Control

5.4.4 Power Device Open Circuit

This type of fault may affect the IGBT or the anti-parallel diode or the IGBT-diode pair. Each of these possibilities will be discussed below.

Diode Open Circuit

Craig [50] acknowledges the possibility of diode open circuit faults but suggests that the effects are machine dependent. The analysis presented here questions this assertion. Renfrew and Tian [46] observe that an open circuit freewheel diode will rapidly lead to further faults but do not consider this situation further.

To investigate the effect of a freewheeling diode failing open circuit, first consider normal, fault-free, operation of the drive phase in Figure 50. Suppose forward current is flowing in the phase, from the supply, via T1 and T4. When T4 turns off, the voltage at the negative inverter output terminal (T4 collector) will rise, being driven up by the winding inductance. When the voltage rises above the DC link by a diode forward voltage drop (approximately 1 to 2 V) D3 will start to conduct, clamping the voltage to this level.

However if the diode, D3, is open circuit, the voltage will continue to rise to about 20 V above the positive rail at which point T3 will suffer reverse breakdown. Here it is assumed that no other part breaks down first - such as T4 suffering forward breakdown or D4 suffering reverse breakdown. Reverse breakdown in the IGBT will rapidly cause the failure of this device. It is most likely that the damaged IGBT will fail short circuit. The resulting situation will be a shorted power switch. Allowing fault propagation is undesirable, but in this instance it is seen as unavoidable since the IGBT will fail very rapidly, making it extremely difficult to detect the diode open circuit in time to prevent fault propagation.

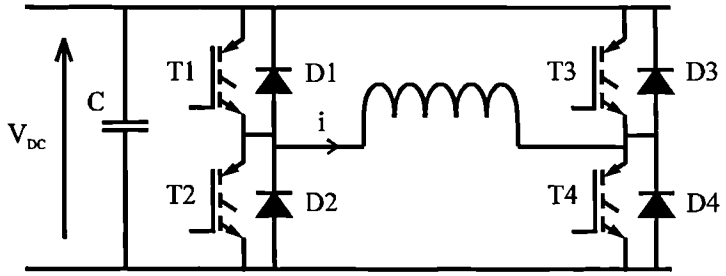

Figure 50 : Simplified Single Phase Drive Schematic

If a detection time of approximately 3 μs were acceptable, then this fault could be detected by the transfer current sensing scheme described in Chapter 6. Upon detection of this fault, the healthy transistor in series with the diode would have to be turned back on. Then the switching of the healthy leg must be controlled to drive the current back through zero At this point the healthy diode in the faulted leg will be conducting and all transistors may be turned off. The faulted phase may be operated in unipolar mode using only the diagonal pair of transistors with healthy series connected diodes.

IGBT Open Circuit

The effect of this fault is to reduce the functionality of the single phase bridge from bipolar operation to unipolar operation. If one of the 2 forward IGBTs (T1 or T4 in Figure 50) is affected, the bridge will lose the capability to drive forward current. If forward current is flowing at the time of the fault, then the current will flow in the diode in anti-parallel with the healthy transistor in the faulted leg. Note that negative currents may still be controlled.

An open circuit IGBT does not require rapid detection. The waveforms in Figure 51 show that while the fault remains undetected, the affected bridge will continue to drive unipolar currents into the phase. This allows half sine current to be driven in the machine winding and results in the average torque in the machine being reduced by half that of a healthy phase.

The ψ-i diagram reveals that the loss of bipolar current operation denies access to one half of the ψ-i plane - thus resulting in the loss of half the previously available torque. The torque waveform shows that the loss in average machine torque is indeed approximately 1/12 of that delivered by the healthy drive. The magnitude of torque pulsation introduced by this fault is the same as that for an open circuit winding, but at half the frequency.

Clearly this fault is less serious than the loss of a phase, as it results in the loss of only half of the torque contribution of the faulty phase, provided that it is considered prudent to continue switching devices in the faulty phase.

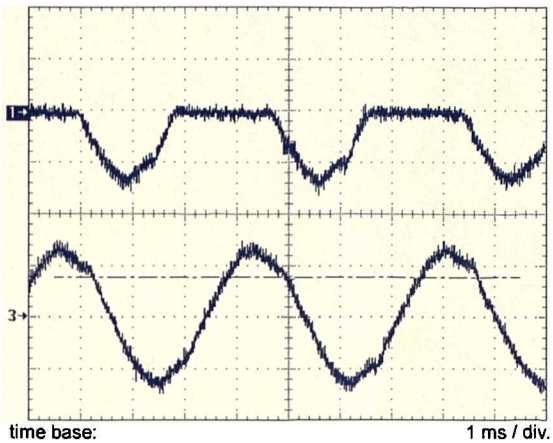

(a) Measured Current Waveforms

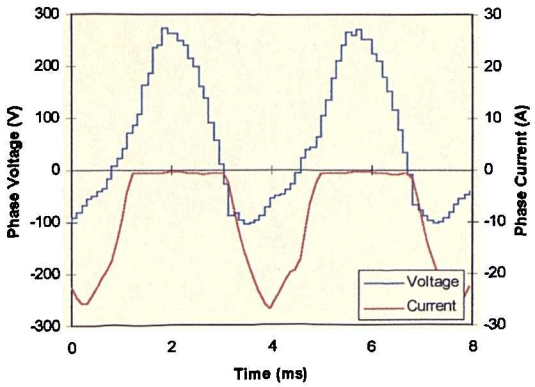

(b) Sampled Phase Current and Voltage Demand

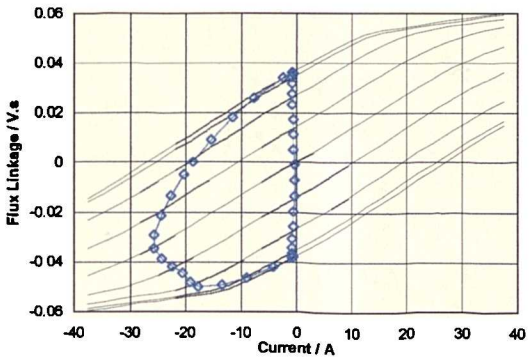

(c) Calculated ψ-i locus for the Faulted Phase

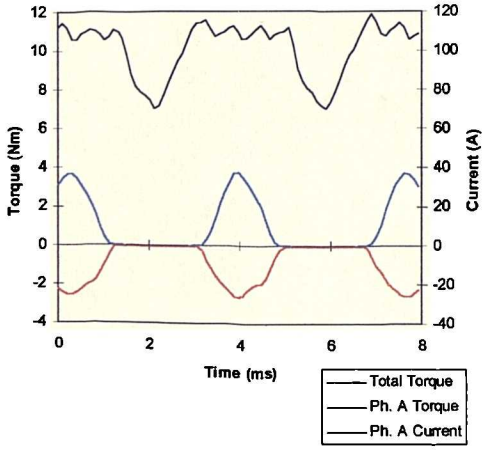

(d) Calculated Torque Output

Figure 51 : Effect of IGBT Open Circuit at 4000 rpm

Switch Open Circuit

The loss of a bond wire or a fracture in a track could result in both an IGBT and the corresponding anti-parallel diode becoming open circuit as the result of a single point failure. This failure will be referred to as a power switch open circuit fault. Whilst the type of mechanical failures necessary to bring about a power switch open circuit fault may seem to be only a remote possibility, some evidence suggests that mechanical failures caused by thermal cycling can be a major cause of power circuit failure. Thermal cycling of the power circuit occurs naturally in applications where the power output is subject to large and frequent variations.

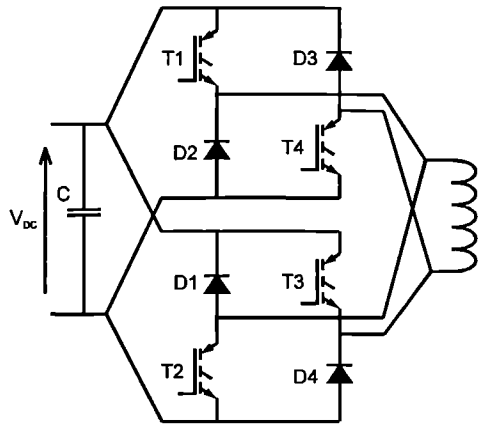

Figure 52 : H-Bridge Arranged to Prevent Power Switch Open Circuit Faults

This fault will result in the healthy IGBT in the faulted leg suffering forward breakdown. The IGBT in which the forward breakdown occurs will rapidly fail short circuit. This type of fault may be eliminated by appropriate arrangement of the tracking and bond wires. Figure 52 shows a connection diagram for a scheme which eliminates the possibility of a power switch open circuit fault resulting from a single connection failure. In this scheme the circuit schematic has not been altered, but the 2 unipolar half bridges inherent in the single phase bipolar bridge have been physically separated. Note that the current only transfers between bridges when the load current polarity changes, so the physical separation of the devices into 2 unipolar bridges does not compromise the inductance of the critical paths in which rapid current changes occur during normal switching operation.

Gate Drive Open Circuit

An open circuit gate drive will naturally result in intermittent operation of the associated power device. If the healthy power device in this limb turns on, the voltage across the faulted switch will rise and the device will tend to turn on due to charging of C_{GE} through C_{CG} (Miller effect). This effect appears to have been over looked by Renfrew and Tian [46] who consider inadequate gate drive faults. Wallace and Spee [34] also look at this fault for bipolar junction transistors (BJTs), suggesting that it is worthy of attention. The intermittent operation resulting from this fault will make detection and post fault control more challenging.

5.4.5 Power Device Short Circuit

A short circuit freewheel diode, a short circuit IGBT or a gate drive circuit failure resulting in a gate drive remaining on, will all have the same effect. All these failures will cause a shoot-through fault when the healthy device in the faulty leg turns on. The result will be that the DC link is shorted. Duong [85] shows that initially the current drawn is limited by the gain of the healthy IGBT. If the fault persists, the device structure will be destroyed by localised overheating. With the device structure destroyed the IGBT gain can no longer limit the current, which will rise rapidly towards a much higher level. Some time later the IGBT will fail open circuit. The rupturing of the IGBT is unpredictable. Very large power levels can be present in the IGBT prior to it rupturing and these will typically lead to violent bursting of the IGBT package, at least in the case of commercial modules.

As excessive current is drawn from the DC link, the DC link capacitance will discharge and the DC link voltage will fall. This fault will have a serious effect on the operation of other phases and indeed other systems sharing the common DC link if the fault is allowed to persist for long enough to significantly discharge the DC link.

In the demonstrator drive the shoot-through current has been measured to be 300 amps. This current would rise higher if it were not limited by the gain of a healthy IGBT in the shoot-through path. Many IGBTs are only capable of withstanding the short circuit condition for 10 μs. Therefore rapid detection of power device short circuits is essential.

If the remaining 3 healthy devices in the bridge are then turned off and the machine continues to rotate then a further problem will result from the faulted power device. The machine back-EMF will drive current round a path formed by the machine winding, the faulted power device and the diode in the healthy limb. The current path resulting from

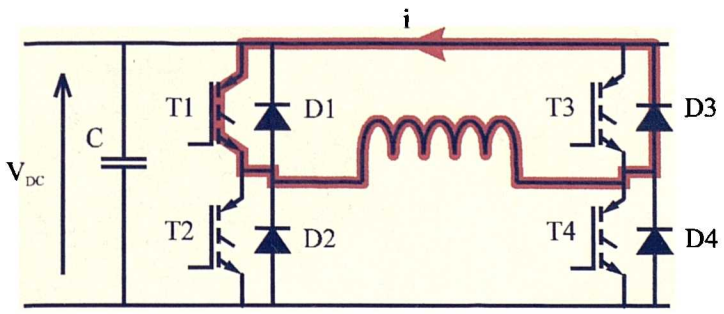

Figure 53 : Unipolar Current Path Resulting from IGBT Short Circuit Fault

T1 failing short circuit is shown in Figure 53. The presence of a diode in the path results in a unipolar current with an rms. value approaching $\sqrt{3}$ times rated current. This situation is shown in Figure 54. Note that the zero current level for the upper trace is marked by '1→' at the left edge of the plot and that the fault current is unipolar. The fault current was measured to be 24.8 A RMS, (compared to the 18.1 A *$\sqrt{3}$ = 31.3 A predicted). Both the lower than predicted value and the discontinuous nature of the fault current are a result of losses in the power devices and machine winding. Figure 55 shows the fault current resulting from the same fault at 12000 rpm rather than 4000 rpm. At the higher speed the reactance has increased but the resistance remains unchanged. Thus the fault current is closer to that predicted. As expected the current in the adjacent healthy phase is seen to be unaffected by the fault.

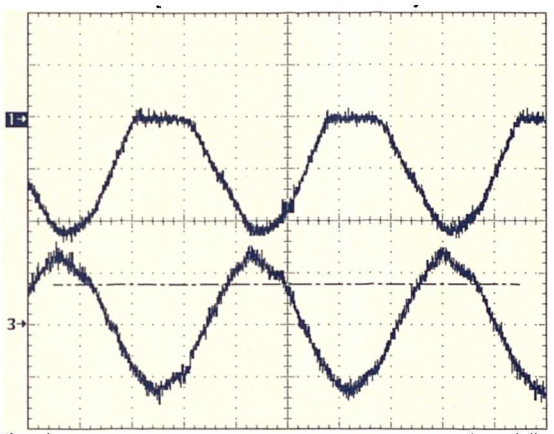

(a) Measured Current Waveforms

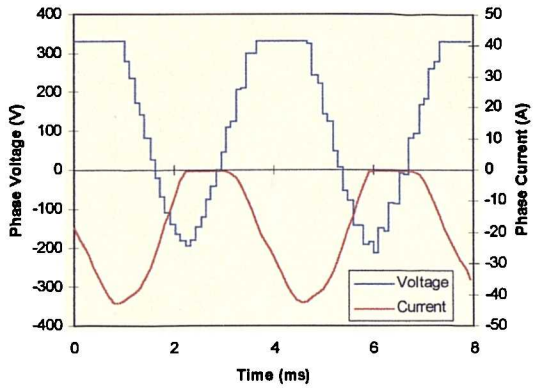

(b) Sampled Phase Current and Voltage Demand

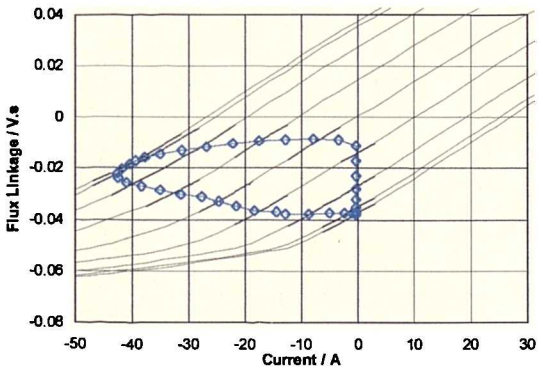

(c) Calculated ψ-i locus for the Faulted Phase

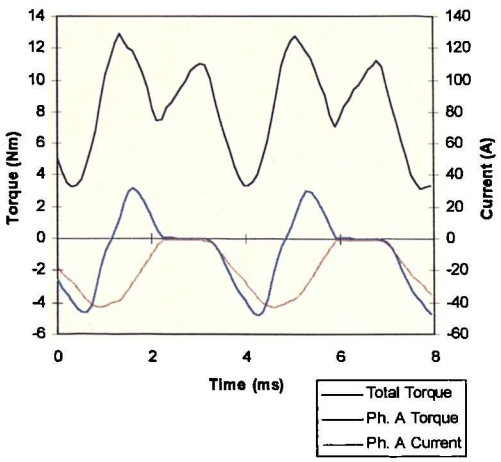

(d) Calculated Torque Output

Figure 54 : Effect of Power Device Short Circuit Fault at 4000 rpm

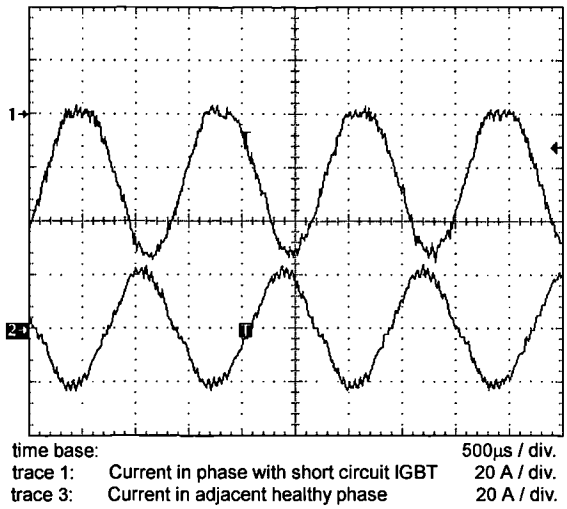

Figure 55 : Effect of Power Device Short Circuit Fault at 12000 rpm

The ψ-i diagram in Figure 54 shows that the fault results in more enclosed area, and thus more braking torque, than the phase short circuit fault. The braking torque is again the result of losses. This braking torque is higher than in the phase terminal short circuit case due to higher losses arising from larger (RMS) current. The torque ripple is very high at 58% of the total torque delivered with this fault.

Even when operating at 4000 rpm, when the phase resistance helps to limit the fault current, the fault current is greater than the rated current. Some post fault control strategy is required to reduce the fault current if operation is to be sustainable without an excessive temperature rise in the machine phase winding.

5.4.6 Reservoir Capacitor Faults

If the fault tolerant capacitor network presented in Chapter 2 were adopted, the effect of an individual capacitor in the network failing would be minimal. A short circuit capacitor would effectively result in a slight increase in the total capacitance of the network. Similarly the open circuit failure of a capacitor would result in a reduction of the total capacitance. The size of the change in capacitance in each case would depend on the level of partitioning and redundancy in the capacitor network. In either case the effect on the supply and the torque output of the affected phases would be minimal.

5.4.7 Fault Summary

In the preceding sections, each fault has been discussed in detail. The effects of each operating condition are summarised in Table 13, along with the torque capability with rated current at 4000 rpm.

Fault	Fault Current / A RMS	Temperature rise in faulted phase / °C	Calculated Torque / Nm	Lost torque	Time before fault propagation
No fault	18.1	115	11.06	0.0 %	indefinite
Machine phase open circuit	0.0	25	9.12	17.5 %	indefinite
Machine phase terminals shorted through the bridge	18.6	103	8.70	21.3 %	indefinite
IGBT open circuit	12.8	66	10.04	9.2 %	indefinite
Single power device short circuit (Gate drives remaining active)	~300	-	-	-	10 μs
Single power device short circuit (Gate drives inhibited)	24.8	N/A.	8.47	23.4 %	several seconds
Single shorted turn	~300	-	-	-	10 ms

Table 13 : Fault Effect Summary for Operation at 4000 rpm

In terms of lost torque, the most serious of the sustainable faults is the terminal short circuit applied through the power converter. As such this is the fault which determines the degree of redundancy which must be built into the drive. It is important to remember that these tests were performed at 1/3 rated speed and that the lost torque resulting from the short circuit faults will reduce for higher speeds.

It is interesting to note that the phase terminal short circuit results in slightly higher current but a lower phase winding temperature, compared to healthy operation at the same operating point. The lower phase temperature is likely to be the result of 2 factors:

- The flux in the faulted phase changes more slowly, as it is no longer subject to PWM switching voltages.
- The fault results in a high degree of flux weakening in the faulted phase, this is clear from the psi-i diagram in Figure 46.

5.5 Conclusion

A short list of faults most likely to occur has been drawn up. It has been demonstrated that the drive here is capable of operating indefinitely with a power device or a phase open circuit or even with a phase terminal short circuit. If this were not the case an additional iteration through the drive design and fault analysis processes would be necessary.

The fact that there are losses means that the introduction of a phase terminal short circuit through the inverter results in a torque reduction of 20.3% at 4000 rpm, rather than 16.7% for a loss-less drive. This means that the drive must be designed with 27 % redundancy (rather than the 20 % redundancy required in a loss-less drive) if the drive is to continue to deliver rated power in the event of this condition. It is important to remember that the torque reduction was measured at one third of rated speed and that the lost torque resulting from the short circuit fault will reduce for higher speeds. Thus 27 % redundancy is a conservative figure.

However a single shorted turn in the machine or an undetected power device short circuit failure results in very large fault currents and correspondingly serious consequences. Without rapid detection and appropriate post fault control these faults will propagate rapidly and/or disturb the operation of healthy phases. It is often argued that these are also the most likely faults [4], [34]. Both of these faults result from the breakdown of a thin insulating layer, formed by the silicon structure in the IGBT and by insulating lacquer in the machine, so the relatively high likelihood of these faults seems reasonable.

Chapter 6

SENSORS

6.1 Objectives

The achievement of fault tolerance depends on the sensing of faults. It is essential to detect and report any failure so that a unit may be replaced before further deterioration occurs. Some of the more serious faults require very rapid detection so that appropriate action may be taken in time to prevent fault propagation and/or a serious impact on the system as a whole.

This chapter will start by considering the types of sensors used for fault detection, then a number of promising candidates will be examined in more detail. Of the sensors qualifying for detailed discussion, some are completely new, whereas others are simply developments of existing ideas.

The aim of this chapter is to develop a set of sensors suitable for fault detection on the demonstrator drive.

6.2 Existing Sensors for Fault Detection

All fault detection schemes require some information about the drive and thus sensors are essential to fault detection. The information required can be provided by sensors which are already included in the drive or by sensors which are added for the sole purpose of fault detection. Most authors examining condition monitoring produce fault diagnosis based on some form of analysis of the phase currents. Spectral analysis of currents is presented as a means for monitoring the health of an induction machine by Schoen et al, [86]. A number of authors investigating the possibility of applying expert systems to fault diagnosis in power converters, such as Aris et al [48], also rely on measurements of phase currents and voltages.

Condition monitoring schemes occasionally use additional sensors to measure the flux in some part of the machine. An example of this approach is the work by Penman et al [83] which places a set of search coils at one end of the machine to measure axial leakage flux. It is claimed that this technique is sufficiently sensitive to detect machine faults with the search coils outside the machine frame.

Whilst detecting faults without any additional sensors is a desirable goal, using even the simplest additional sensors in the machine or power electronics can be a very powerful tool in enabling fault detection. This is particularly true if rapid detection of some of the more serious faults is to be achieved. For example one type of additional sensor, the on-state voltage sensor, is commonly used to protect power devices from externally applied short circuits. This sensor is discussed in more detail when the principle is developed later in this chapter.

Whilst examining fault detection in a SRM, Stephens [24] observes that certain faults have no effect on the phase current and therefore cannot be detected by the phase current sensor. Stephens proposes additional sensors to provide more comprehensive fault detection including a differential current detector, a rate of rise detector for phase current and a differential flux detector. The last of these is based on a search coil measuring flux in opposite teeth in an SRM. Under normal operation the flux in opposite teeth should be equal. This idea is developed by Miller [25] who replaces search coils measuring the flux in each tooth by measurement of the voltage developed across each series connected coil.

In 1981 Evans et al [87] split the familiar phase leg comprising two power transistors and two anti-parallel diodes into forward and reverse legs. A choke is inserted between the outputs of the two resulting legs as shown in Figure 56. Normally the choke will carry the phase current. However in the event of a shoot-through fault caused by the simultaneous conduction of both transistors in the modified phase leg, the rate of rise of fault current will be limited by the choke. The fault may be detected by sensing current in the choke. Modern IGBTs and gate drive circuits, which limit the short circuit current and allow the device to be turned off safely, make the use of the choke unnecessary. However, this scheme is still important as it recognises that the phase leg is really

composed of two independent unipolar legs, a principle exploited by the transfer current sensor presented later.

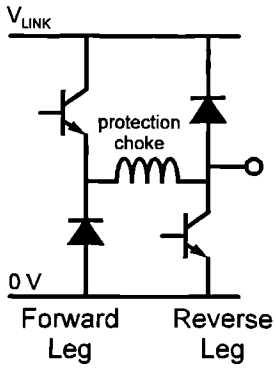

Figure 56 : Modified Leg Arrangement Proposed by Evans

Carvalho et al [88] incorporate a current sensor into each limb in a novel fashion, shown in Figure 57. The current sensor measures the phase current but also detects any shoot-through current. It is claimed that positioning the phase current sensors in this way provides shoot-through fault detection without requiring any additional sensors. However it should be noted that this scheme uses one sensor per limb, in contrast to the conventional phase current sensing configuration in which the current in one phase is calculated from the sum of the currents in the other phases, providing a saving of one sensor. Furthermore the inclusion of the current sensor in the position proposed will increase the limb inductance, forcing slower device switching in normal and fault conditions.

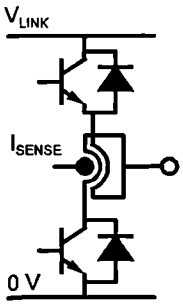

Figure 57 : Combined Sensing of Phase Current and Shoot-Through Current

Finally, Craig et al [50] place an additional sensor in the bridge to measure the freewheeling and shoot-through currents. The output of the sensor is correlated with the phase current sensor output and the gate drive commands to determine the health of the bridge, an idea which is developed later in this section.

6.3 Phase Current Sensor

It is possible to detect many faults in the drive simply by monitoring the phase currents. The phase current monitoring for fault detection is particularly attractive because phase current sensors are already present in most drives. In conjunction with a suitable detection scheme this sensor can be utilised for detecting turn to turn faults and phase to ground faults as well as open circuit phase or open circuit power switch. Schemes for fault detection which operate by comparing the measured phase current to either an expected value or currents measured in other phases are described in Chapter 7.

6.4 On-State Voltage Sensor

6.4.1 The Principle

The voltage drop across an IGBT or MOSFET varies with the current flowing in the device. In an IGBT, the variation of on-state voltage with current is highly non-linear, as shown in Figure 58. However, in the normal operating region the voltage drop is almost constant, rising only a little with increasing current. Some schemes use the on-state voltage drop as an indication of the current flowing in the device. When applied to the IGBT, this technique is often refereed to as V_{CESAT} sensing, since it senses the collector-emitter voltage when the device should be fully turned on. If the scheme is used to detect short circuits it is more correctly referred to as de-saturation sensing since, during a short circuit, the IGBT is operating in its linear region. In the linear region the device supports the full DC link voltage with the current limited by device gain. For a typical 100 A IGBT, with a gate drive voltage of 15 V, the device gain will limit collector current to

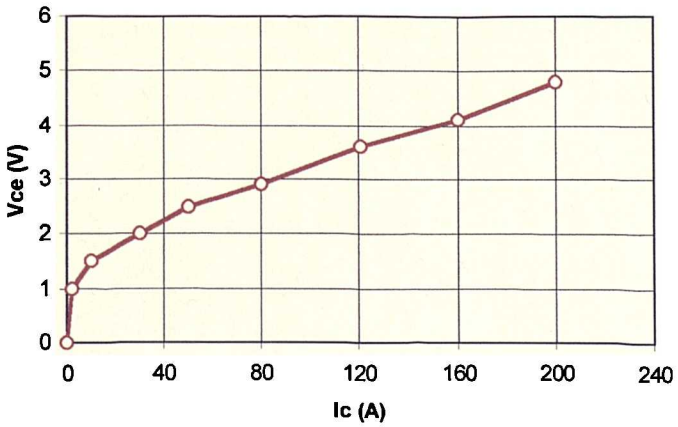

Data taken from Semikron data sheet for SKM100GB101D.

Figure 58 : On-State Voltage verses Current Characteristic for a 100A, 1000V IGBT

between 300 A and 600 A.

On-state voltage sensing is widely used to protect IGBTs from externally applied short circuits. There do not appear to be any examples in which on-state voltage sensing has been used to detect faults occurring within an inverter.

Castino et al [89] suggest the use of V_{CESAT} sensing to detect over-current. They present a scheme for reducing the gate voltage to extend the short circuit withstand time. This loses some of its significance if 10 μs short circuit rated devices are used. However a slow turn off, with the possibility to turn on again if the fault clears in the turn off time, may be useful. Chokhawala et al [63] revisit the V_{CESAT} sensing technique, observing that there are some differences between the situation where a power device turns on into a short circuit and the situation where a short circuit is applied to the output of a conducting IGBT. This work recognises that the shorted IGBT can fail through exceeding the thermal limit, latching or exceeding the voltage rating.

Ekel et al [90] present a number of gate drive schemes which use either V_{CESAT} sensing or current shunt sensing to provide device short circuit detection whilst at the same time reducing the short circuit current and voltage overshoots. The possibility of using the V_{CE} voltage monitor to control the rate of turn off in normal and short circuit operation is explored by Chokhawala et al [91]. The same idea is further developed by Gediga [92] to control the rate of gate charging from V_{CE}.

The ultimate vote of confidence for V_{CESAT} sensing schemes is provided by the high take up of these schemes by the manufacturers of gate drive circuits. Valentine [93] protects his smart power module with V_{CESAT} sensing, using a larger value of resistance in the gate discharge path, R_{OFF}, to provide slow turn off in the event of de-saturation. Semikron [94] use V_{CESAT} sensing to protect IGBTs driven by their gate drive modules, but do not appear to make any allowance for reduced turn off rate required to avoid over voltage transient.

In spite of all the enthusiasm for the use of on-state voltage sensing in the control of IGBTs, the future of on-state voltage monitoring as an indication of device current is threatened by developments in current sensing IGBTs. Liang et al [95] present a design for a current sensing IGBT using one cell of the device as a current mirror. The research

into these current sensing devices is clearly at an advanced state, since Shen et al [96] present results from such a device, including performance under short circuit conditions.

6.4.2 Implementation

Commercially available gate drive modules often implement V_{CESAT} sensing. These circuits are configured to provide effective protection of the power devices in the event of a short circuit load. When used to detect faults within a fault tolerant drive they impose some undesirable limitations. The two principle limitations arise because the action taken in the event of a fault is fixed within the driver circuit. The result is that both devices in the faulted phase leg are forced off for a period after any fault detection and that there is often a significant delay between fault detection and annunciation.

In the early stages of demonstrator drive development the gate drive signals were supplied by gate drive modules manufactured by Concept, [97]. The short circuit protection function of these modules was tested and the delay from fault occurrence to fault annunciation was measured to be 120 μs.

A gate drive circuit has been designed without the limitations described. This gate drive circuit embodies an on-state voltage sensor and the output of this sensor is fed out of the circuit with a delay of 0.2 μs. The gate drive circuit does not include a device protection mechanism, such as that found in the commercial gate drive circuits using on-state voltage sensing. Instead a central detection circuit will receive fault signals from all the power switches in a bridge, enabling device protection at bridge level rather than at phase leg level. Monitoring faults for a complete bridge also permits more flexible fault diagnosis.

The gate drive circuit design fulfils three functions as shown in Figure 59. These functions are to:

- Provide an isolation barrier between the power circuit and the control circuits. The isolation is required because the gate drive circuit is referenced to the emitter of the IGBT which it controls and the emitter voltage swings between the positive and negative DC link rails during the normal switching operation of the bridge. The power supplies for the gate drive and on-state voltage sensing circuits must be transferred across the isolation barrier.

- Transfer the gate drive command across the isolation barrier to a driver circuit capable of driving the gate between its on and off voltage levels.
- Compare the voltage across the power device whilst it is turned on to a reference level and transfer the resulting fault indication back across the isolation barrier.

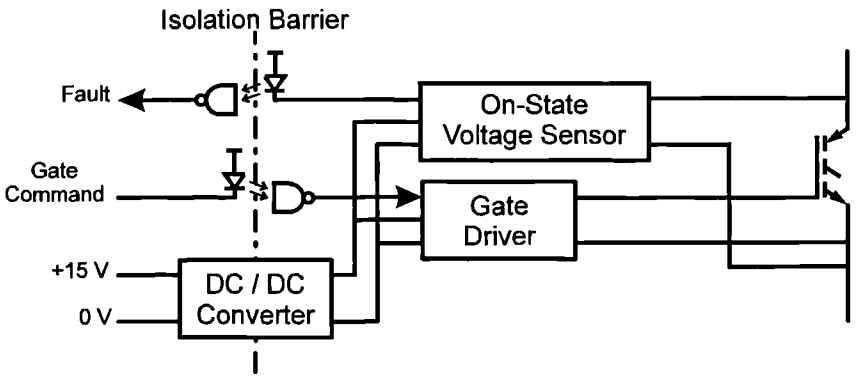

Figure 59 : Functional Blocks of a Floating Gate Drive Circuit

The gate drive circuit has been designed to provide the correct functionality. The design has not focused on providing high reliability or fail-safe characteristics. Each element of the resulting gate drive circuit is described below, beginning with the gate driver block.

The gate driver must provide a low impedance floating supply to drive the gate emitter voltage between the voltage levels required to turn the IGBT on and off. The voltage level required in the on-state is set by the device manufacturers at 15 V. The voltage used to turn the device off could be 0 V, but it is more common to choose a voltage in the region of -5 V to -15 V. Adopting a turn off voltage below zero results in more decisive turn off and better rejection of gate charging, through the Miller capacitance, as the output voltage changes.

The schematic in Figure 60 shows how the gate driver has been realised. The gate drive command is transferred across the isolation barrier by an opto-isolator, which introduces a delay of 0.1 μs. This device is selected for its speed and high common mode dV/dt immunity of 15,000 V/μs. The resulting floating gate command is converted to a gate drive signal by a pair of MOSFET drivers, MAX4420 and MAX4429. The use of a complimentary pair of MOSFET drivers to provide a differential gate drive allows the gate to be driven to -15 V with respect to the emitter. Note that the 0V of the drive

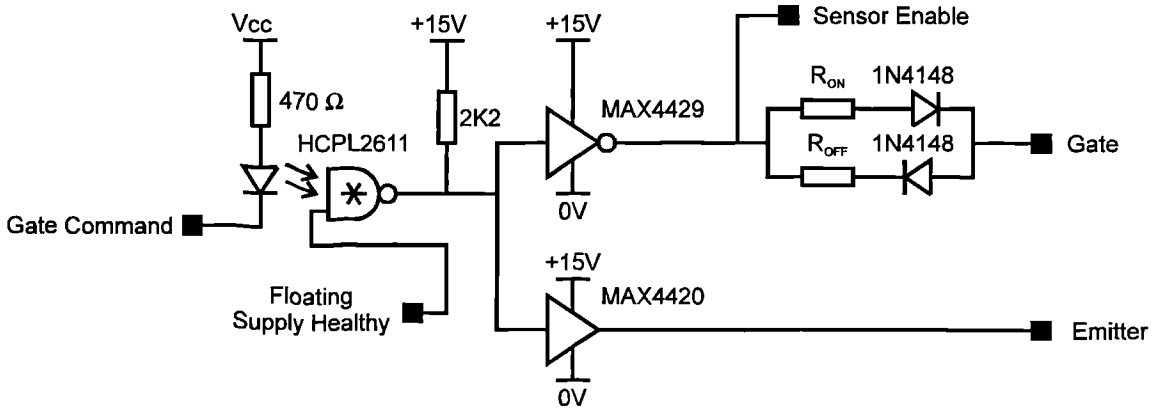

Figure 60 : Gate Driver Circuit Schematic

circuit is no longer tied to the emitter of the IGBT it drives. Instead when the IGBT is driven into the off-state the 0V of the driver circuit is driven to 15V below the IGBT emitter.

The author is not aware of any existing gate driver circuits using this technique of driving both the gate and emitter voltage. All the gate drive modules discussed earlier are referenced to the IGBT emitter and only drive the gate voltage. As a consequence these circuits require both positive and negative isolated voltage supplies whereas the circuit presented here requires only a positive isolated voltage supply.

The on-state voltage monitor differs little from existing designs. The differences arise because it must work in conjunction with a gate drive circuit which is not referenced to the IGBT emitter. The on-state voltage monitor, shown as a schematic in Figure 61, is based on a comparator which compares the IGBT collector-emitter voltage drop to a

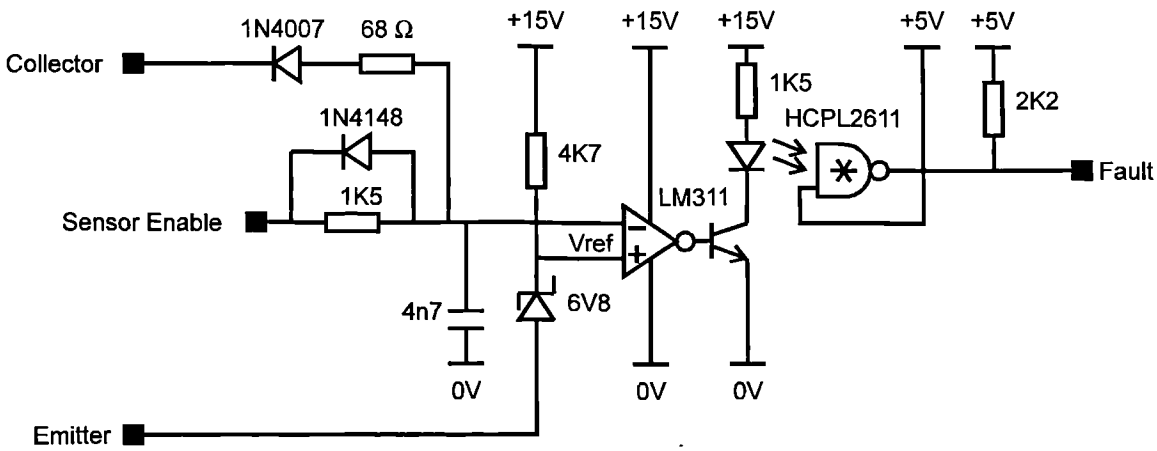

Figure 61 : On-State Voltage Monitor Circuit Schematic

reference voltage. The circuit is configured to detect a fault when the IGBT is turned on and V_{CE} exceeds the 6.8 V reference level. The fault indication is transferred back across the isolation barrier by an opto-isolator, again selected for its speed and high common mode dV/dt rejection.

The power required by the gate drive and on-state voltage monitor is transferred across the isolation barrier by a simple DC/DC converter which provides a stable +15 V floating supply. This supply drives all functions directly except for the opto-isolator whose 5 V supply is derived by a linear regulator. Under-voltage monitoring is provided to ensure that the IGBT is not turned on with a gate voltage of less than 11.6 V.

6.4.3 On-State Voltage Monitoring during Normal Operation

Before examining the operation of on-state voltage sensing under fault conditions it is important to understand how the power circuit and sensor behave under normal operating conditions. Figure 62 shows the circuit locations at which voltage and current measurements are taken. Note that whenever IGBT collector current, I_C, is measured, the current has actually been measured at the point marked i_{dc} in the figure.

Looking first at the IGBT switching characteristics whilst driving current into a phase with the machine operating at 4000 rpm, Figure 63(a) shows the IGBT switching characteristics at turn on. To allow for the turn off of short circuit current the gate drive resistances, R_{ON} and R_{OFF} have been set to 33 Ω and 47 Ω respectively. The gate is charged through the turn on resistor until it reaches the threshold voltage for the IGBT, approximately 9V. At this point the voltage across the device falls rapidly, reverse biasing the previously conducting anti-parallel diode. As the diode recovers its blocking

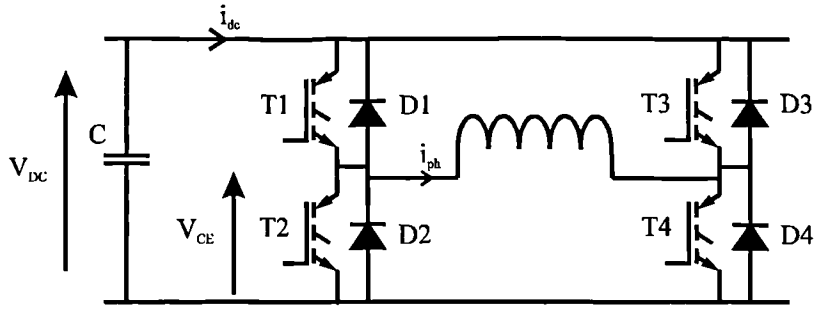

Figure 62 : Single Phase Bridge Showing Voltage and Current Measurement Points

ORIGINAL IN COLOUR

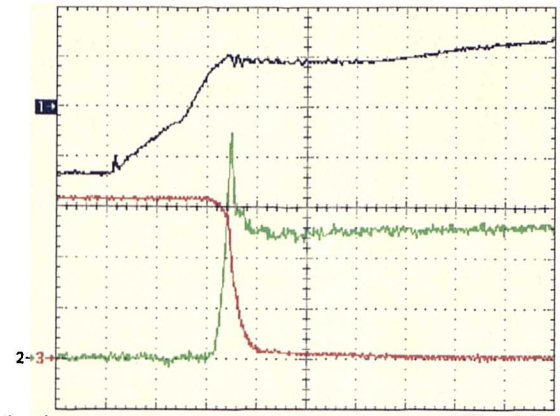

(a) Normal Operation - Forward Lower Device Turn On

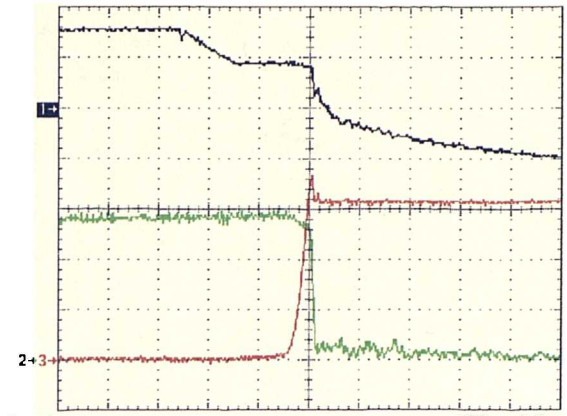

(b) Normal Operation - Forward Lower Device Turn Off

Figure 63 : IGBT Switching Behaviour under Normal Operating Conditions

capability the current in the IGBT rises. Note the current overshoot, due to a combination of the diode reverse recovery and the DC link power supplied into the motor phase capacitance. Also note the just visible tail voltage across the device after it has started conducting but before the gate voltage has reached the final value. The turn on delay is 1.1 μs and the fall time is 0.3 μs (90% to 10%).

Figure 63(b) shows the IGBT switching characteristics at turn off. The gate voltage is driven down to the threshold level at which point the IGBT starts to come out of saturation. V_{CE} rises to a level just above the DC link voltage, at which point the series connected diode is forward biased and starts conducting. Inductance in the connections from the DC link (i.e. from the reservoir capacitance) to the IGBT cause the falling current to generate a voltage overshoot. The turn off delay is 1.1 μs, the rise time is 0.2 μs (10% to 90%) and the voltage overshoot is 60V.

The behaviour of the on-state voltage monitor whilst the drive is functioning normally is shown in Figure 64. Figure 64(a) shows the operation of the sensor over a PWM cycle with 25 A flowing in the phase. The reference threshold falls from 15 V to 6.8 V at the instant that the gate drive is applied and returns to 15 V at the instant that negative gate drive is applied. The V_{CE} monitor signal is forced to zero when negative gate drive is applied and the slew rate of the V_{CE} monitor signal is limited to approximately 2 V/μs

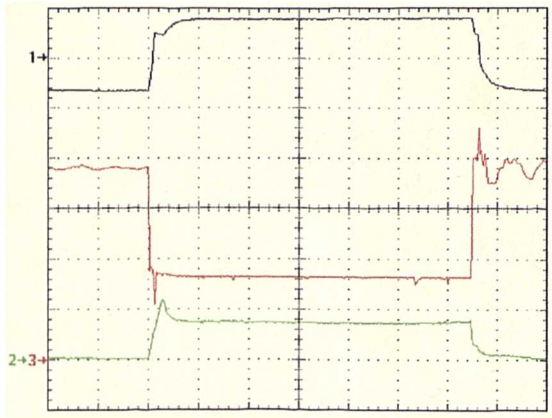

(a) Comparator Inputs over a PWM Cycle for a Phase Current of 20 A

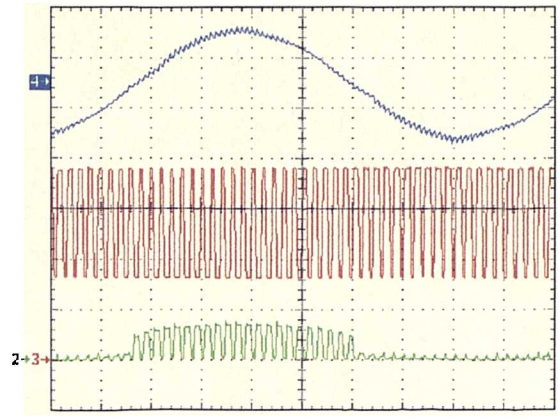

(b) Comparator Inputs over a Full Electrical Cycle

Figure 64 : On-State Voltage Sensor Behaviour under Normal Operating Conditions

when positive gate drive is applied. These measures prevent a fault indication from being generated when the IGBT is off or before the IGBT turn on is complete. From the figure it is apparent that the V_{CE} monitor signal rises at approximately 2 V/µs until equal to the voltage across the device. After this point the V_{CE} monitor signal voltage tracks the voltage across the device, which falls a little as turn on completes. Finally, a comparator generates a fault signal if the V_{CE} monitor signal rises above the reference level.

The noise immunity of the sensor has proved satisfactory, whilst operating the drive in a healthy state at rated current, with no instances of nuisance tripping. However, it is clear that the reference signal includes a noise component. Furthermore it is not actually necessary to force the reference voltage up to 15 V when the IGBT is turned off as the V_{CE} monitor signal is forced to 0 V at these times. Therefore a fixed 6.8 V reference would be more appropriate. This could be achieved simply by connecting the zener diode anode to 0 V rather than to the IGBT emitter.

6.4.4 Fault Coverage

A detailed study of fault detection using on-state voltage monitoring is provided in Chapter 8. However, it should already be clear that this sensor is capable of detecting load short circuits and shoot-through faults.

6.5 Transfer Current Sensor

6.5.1 The Principle

This scheme, developed by the author, recognises that the H-bridge can be split into two unipolar bridges, one for forward current and one for reverse current, Figure 65. In normal operation these two bridges function independently. When positive current is flowing in the load it is delivered from or to the DC link by the positive bridge or the current is freewheeling in the positive bridge. In any case, no current will flow in the negative bridge. Similarly if the load current is negative then this current will be carried solely by the negative bridge. Clearly the current transfers from one bridge to the other as the phase current crosses zero. Using the currents defined in Figure 65, there are only two possibilities for healthy operation:

Positive load current: $i_1 = i_2 \geq 0$ and $i_3 = i_4 = 0$ (24)

Negative load current: $i_1 = i_2 = 0$ and $i_3 = i_4 \leq 0$ (25)

Any deviation from these relationships is an indication that a fault has occurred. If an additional current sensor is added, the current $(i_1 - i_2)$ can be measured. This current will always be zero during normal operation. This current sensor detects any difference between the current flowing from the DC link and the current returning, thus detecting any current flowing to ground. This additional sensor acts exactly like a residual current device (RCD).

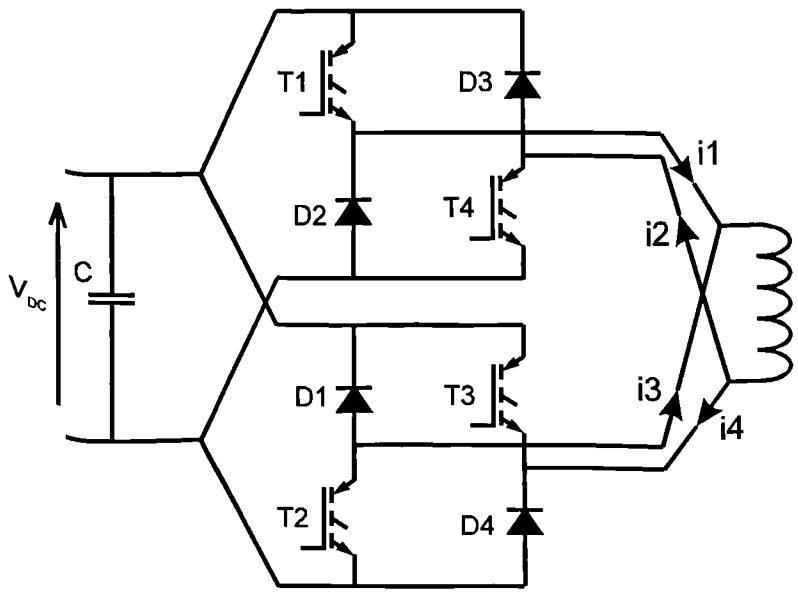

Figure 65 : Transfer Current Sensing

A further additional current sensor could allow measurement of (i_3 - i_4) which should also be zero. However, by arranging the four conductors with the correct polarities a single sensor can be used to measure (i_1 - i_2 - i_3 + i_4) which again should be zero. This single additional current sensor senses current delivered by one bridge but returning through the other. This current must transfer from one bridge to the other and hence gives this sensor its name - the transfer current sensor.

The suggestion here is that the sensor be placed between the inverter and the machine phase. Another possible location would be between the DC link capacitor and the inverter.

6.5.2 Fault Coverage

This sensor was originally conceived for its ability to detect shoot-through current resulting from a short circuit IGBT. However the sensor actually does much more than simply detecting current transfer from one unipolar bridge to the other. Any deviation of the sensor output from zero will indicate that current supplied to the load is failing to return by the appropriate path. Thus this sensor is capable of detecting phase to ground faults and phase to phase faults.

An IGBT short circuit fault will result in shoot-through when the healthy device in the faulty limb turns on. The fault current will be supplied by one bridge but return through the other. Thus the sensor can detect an IGBT short circuit - however a diode short circuit will remain undetected as the resulting shoot-through current in this case does not transfer from one bridge to the other.

Of all the sensors considered, this is the only one capable of detecting diode open circuit faults. Recall, from Chapter 5, that an open circuit failure in diode, D3, will result in reverse breakdown of IGBT, T3. When T3 breaks down, current delivered to the load by the forward bridge will return via the reverse bridge. Whilst this transfer current will be detected by the transfer current sensor, it is doubtful whether detection will be fast enough to save the parallel IGBT from fatal reverse breakdown.

6.6 Freewheel Current Sensor

The sensor presented here is taken from a fault detection procedure for single phase bridge converters, proposed by Craig, [50]. Craig discusses the principle behind the

sensor and the theory of operation at length. The work here will focus on how this sensor can be implemented and on the application of the resulting scheme to sensing faults in a bridge used to control the current in one phase of the demonstrator drive.

6.6.1 The Principle

The term 'freewheel current sensor' is used to refer to an additional current sensor placed so as to measure only freewheeling and shoot-through currents in the H-bridge as shown in Figure 66. By correlating the measurement from this fault detection sensor with the phase current sensor output and the switching state of the H-bridge, the health of the bridge may be determined. Whilst this technique requires one additional sensor per phase, Craig claims that it can be used to identify the power device which has failed and the failure mode. If this is true then this simple additional sensor could prove to be invaluable in fault detection and identification.

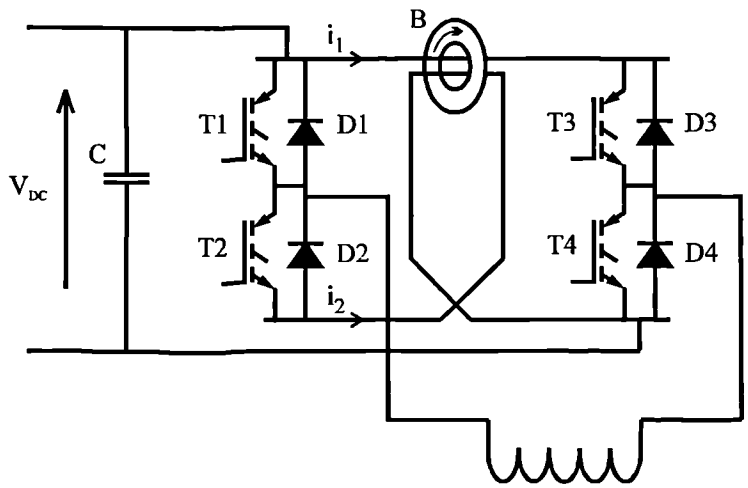

Figure 66 : Freewheel Current Sensor Location

The polarities of the current measured by the freewheel current sensor can be confusing. Thus the sign convention used by Craig is adopted here. The freewheel current sensor is arranged to provide a positive output when positive phase current is freewheeling in the upper path. This implies that positive phase current freewheeling in the lower path will result in a negative output from the freewheel current sensor. Note that the DC link physical connection must be made to diametrically opposed corners of the bridge. The choice of which pair of corners is immaterial and does not affect the sign of the freewheel or shoot-through currents seen by the freewheel current sensor.

6.6.2 Circuit Realisation

A circuit design has been prepared for a single phase bridge with an embedded freewheel current sensor. Layout and assembly drawings for a power hybrid design were produced from the circuit design, with assistance and support from Welwyn. A number of power hybrids were manufactured to provide a vehicle for testing the freewheel current sensor. Figure 67 shows a photograph of one such hybrid. A close up of the sensor arrangement is shown in Figure 68.

The use of a hybrid power stage provides two distinct advantages over alternative power electronics packaging techniques. The sensor can be integrated into the power stage without compromising the power circuit layout. The power circuit layout can be arranged with sufficient separation between components to minimise the risk of fault propagation.

The design uses 2 parallel IGBTs and a fast soft recovery diode to form each power switch. The IGBTs are Siemens die and the diodes are supplied by IXYS. Gate resistors and gate clamping zener diodes are included in the hybrid but the gate drive circuitry is not. All these parts are mounted on a directly bonded copper (DBC) substrate. There is some separation between each power switch to provide thermal and physical isolation between switches.

The freewheel current sensor is implemented as follows: the 2 power tracks which carry the freewheeling currents pass through a ferrite ring clearly visible in the photographs. One of the freewheeling paths can be seen tracked in from the right of the upper right switch in Figure 67, through the flux channel and back out to the lower right switch. The other freewheeling path is on the reverse side of the substrate so the only visible sign is the pair of via holes on the left of the substrate in Figure 67.

The flux in the ferrite ring is proportional to the sum of the currents in the tracks which pass through it. A Hall effect sensor is inserted into a gap in the ferrite ring so as to measure the flux. The Hall effect sensor is mounted on a separate thick film substrate along with the amplifier circuit shown in Figure 69, to provide an output of 100 mV / amp. The thick film substrate on which the amplifier is constructed dominates the close up view in Figure 68. The Hall effect sensor is difficult to see as its small black body is buried in the slot in the black ferrite ring. However the 4 silver legs of the Hall effect

sensor are visible, projecting 'down' from the ferrite ring, across the surface of the thick film hybrid substrate.

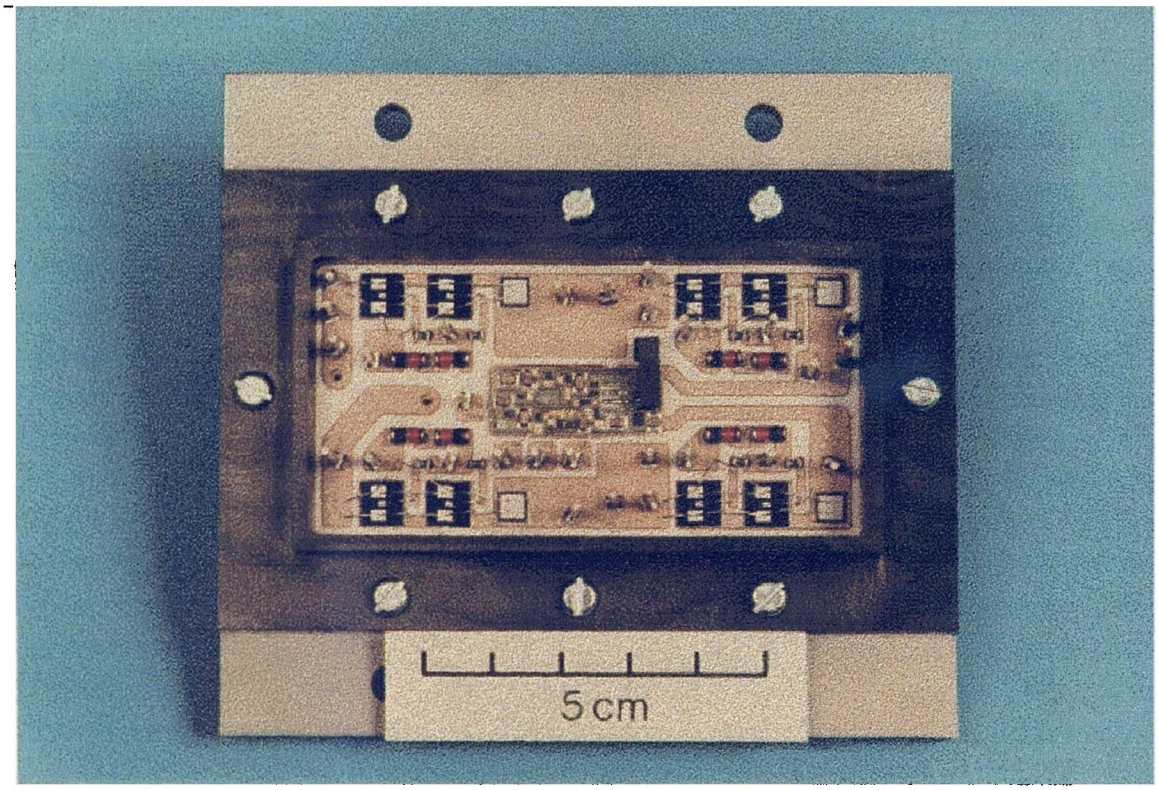

Figure 67 : Power Hybrid Incorporating Freewheel Current Sensor

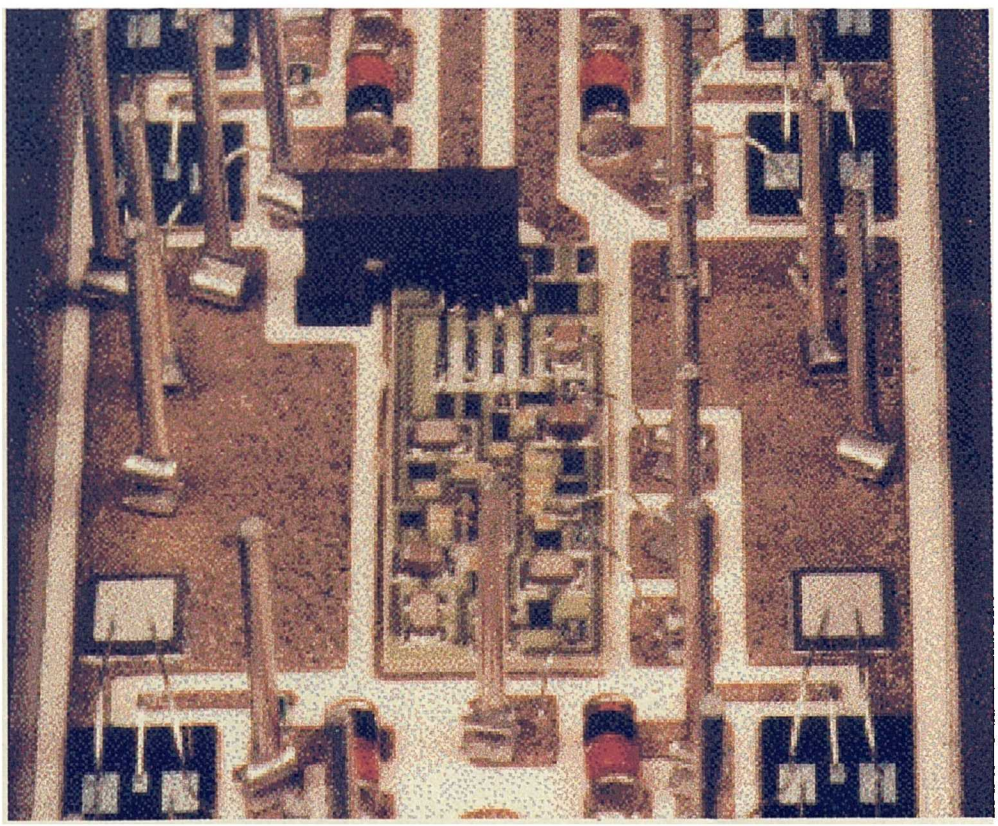

Figure 68 : Close Up on the Freewheel Current Sensor

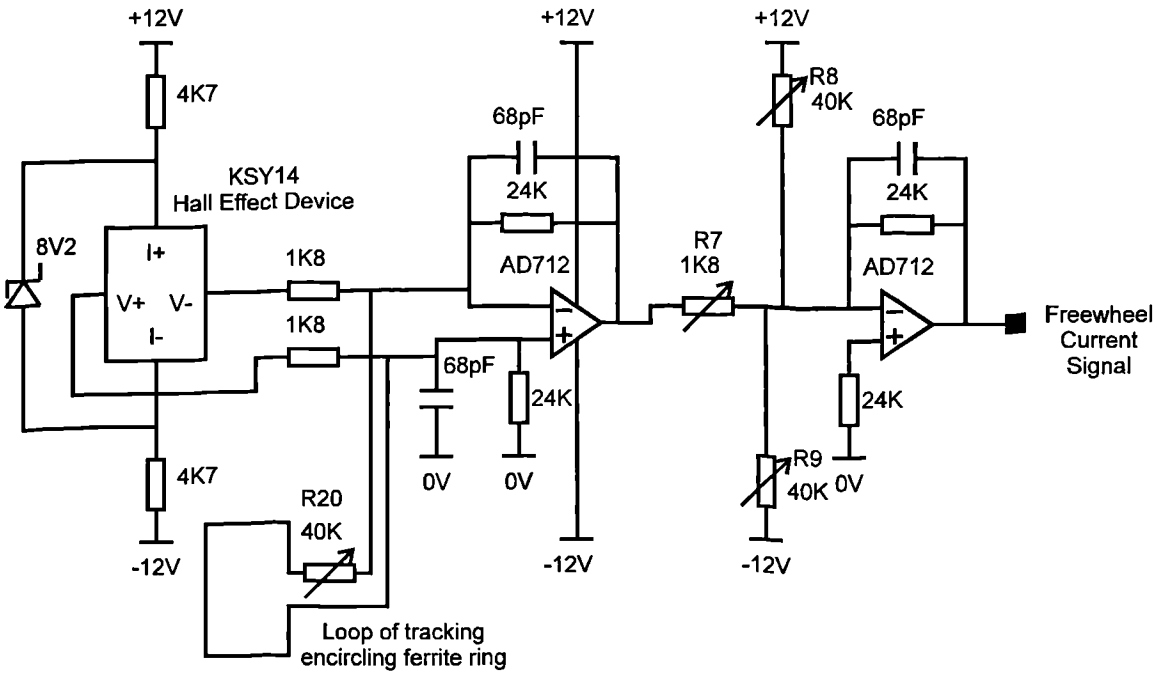

Figure 69 : Freewheel Current Sensor Circuit Schematic

Experiments on a prototype current sensor circuit using the KSY14 Hall effect device revealed that the output voltage of this device included two components. The first of these is a voltage generated by the Hall effect and is proportional to the flux density, B, in the channel. The second component was produced by a parasitic search coil in the device, formed by the device leads. The second component was proportional to the rate of change of flux, dB/dt. The freewheel current sensor is required to measure rapidly changing currents, and the resulting dB/dt meant that the component of voltage generated by the parasitic search coil swamped the component of voltage proportional to the current.

An effective and perhaps novel solution to the parasitic search coil problem was to add a second search coil to the circuit. The voltage from this second coil was used to cancel the voltage from the parasitic search coil as shown in the schematic, Figure 69. The second search coil is formed by a track on the thick film substrate which encircles the flux channel. The go and return conductors of the encircling track are just visible along the left edge of the thick film substrate in Figure 68.

6.6.3 Power Hybrid Performance

A power stage using one of these hybrids has been built and used to drive rated current into one phase of the machine. Short circuit tests have been performed on the hybrid.

The stray inductances within the hybrid are low enough to allow rapid turn off of the short circuit current. Tests indicate that most of the stray inductance results from bond wires. The tracking and ferrite ring associated with the freewheel current sensor only increases the stray inductance by approximately 20%. The short circuit current was measured to be approximately 250 amps.

The output of the freewheel current sensor has been examined. This output is an acceptable indication of the steady state current in the freewheeling paths. The currents in the two paths add in the correct proportions to ensure that the sensor measures only freewheeling and shoot-through currents.

6.6.4 Freewheel Current Sensing during Normal Operation

The response of the phase current, the current delivered to the phase from the DC link and the freewheel current sensor output during normal operation of the drive are shown in Figure 70. Note that the freewheel current transducer output is subject to some disturbance, usually manifested by an overshoot, after each transition. This disturbance is a result of capacitive coupling between the power circuit and the sensor. The size of the overshoot corresponds to a current indication of approximately 5 A, i.e. 17% of peak rated current.

6.6.5 Fault Coverage

The details of how faults can be detected and identified will be discussed at length in the following chapters. At this stage it is helpful to understand that the faults which may be detected by this sensor fall into 2 groups, those which result in shoot-through current, and those which result in a difference between the phase current and the freewheeling current during the freewheeling periods in the PWM switching cycle.

This sensor is obviously capable of detecting power device short circuits as these faults result in shoot-through current. Craig claims that the sensor is also capable of detecting power device open circuit faults as these faults result in a difference between the phase current and the freewheeling current during the freewheeling periods in the PWM switching cycle. The sensor is not able to detect short circuits at the machine terminals, nor is it able to detect an open circuit winding.

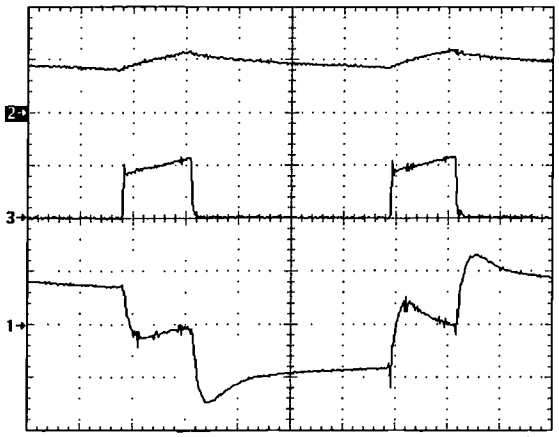

(a) Freewheel Current Signal for Positive Phase Current over a PWM Cycle

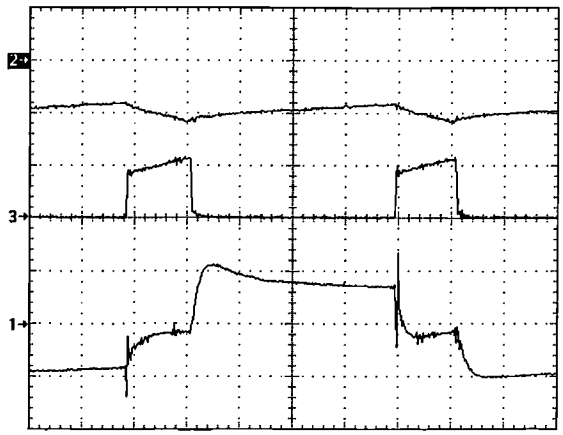

(b) Freewheel Current Signal for Negative Phase Current over a PWM Cycle

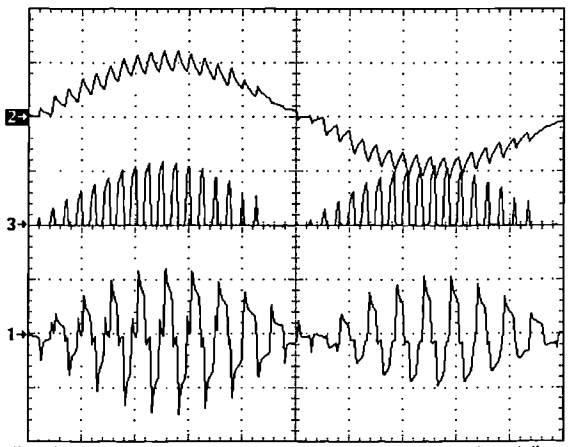

(c) Freewheel Current Signal over a Full Electrical Cycle

Figure 70 : Freewheel Current Sensor Output during Normal Operation

6.7 Sensor Combinations

When selecting a set of sensors, the first requirement is that, in combination, they must be capable of detecting all of the faults considered. The fault coverage of each of the sensors discussed in this chapter is presented in Table 14. The rationale behind some of the table entries will not be completely clear until after fault detection schemes based on the sensors are developed in later chapters. However the comparison must be made at this stage to allow those sensors most worthy of further attention to be selected.

From the table it is clear that, with the exception of the power diode open circuit, all possible faults may be detected using a combination of on-state voltage sensing and phase current sensing. The last line of the table indicates whether or not a detection scheme based on each sensor would be capable of detecting the failure of the sensor itself during normal operation. It appears that fault detection schemes should be developed based on the phase current sensor and the on-state voltage sensor as a minimum. The availability of power hybrids with an integrated freewheel current sensor make this an attractive sensor for further investigation. Thus fault detection schemes based on these 3 sensors will be developed in the following chapters.

	Phase current sensor	On-state voltage sensor	Transfer current sensor	Freewheel current sensor
Phase open circuit	√			
Phase short circuit (at machine terminals)	√	√		
Phase connected to frame	√	√ (1)	√	
Phase connected to another phase	√	√ (2)		
Turn to turn fault	√			
Power switch open circuit	√			√
Power switch short circuit		√	√	√
Power diode open circuit			√ (3)	
Power diode short circuit		√		√
Gate drive stuck off	√			
Gate drive open circuit	√			
Gate drive stuck on		√	√	√
Fault detection sensor failure	√	√ (4)		√ (4)

1. Low impedance faults only, provided neither stator frame nor DC link are isolated from ground.
2. Low impedance faults only.
3. This fault would be detected but possibly not fast enough to prevent fault propagation.
4. The failure of this fault detection sensor will only be detected if careful consideration is given to the fault detection scheme (see Chapter 8).

Table 14 : Sensor Combinations

6.8 Conclusion

A number of sensors which could be used for fault detection have been considered. Each sensor is capable of detecting a variety of faults. A combination of sensors will permit the detection of all faults. As more sensors are added, the level of redundancy in the sensing scheme increases which may be advantageous. However the integration of the information also becomes more difficult as the sensor count increases. At this stage it is possible to see that the sensors presented are capable of detecting the target faults in appropriate time scales.

Use of information from the phase current sensor, for fault detection, is extremely attractive, as it offers the possibility of fault detection without requiring any additional sensors. Moreover, any detection scheme based on the phase current sensor will naturally detect the failure of the phase current sensor.

In general, fault detection sensors whose outputs are exercised during normal operation are preferred, as these types of sensor are naturally subject to continuous self checking. The on-state voltage sensor may or may not fall into this category depending on the design. This sensor is electrically connected to the power switch it monitors. The lack of isolation is not a serious concern as power switch failures are detected by the on-state voltage sensor connected to the remaining healthy device in the faulty leg. On-state voltage sensing is already used in many gate drive circuits to protect power devices from damage due to excessive current, so it can be regarded as an existing, rather than additional sensor.

In appears that the transfer current sensor is capable of detecting many of the faults which could be detected with either the on-state voltage sensor or with the freewheel current sensor. It appears that the transfer current sensor is the only sensor capable of detecting high impedance phase to phase and phase to ground faults and possibly of detecting open circuit diode faults. However the sensor output only changes state when a fault occurs which means that sensor failures may go undetected and this represents a serious drawback.

The freewheel current sensor appears to offer the possibility of detecting most bridge faults with only one additional sensor. It is galvanically isolated from the bridge so is

much less prone to damage in the event of a power device failure. This sensor is naturally self checking as the current in the bridge alternates between the freewheeling paths during normal operation.

In the next two chapters fault detection schemes will be developed, based on the phase current, on-state voltage and freewheel current sensors.

Chapter 7

FAULT DETECTION - Part I

7.1 Introduction

Earlier chapters have established that, if a drive is only able to survive a single failure, it is essential to replace a faulty drive before a second failure occurs. The inability to survive more than a single point failure means that fault propagation must be prevented. Furthermore, certain faults are likely to result in more than 10 times rated current flowing in the machine. One of the questions posed here is 'Is it possible to detect these faults fast enough and can appropriate action be taken to allow continued operation of the drive in the presence of these faults without significant disturbance to the supply or to the load?'

The requirement to replace a faulty unit before a further failure occurs means that all possible faults must be detected, so that a faulted drive can be identified and reported. This requirement has been alluded to by Nicolaidis et al [98] who point out that the design of strongly fail safe digital systems is less intractable if the system is also self checking. The requirement to prevent fault propagation will require the very rapid detection of the most severe faults.

If a failed unit is to be repaired by component replacement then it is desirable to know the exact nature of the failure. In practice small drives are usually repaired by replacing the entire drive. Thus in a small fault tolerant drive it is only important to know if a fault has occurred and what the effects of the fault are, so that appropriate action can be taken to enable continued operation of the faulty drive. It is much less important to know exactly which component has failed. If several different component failures each result in the same symptoms and demand the same post fault operating strategy, then it is not important to distinguish between these failures.

In general, fault detection schemes can be realised in hardware or software. Fault detection schemes implemented in software will naturally operate in discrete time, using sampled data. It is convenient to run these detection schemes at the same sampling rate as the current loop. In modern drives the current loop update period is typically of the order of 100 μs. In contrast, schemes implemented in hardware can easily operate directly in response to continuously monitored values. In principle these hardware based schemes could detect faults instantly, however in practice delays in the signal path will typically be of the order of 1 μs. This chapter will consider fault detection schemes implemented in S/W. Fast operating, hardware based fault detection will be covered in Chapter 8.

In the course of this chapter, shorted turn faults and power device open circuit faults will be imposed on the drive. The techniques used to introduce these faults were presented in Chapter 5.

7.2 Objectives

This chapter will consider fault detection schemes which operate on data sampled once every current loop control cycle. Two fault detection schemes will be considered, both based on information from the phase current sensors. The objective of the fault detection schemes developed will be to provide complete fault coverage, i.e. the fault detection scheme will attempt to detect all the faults discussed in Chapter 5. To be of use, the faults must be detected before fault propagation occurs. One area that was not discussed in the section on sensors was the possibility of developing fault confirmation and fault diagnosis techniques, so these will now be examined here.

7.3 Fault Detection by 'Phase to Model' Comparison

7.3.1 Theory of Operation

This fault detection scheme is dependent on the current control technique used. In Chapter 3 a model based current control scheme was described. This current control technique requires that the machine be characterised in terms of measured flux linkage with respect to current and rotor angle. The ψ, i, θ characteristic is held in the controller as a lookup table. The current controller uses the sampled quantities in conjunction with the PWM demand for the present PWM interval and the ψ, i, θ data to calculate the

required PWM demand. This is applied to the inverter to drive the current to the exact value desired at the end of the following PWM interval. The calculation is repeated for each phase, every PWM cycle.

From the ψ, i, θ lookup table the current controller knows what current it expects to achieve at the end of each PWM cycle. Any difference between the current subsequently measured and the current expected is an indication that the machine has changed, i.e. the machine has developed a fault.

Open circuit power devices and open circuit phases both result in zero phase current during part or all of the machine cycle. Does this mean that phase current sensing can detect these faults? Turn to turn faults result in a more subtle change in the phase current. Is this change sufficient to be detected? Can phase current sensing be used to detect failure of the phase current sensor itself? The answers to these questions are prerequisites to assessing the usefulness of phase current sensing. Each of the questions will be addressed in the following sections, beginning with those related to shorted turn detection.

7.3.2 Shorted Turn Detection

Some work has been directed towards developing techniques for detecting shorted turns. Most of these techniques may be classified as flux based, [83] and [99], or current based [100]. A variety of off line diagnostic tests for shorted turns in a super conducting magnet have also been described, [101]. It appears that none of the work provides a simple detection technique which would enable a PWM inverter drive to self monitor without the need for additional sensors.

This section will examine whether it is possible to detect a single shorted turn without greatly changing the control scheme described in Chapter 3. The process of insulation breakdown is a complex one, beginning with partial discharge and very localised heating, then progressing to a low resistance short circuit capable of carrying considerable current. This discussion is centred on detecting the fault as soon as significant fault current begins to flow.

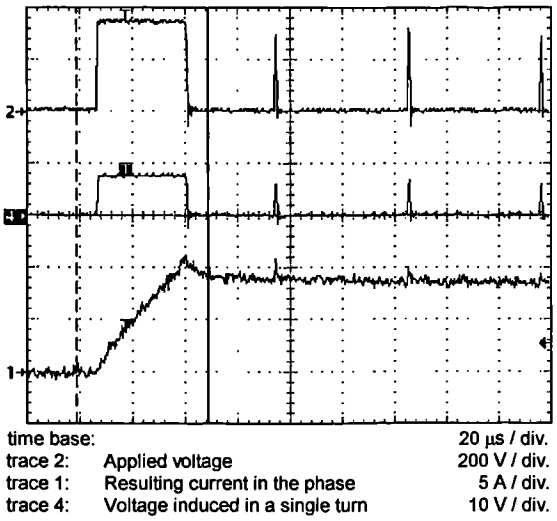

(a) A Healthy Phase

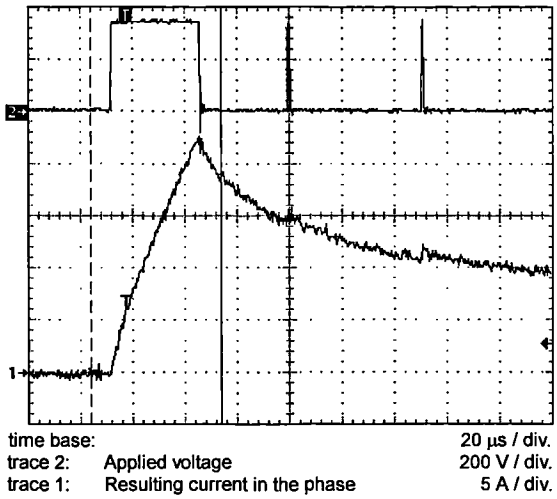

(b) A Phase with a Shorted Turn

Figure 71. The Effect of a Single Shorted Turn in the Demonstrator Machine

Effect of a Shorted Turn on Current Control in the Demonstrator Machine

Chapter 5 hinted at how shorted turn detection might be possible in this machine. The key is that the presence of a shorted turn reduces the inductance seen looking into the phase. At, and below, the machine operating frequency the current in the shorted turn is largely limited by its resistance and it has only a small effect on the fundamental phase flux and therefore little impact on the fundamental phase current or voltage. However at the PWM frequency this is no longer true. At this frequency the single turn reactance is 64.1 mΩ, whilst the resistance, neglecting skin effect, remains at 3.12 mΩ. Thus the shorted turn does have a large effect at the PWM frequency.

Compare Figure 71(a) and Figure 71(b) which show the effect of a shorted turn over about three PWM cycles. Each oscillogram shows voltage and current waveforms measured for a particular phase of the machine. In these tests the machine is at rest so no back-EMF is present. The plot in Figure 71(a) shows operation for a healthy phase. The waveforms correspond to a demand for a step in current, from 0 A to 9 A. Trace 2 shows the terminal voltage applied to the machine. This voltage is calculated by the controller to drive the phase current from 0 A to 9 A and to subsequently hold the current at 9 A. Trace 1 shows the resulting phase current. Trace 4 shows the voltage appearing across the individually monitored single turn. This voltage would drive current in that turn if it were shorted.

A cursor bar has been placed on the oscillogram at the start of each of two consecutive PWM cycles. These cursor bars appear as vertical lines, immediately before and after the current step. The controller samples the phase currents at these instants and measures the change in current to be 9.0A. This matches the current change that the controller expected to achieve so all is well.

The plot in Figure 71(b) shows waveforms with a short circuit applied across the individually monitored single turn, i.e. between points X and Y in Figure 47 of Chapter5. Trace 2 shows the terminal voltage applied to the machine and trace 1 shows the resulting phase current waveform. The terminal voltage applied to the machine is identical to that applied in Figure 71(a) and again the two cursor bars have been placed on the plot at the current sampling instants immediately before and after the current step. Now the controller measures the change in current to be 17.0 A, which is 89% higher than the anticipated change in current. This is a sufficient and clearly detectable indication of a turn to turn fault. The percentage by which the measured change in current exceeded the expected change in current is given the symbol 'F'.

It is possible to make an estimate of the current in the single shorted turn by subtracting the phase current in the healthy phase (Figure 71(a)) from the phase current with a shorted turn (Figure 71(b)) and multiplying the result by the turns ratio, 50/1. This suggests a peak current of 575 A in the shorted turn and a time constant of 35μs for the rise and decay of this current. This calculation is valid because the same volt seconds have been applied to the healthy and the faulted phase and therefore the same total flux is present in both cases. Thus the MMF resulting from the additional current in the faulted phase must be balanced by an equal but opposite MMF resulting from current in the shorted turn.

Effect of a Shorted Turn on Current Control in the Small Prototype Machine

Figure 72 shows the same comparison for the small prototype machine. Again the 2 cursor bars have been placed on the plot at the sampling instants immediately before and after the current step. In Figure 72(a) the controller samples the phase currents at these instants and measures the change in current as 9.4 A. This matches the current change that the controller expected to achieve so all is well.

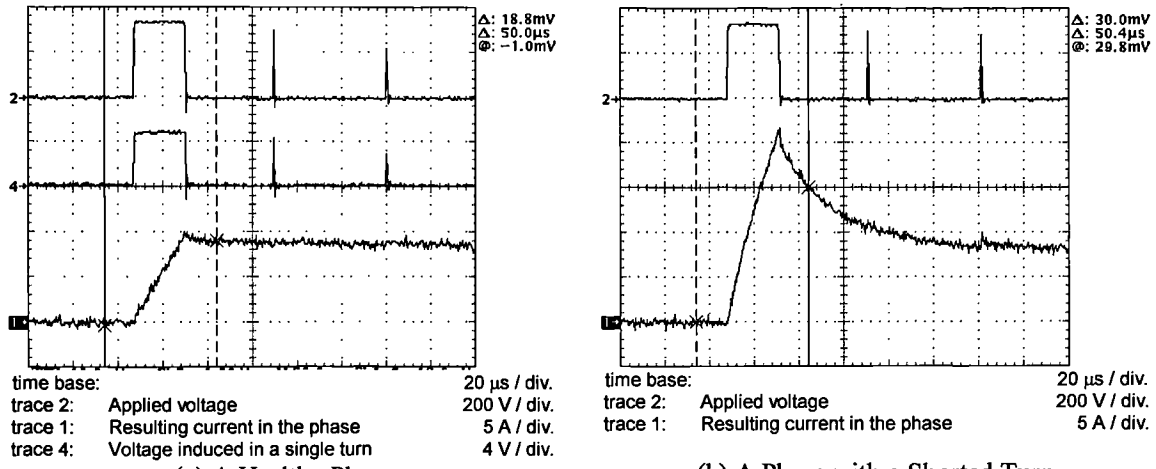

(a) A Healthy Phase (b) A Phase with a Shorted Turn

Figure 72 : The Effect of a Single Shorted Turn in the Small Prototype Machine

The plot in Figure 72(b) shows waveforms with a short circuit applied to the terminals of the individually monitored turn. The terminal voltage applied to the machine is identical to that applied in Figure 72(a) but now the controller measures the change in current to be 15.0 A, 60% higher than the expected change. Clearly phase current sensing is also capable of detecting a single shorted turn in this machine, even though there are more turns per phase than in the demonstrator machine.

7.3.3 Factors Affecting Detectability

The timing of sampling instants within the PWM cycle has important consequences for the sensitivity of detection. The sampling instants shown in Figure 71 and Figure 72 are at the start of each PWM period, when the current passes through its average value for the PWM cycle. Selection of these points means that the current controller does not see the PWM ripple, therefore it controls the average value of the current. It would be better for shorted turn detection to use a measurement of the ripple in the phase current over each PWM cycle. This current ripple is present, even for a current demand of zero, provided that the machine is operating under PWM control. Even using the sub-optimal sampling instants, the shorted turn is clearly detectable for a modest current demand, using the following criterion for shorted turn detection:

$$| i_{Measured} | > | i^* | + | i_{MeasError} | + | i_{CtrlError} | \quad (26)$$

There is a minimum change in current demand between sampling instants which will make the shorted turn detectable. The current overshoot must exceed the sum of the

maximum measurement error, $i_{MeasError}$, and the maximum control error, $i_{CtrlError}$, before it becomes an unambiguous indication of a shorted turn.

An estimate of the maximum current control error can be made from the results in Chapter 3. The measurement error is small compared to this control error. If 4.0 A is taken as a conservative estimate for the sum of the errors then a current overshoot of greater than 4.0A is required to provide an unambiguous shorted turn indication. From Figure 71 the shorted turn results in a current overshoot, F, of 89%. A more conservative figure, F = 80%, is use to calculate the minimum step in current demand required to allow shorted turn detection:

$$\Delta i = 4 \text{ A} / 0.80 = 5 \text{ A} \quad (27)$$

If the winding fault results in a sufficient reduction in the impedance looking into the terminals then the current will overshoot by more than 100 % each PWM cycle. In this case a current oscillation of increasing magnitude will result from the smallest change in current demand. Although this situation would make shorted turn detection easier, it is not a requirement for detectability.

It is clear that larger changes in phase current demand from one sampling instant to the next will produce larger current overshoots and better detection. So shorted turns will be more detectable when the electrical frequency increases (faster speed) or when the magnitude of the current increases (higher load).

So far it has been assumed that the percentage overshoot is affected only by the characteristics of the machine. However, the current is supplied from a power converter comprising a set of diodes and IGBTs. The voltage drop across these components is not proportional to current. To a first approximation the voltage drop across a conducting IGBT or diode is constant, thus for low voltage demands a higher proportion of the voltage is dropped across the power devices. The result is that the power converter tends to reduce the percentage current overshoot for lower voltage demands, i.e. for more slowly changing currents. The effect of this non-linearity is illustrated in Figure 73. The lower pair of traces in Figure 73(a) show the sampled current superimposed on the sampled current demand for a phase with a shorted turn. Sustained oscillation of the current is clearly visible. The upper pair of traces show the sampled current

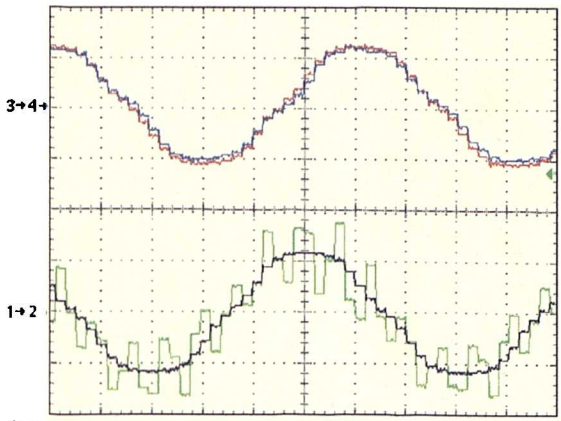

(a)

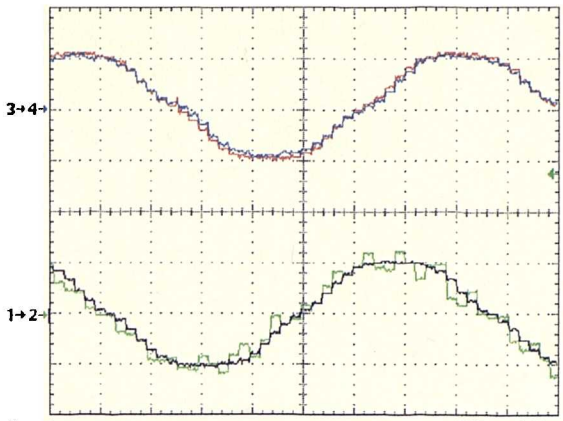

(b)

Figure 73 : The Effect of a Reduction in Current on Detectability

superimposed on the sampled current demand for an adjacent healthy phase. Good current control is maintained in the adjacent healthy phase. Figure 73(b) shows the effect of a small reduction in the current demand. Each of the traces represent the same quantities as in Figure 73(a). By comparing the two oscillograms, it is apparent that a slight decrease in current demand results in a transition from a large sustained current oscillation to a small current oscillation. This is particularly significant since the reduction in current demand has resulted in a reduction in the amplitude of oscillation to below the level of detectability using the detection criteria proposed.

The reduction in inductance seen looking into the terminals of a phase containing a shorted turn depends on the coupling between the shorted turn and the rest of the phase. The leakage inductance and resistance of the short are important in determining the time constant for the additional current. Other factors also play a part in the ease of detection, such as PWM strategy and PWM frequency.

For larger machines operating from the same supply voltage the number of turns per phase will tend to be lower, making the detection of shorted turns easier.

7.3.4 Adapting the Current Profile

In the last section the factors affecting detectability were examined. One of the key factors is the magnitude of current change over each PWM cycle. Figure 73 showed that

for a given current profile there is a critical level of current in the machine, below which a single shorted turn is no longer detectable using the detection criteria proposed.

The current profile chosen will affect the minimum level of peak current required for a shorted turn to be detected. A square wave current profile is best in this respect since it includes fast current transitions. However it would be a severe penalty if the detection scheme required a particular current profile to work effectively. Furthermore, even with a square wave current demand, there is still a minimum current demand below which the single shorted turn is not detectable.

One solution would be to accumulate error over many PWM cycles. Investigation of this possibility showed that it gave some improvement in sensitivity to the shorted turn, but still broke down at low speed and light load.

A solution to the problem of shorted turn detection at light load is found by using an adapted current waveform with current pulses inserted. By driving a short pulse of current into the phase, a shorted turn can be detected even when the machine is not rotating. The current pulse should be as short as possible and in this case the PWM frequency sets the pulse duration at 100 μs. The current pulse magnitude is set by the maximum control error, $|i_{CtrlError}|$ the maximum measurement error, $|i_{MeasError}|$ and the level of current overshoot caused by the shorted turn, F.

$$\left|\Delta\hat{i}_{Pulse}\right| = \frac{\left|i_{CtrlError}\right| + \left|i_{MeasError}\right|}{F} \tag{28}$$

For example taking some very conservative figures for the demonstrator drive, $|\, i_{CtrlError} \,| = 1.0$ A, $|\, i_{MeasError} \,| = 6.0$ A and F = 80 %, a current pulse with a peak amplitude of 8.75 A is required.

If a current pulse is added to the current demand once per machine cycle then a shorted turn will be detectable within one machine electrical cycle, even at no load. It is best to insert the current pulses when the back-EMF passes through zero in the positive and/or negative sense, as at these points they do not affect the machine torque. The polarity of the current pulses is chosen in the same sense as the natural direction of current change at these points so as not to hamper the current controller in its task of rapidly slewing the current at these positions.

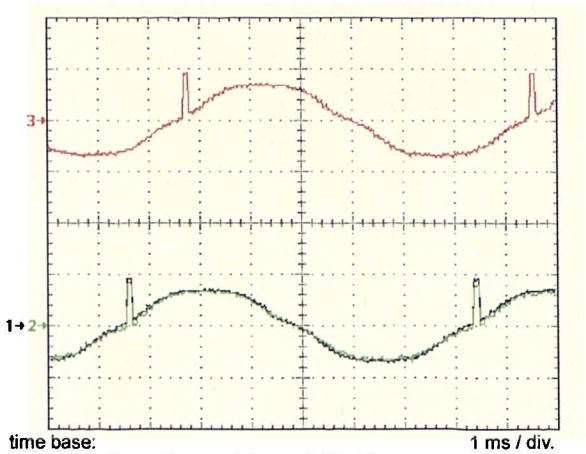

Figure 74 : Current Profile Adapted by the Addition of Current Pulses

Figure 74 shows operation of the machine with current pulses inserted at the unaligned positions. The lower pair of traces show the sampled current superimposed on the sampled current demand for phase A. The upper trace shows the sampled current demand for phase B. The algorithm used in the control software detects the arrival of the unaligned position. If the change in current over the PWM cycle corresponding to the arrival of the unaligned position is to be less than 8.75 A then the demanded change in current is increased to 8.75 A without changing the polarity of the current change. In the subsequent PWM cycle the current is driven back to the level demanded by the current profile.

Figure 75 shows the same current pulse injection scheme applied with a single shorted turn in phase A. In Figure 75(a) the traces represent the same quantities as those in Figure 74. The current pulse causes an initial detectable phase current overshoot and decaying oscillation. It is clear, from the absence of oscillation during the negative half cycle of phase current, that the shorted turn would be undetectable at this current level without current pulse insertion. Figure 75(b) shows exactly the same situation but in this oscillogram the traces represent continuously currents rather than sampled values. The lower two traces show the current in the faulted phase and the current in the shorted turn. The upper trace shows the current in an adjacent healthy phase, also showing an injected current pulse but with no subsequent oscillation.

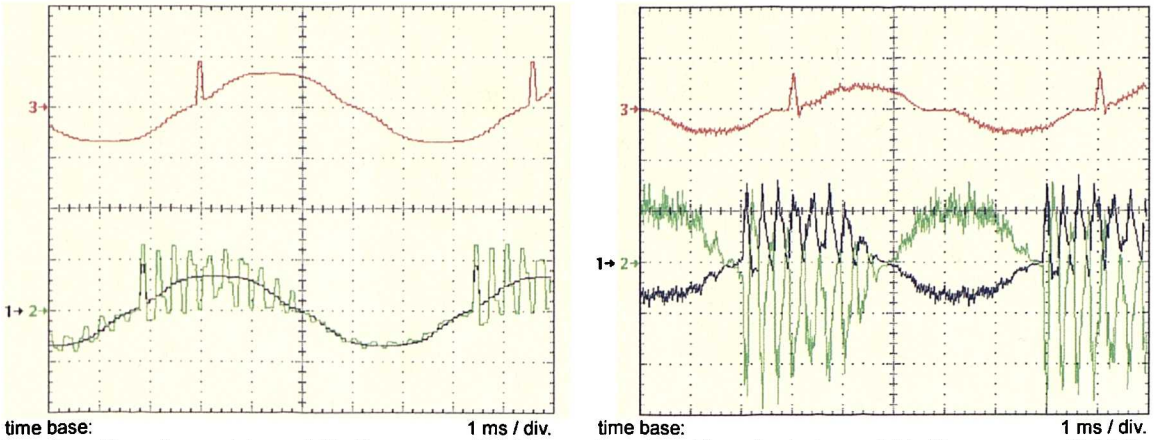

(a) Controller Sampled Currents (b) Measured Currents

Figure 75 : Effect of Current Pulse Insertion in Shorted Turn Detection

Inserting a current pulse of 8.75 A for one PWM cycle at the unaligned positions increases the rms current in the machine. The increase in rms current is at a maximum when the machine is at full speed but no load. Then the rms current rises from 0.0 A to 1.07 A, which is still only 5.1 % of the rated current of 21.1 A rms. In practice this operating point is an unrealistically harsh worst case for a fuel pump drive, which will never operate at full speed with no load.

7.3.5 Open Circuit IGBT and Open Circuit Phase Winding Detection

An open circuit IGBT will result in zero phase current during either positive or negative half cycles of current demand and may be detected using the criterion:

$$| i_{Measured} | < | i^* | - | i_{MeasError} | - | i_{CtrlError} | \quad (29)$$

The half cycle for which this criterion is met will indicate whether a forward device or a reverse device is affected.

An open circuit winding will result in zero phase current and may be detected using the same criterion. To avoid confusion with the open circuit power device fault the criterion must be satisfied for positive and negative current demands.

For the detection of any of these open circuit faults, the current demand must exceed the sum of the maximum measurement and control errors, i.e. the condition for detectability is:

$$|i^*| > |i_{CtrlError}| + |i_{MeasError}| \quad (30)$$

The current pulse insertion technique outlined in the previous section will ensure detectability of these faults. For the detection of all the open circuit faults, both positive and negative current pulses are required. This is simple to achieve by adding pulses at both the aligned and unaligned positions. Finally note that the requirements for the pulses for open circuit detection and for shorted turn detection differ in that a minimum peak current is required for open circuit detection, whereas a minimum change in current over the PWM cycle is required for shorted turn detection. Nevertheless the same pulses can fulfil both sets of requirements.

7.3.6 Power Device Short Circuit and Phase Terminal Short Circuit Detection

Power device short circuit faults and phase terminal short circuit faults will both result in large fault currents flowing when the inverter switches the DC link voltage onto the fault. The resulting fault current will only be seen by the phase current sensor in the case of a phase terminal short circuit, and even then only if the fault occurs on the machine side of the phase current sensor. Whilst the fault current resulting from certain short circuit faults will not flow through the phase current sensor, all short circuit faults will result in a change in the phase current. Thus it should be possible to detect power device short circuit faults and phase terminal short circuit faults by monitoring the phase current sensor output and comparing this to the expected value of phase current.

In the case of short circuit faults, very rapid fault detection is required so that fault propagation can be prevented. Using processors currently available, the current loop update rate is of the order of 100 μs. Given a typical IGBT short circuit withstand time of 10 μs, it will not be possible to use 'phase to model' comparison to detect power device short circuits in time to prevent fault propagation. This is true for all fault detection schemes which run at the current control loop update rate, i.e. none of the schemes described in this chapter will be capable of protecting power devices from short circuit faults.

7.3.7 Phase Current Sensor Failure Detection

Phase to model comparison will detect phase current sensor failures such as phase current sensor output stuck at zero, or sensor output open circuit. However unless a

fault diagnosis scheme is adopted the current sensor failure will be indistinguishable from other faults. This is not such a serious problem as might be imagined, as will be shown in Chapter 9.

7.4 Fault Detection by 'Phase to Phase' Comparison

The last section showed how it is possible to detect faults by comparing the current measured in a phase to the current predicted by a simple model. The principle of phase to phase comparison is similar at least to the extent that it also uses the phase current sensor as an information source.

In any machine with an even phase number, the opposite phases are equivalent. The machine presented here has 6 phases and the separation between adjacent phases is 60 electrical degrees. Thus the separation between opposite phases, such as A to D, B to E, and C to F, is 180 electrical degrees. When the drive is operating the current in opposite phases should be equal but opposite. This presents two possibilities:

1. A single control loop could control the current in a pair of opposite phases.
2. The current in a pair of opposite phases should always sum to zero. Any significant deviation from zero is an indication that one of the phases is faulty.

These two could be used in conjunction but need not be. Additionally, either or both of these may be used in conjunction with model based current control and fault detection by 'phase to model' comparison.

As a sensing scheme, comparison of opposite phases is capable of detecting turn to turn faults, phase to ground faults as well as open circuit phase or open circuit power switch. One difficulty with this detection scheme is that is not easy to establish which of the two phases is at fault when the phase currents do not agree. However, this scheme is very well suited to SR drives which naturally have two equivalent coils in each phase [24].

Owing to time pressure and the successes with fault detection based on 'phase to model comparison', no experimental investigation of 'phase to phase comparison' was undertaken.

7.5 Conclusion

Two fault detection schemes have been presented, both of which work by comparing the sampled phase current to an expected value. The schemes differ in that phase to model comparison uses an expected value derived from a mathematical model, where as phase to phase comparison uses an expected value obtained from an equivalent phase. Both detection schemes should be capable of detecting a similar range of faults but with phase to phase comparison it is not possible to know which phase is faulted if only two phases are compared.

Phase to model comparison can be used to detect a number of faults including open circuit faults in the switching devices and in the motor phase and turn to turn short circuits in the motor phase. It has been shown that this detection scheme is sufficiently sensitive to detect a single shorted turn in the machine phase. This is particularly important as the turn to turn fault is widely regarded as a major cause of electric machine failure. Also of great importance is the ability of the detection scheme to detect the failure of the phase current sensor itself. Thus this scheme is self checking.

Using phase to model comparison, the clearest fault indications are given at full load. It has been shown that the addition of small current pulses to the current profile will allow a single shorted turn to be detected, even at no load.

Using processors currently available the current loop update rate is of the order of 100 μs. Therefore it is not possible to use the current control based schemes developed in this chapter to detect power device short circuits before fault propagation occurs.

Chapter 8

FAULT DETECTION - Part II

8.1 Objectives

In the last chapter the key aspects of fault detection were presented. The fault detection schemes developed were based in software and operated on data sampled once per PWM cycle. These schemes were capable of detecting many faults in the drive. However it was pointed out that the sample rate of once per PWM cycle is not fast enough to protect the power electronics from power device short circuits or phase terminal short circuits.

The aim of this chapter is to develop fault detection schemes capable of detecting faults and taking appropriate action, to prevent fault propagation, within the short circuit withstand time of the power devices, i.e. in substantially less than 10 μs. The rapid response required from these fault detection schemes demands that they are based in hardware. These schemes could take information from any of the sensors discussed in Chapter 6, but will concentrate on processing information from on-state voltage monitors and the freewheel current sensor.

The following sections will examine how the information from the on-state voltage monitor and the freewheel current sensor should be interpreted and demonstrate the fault detection capabilities of each sensor. Two areas that were not discussed in the earlier section on sensors were those of fault diagnosis, and fault confirmation techniques, so these will also be examined here.

8.2 Power Device Fault Imposition

A shorted power device fault is effectively imposed in the demonstrator drive by turning the chosen power device on with a 15 V gate drive. This is simple to achieve by over-riding the gate command from the controller. If on-state voltage monitoring is employed then it must be disabled on the 'faulted' device to ensure that the device remains on, even with excessive current.

An open circuit IGBT fault is imposed by forcing the gate drive of the selected IGBT to -15V. To allow for the manipulation of the gate drive signals, Figure 76 shows that the gate drive commands are not fed directly from the controller to the IGBT driver cards. Instead the gate commands are fed first to a fault imposition circuit where they can be overridden by fault commands, before being passed on to the IGBT driver cards. The desired fault states are instructed to the fault imposition circuit by a fault command box which allows the state of each IGBT to be altered between 'healthy' (operating under the control of the gate commands from the controller), 'forced on' and 'forced off' simply by a set of switches. Note that the fault signals from the IGBT driver cards pass, unaltered, through the fault imposition cards en route to the fault detection card.

8.3 Fast Fault Detection Hardware

It is possible to implement fault detection schemes either on the bridge or in the control electronics. In the case of phase current based detection schemes, their long sampling interval of 100 μs and their strong links with the current control software, mean that it is most natural to implement these schemes in DSP software within the controller.

In contrast, on-state voltage monitoring and freewheel current sensing both operate on a faster time scale, performing fault detection in a small fraction of the controller sampling period. Thus it is essential to implement these schemes in hardware. Figure 76 shows

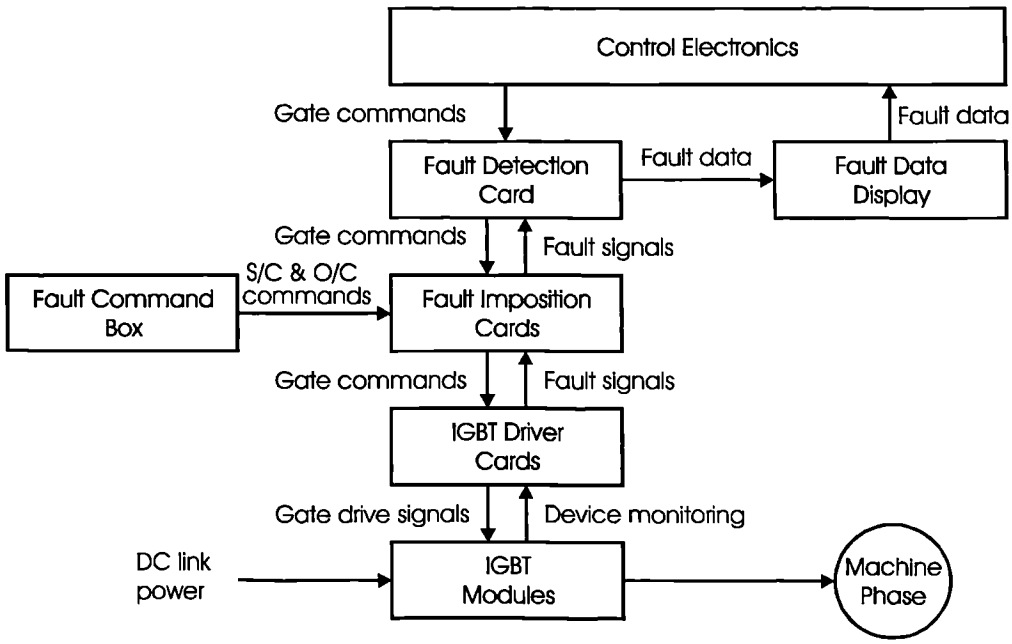

Figure 76 : Fault Imposition and Fault Detection Circuit Organisation

how the hardware based fault detector fits in with the controller, the fault imposition cards and the gate drive cards for a scheme based on on-state voltage monitoring. The scheme is the same for freewheel current sensing except that the sensor inputs to the fault detection board come from current sensors rather than from on-state voltage monitor circuits.

Note that, because the fault detection card is not informed of the fault commands, the faults can be detected only on the basis of the feedback from the fault sensors. Also note that the gate commands from the fault imposition circuit take precedence over the gate commands from the fault detection circuit which in turn take precedence over the gate commands from the controller. Thus the fault detection circuit is relied upon to detect any imposed power device short circuit and take appropriate action, by overriding switching commands from the controller to the remaining 'healthy' devices, in order to avert failure of the power circuit.

The fault detection card samples the sensor signal inputs at 4 MHz. In principle, fault detection in a single 250 ns clock cycle is possible.

8.4 Fault Detection using On-State Voltage Monitors

8.4.1 Introduction to Fault Detection using the On-State Voltage Monitor

The on-state voltage monitor was presented in detail in Chapter 6. The sensor provides a single two-state output, referred to as 'desat'. For the sensor circuit in Chapter 6, desat goes active (i.e. true) when the gate command indicates that the device is on and the voltage across the device exceeds a predetermined reference voltage. At all other times desat is inactive (i.e. false).

The following sections will describe how the on-state voltage monitor may be used in a fault detection scheme. First a black box model will be developed to overcome the difficulties of fault analysis in the gate drive and sensor circuit. Then the type of faults which could be detected by the sensor will be discussed. A fault detection scheme based on the on-state voltage monitor will be presented. The performance of this scheme in detecting power device short circuit failures will be demonstrated. A section is dedicated to answering the important question: 'Can failure of the on-state voltage monitor itself

be detected?'. Finally some improvements to the sensor and detection scheme are suggested.

8.4.2 A Black Box Model for the Power Switch Element

In the last chapter the on-state voltage monitor was examined at a component level. Consideration of faults in the power electronics should include faults in the gate drive circuit. However the gate drive circuit and integrated on-state voltage monitor contain a large number of components, making analysis of all the possible failures difficult. The solution to this problem is to consider the power switch and sensor as a black box.

Viewing the IGBT, anti-parallel diode, gate drive circuit (including the isolated floating power supply) and on-state voltage monitor as a single unit makes it possible to consider faults in any part of the switch, drive or sensor. The black box switch element is shown in Figure 77(a). In essence each power switch is considered as a separate independent unit in the same way that each phase of the electric machine is an independent unit.

The black box model in Figure 77(a) has two power terminals, denoted by 'POWER IN' and 'POWER OUT'. The switch is controlled by a single control terminal, 'ON' and the sensor output is presented at the single monitor terminal, 'DESAT'. Figure 77(b) shows that a single phase power circuit comprises four switch elements and a machine phase element. The power switch control, and monitor terminals have been renamed to give a unique reference for each terminal. The 'ON' input for the Upper device in the Forward leg becomes 'Forward Upper ON' or 'FUON'. Similarly the DESAT terminal on the Lower device in the Reverse leg becomes 'Reverse Lower DESAT' or 'RLD'.

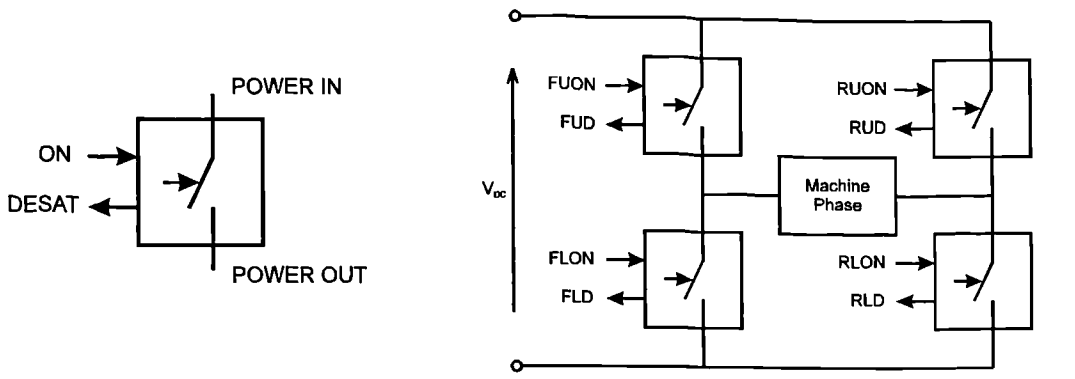

(a) Black Box Power Switch (b) Single Phase Constructed from Black Box Elements

Figure 77 : Black Box Model for a Power Switch

8.4.3 Fault Coverage using the On-State Voltage Monitor

To establish which faults may be detected by the on-state voltage monitor, each fault must be considered in turn and its effect on each on-state voltage monitor output determined. This process is best accomplished by asking which faults will cause a selected on-state voltage monitor output to go active. Symmetry in the bridge ensures that it is a simple matter to extend the analysis for a single device to any other device in the bridge. The possible causes of the on-state voltage monitor within the reverse lower switch, the lower right device in Figure 77(b), going active have been listed in the first column of Table 15.

Clearly the reverse lower on-state voltage monitor output could go active as a result of any one of a wide range of faults. This is the set of faults which may be detected by the sensor. Given such a wide range of faults resulting in an active on-state voltage monitor output, it will be difficult to distinguish between each fault without resorting to a complex fault diagnosis scheme. Any diagnostic procedure is likely to depend on information from other sensors, e.g. the phase current sensor, resulting in a further increase in complexity.

In general, if a single common action were acceptable in the presence of any one of a set of faults then it would not be necessary to distinguish between those faults. The simplest situation would be achieved if a single common post fault operating strategy could be selected for all faults detectable by the reverse lower on-state voltage monitor. The first step towards realising such a strategy is to consider the ideal response to each fault. The ideal response to each fault, detectable by the reverse lower on-state voltage monitor, is presented as the second column of Table 15.

Most of the faults are straightforward and demand a simple, unconditional, response. However the 'reverse upper gate drive open circuit' and 'inadequate gate drive to the reverse lower device' both require a more complicated response. Good circuit design should ensure that the possibility of a fault indication resulting from induced noise is remote but even so this possibility should not be ruled out. The type of disturbance which may result in 'erroneous' noise induced fault detection is likely to affect more than one phase so it is essential that it does not result in a permanent shutdown of the phase.

Fault	Best action	Result of switching all devices off	Result of switching both upper devices on and both lowers off
Ground fault	Switch all devices off	No current flows - OK	Current is delivered into the ground fault - NOT SUSTAINABLE
Machine phase terminal short circuit	Switch all devices off	Rated current is driven in the machine phase - OK	Rated current is driven in the machine phase - OK
Large shorted turn group	Switch both uppers on *(or switch both lowers on)*	Current in excess of rated current is driven in the shorted turn group - NOT SUSTAINABLE	Rated current flows in all turns of the phase and through the power converter - OK
Reverse upper switch short circuit	Switch both uppers on	Current in excess of rated current is driven in the machine phase and faulted inverter - NOT SUSTAINABLE	Rated current flows in the machine phase and through the power converter - OK
Reverse upper gate drive open circuit or Reverse upper has inadequate gate drive *In practice both of these faults can lead to a device turned on with inadequate gate volts. The resulting increased losses will lead to thermal failure.*	Switch all devices off initially, then switch both uppers on if any current flows, then revert to switching all devices off if unipolar current flow results. Repeat the cycle as required.	If the faulted device remains off no current flows - OK If the faulted device conducts then current in excess of rated current is driven in the machine phase and faulted inverter - NOT SUSTAINABLE	If the faulted device remains off, then current in excess of rated current is driven in the machine phase and faulted inverter - NOT SUSTAINABLE If the faulted device conducts, then rated current flows in the machine phase and through the power converter - OK
Reverse lower IGBT open circuit	Switch all devices off *(or continue to operate reverse upper and forward lower switches)*	No current flows - OK	Rated current flows in the machine phase and through the power converter - OK
Reverse lower monitor stuck active	Switch all devices off	No current flows - OK	Rated current flows in the machine phase and through the power converter - OK
Induced Noise	Switch all devices off initially, but continue normal operation when disturbance has passed	No current flows - OK	Rated current flows in the machine phase and through the power converter - OK

Table 15 : Faults Causing an Active Reverse Lower On-State Voltage Monitor Output

Column 2 of Table 15 shows that there are only two different 'best actions' specified to cover the complete list of faults in column 1. Either all the devices should be turned off or both the upper devices should be turned on, depending on the exact fault type. The two final columns in the table give the consequences of applying each of these two post fault control strategies to each fault.

It is apparent from the fourth column of Table 15 that the consequences of turning both the upper devices on and both lower devices off is acceptable for all faults except for the 'reverse upper gate drive open circuit' fault and the ground fault. Leaving these two faults aside for the moment, this one simple action is a suitable response to the active reverse lower on-state voltage sensor signal, irrespective of the exact cause, so no further fault diagnosis is required.

8.4.4 A Simple Fault Detection Scheme Based on the On-State Voltage Monitor

From the analysis of a single power switch, it is now possible to devise a very simple fault detection scheme based on the output of the on-state voltage monitor attached to that switch. Again using the reverse lower switch as an example, during normal operation the gate commands calculated by the current controller are applied to the power switches. If at any instant the reverse lower on-state voltage monitor output goes active then all the switches are turned off for a 10 μs dead-time, after which both upper devices are commanded on and both lower devices are commanded off.

The symmetry of the single phase bridge makes it easy to extend this fault detection scheme to include on-state voltage monitor inputs from all four power switches. If the output of the on-state voltage monitor on either lower device goes active then the gate commands calculated by the current controller must be replaced by gate commands to turn both upper switches on and both lower switches off. Conversely the output of the on-state voltage monitor on either upper device going active should result in both upper switches being turned off and both lower switches being turned on. This fault detection scheme is presented as a simple state transition diagram, shown in Figure 78. The fault detection scheme has been realised as a state machine embedded in the fast fault detection hardware introduced at the start of the chapter. The effectiveness of the resulting system is the subject of the next section.

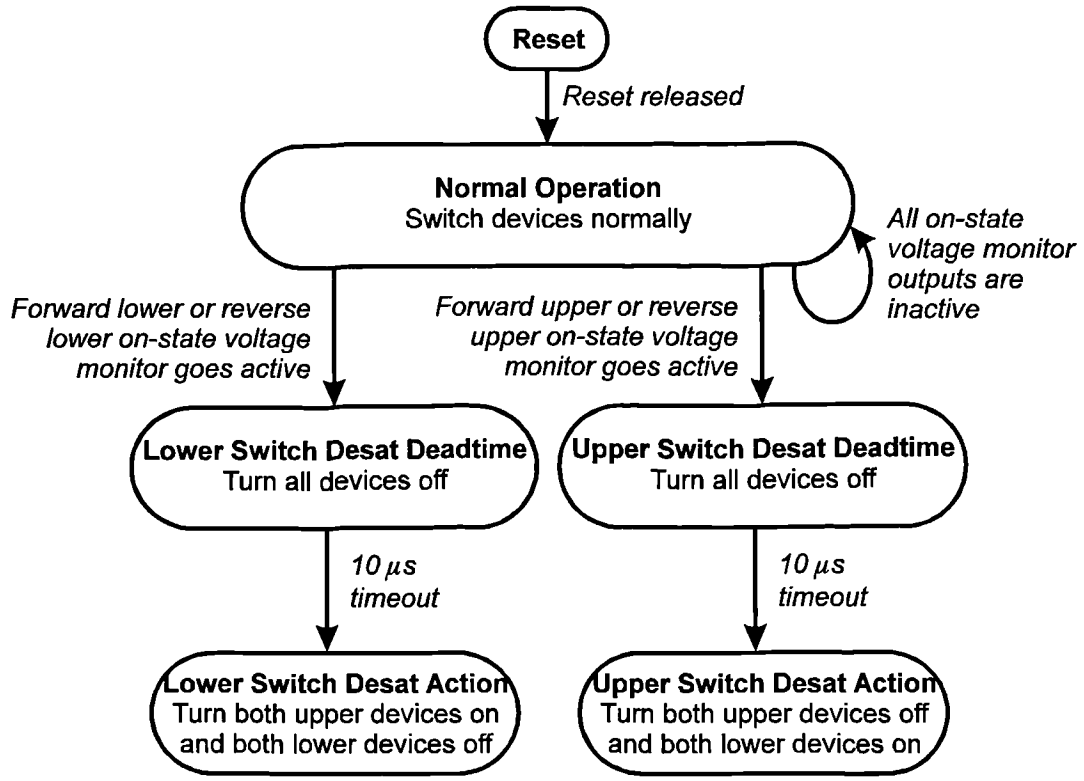

Figure 78 : Flow Diagram for Fault Detection Based on the On-State Voltage Monitor

8.4.5 Short Circuit Power Device Detection

There are two fault conditions under the heading 'power device short circuit'. These are termed 'hard fault' and 'fault under load' when discussed in a paper by Chokhawala et al [63]. In this paper Chokhawala considers schemes for protecting IGBTs from short circuits applied to the output of the inverter. A 'hard fault' is one which is already present at the time when the IGBT turns on. A 'fault under load' is a short circuit which arises when the IGBT is already conducting. Chokhawala concludes that the 'fault under load' is the more serious of the two faults because for this fault the gate emitter voltage is added to by charge supplied via the Miller capacitance. When examining short circuit faults, most effort will be focused on the 'fault under load' case and the effect of the Miller capacitance will be noted. For an IGBT short circuit fault, the 'hard fault' results from a device failing to turn off, whereas the fault under load results from the unexpected turn on of a device.

The oscillogram in Figure 79 shows the effect of an imposed IGBT short circuit fault. The left hand half of the oscillogram shows the reverse upper and reverse lower gate drives in normal operation for 2 PWM cycles. The gate drives to the inverter leg form a

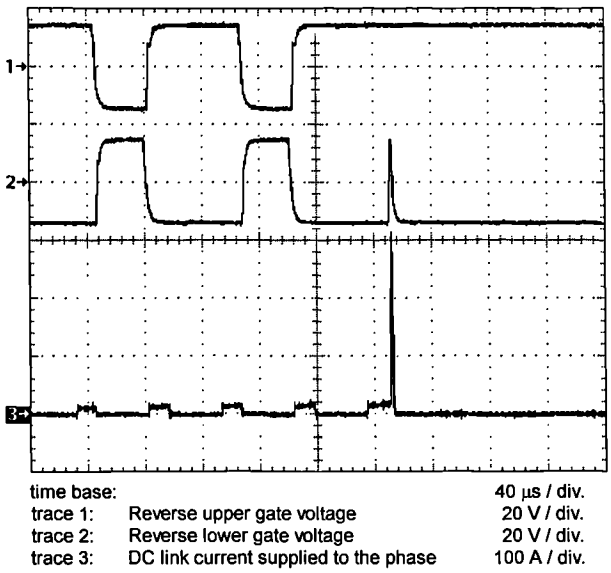

Figure 79 : Gate Drives and DC Link Current during a Power Device Short Circuit Fault

complimentary pair with the upper drive off whenever the lower drive is on and vice versa. Current is supplied to the phase in pulses of 15A. In the right hand half of the oscillogram, the gate drive to the reverse upper IGBT has been forced to remain high. When the lower gate is driven high, by the continued PWM operation of this device, the DC link is shorted through the leg and a fault current of 320 A flows. The fault is detected and the fault current is interrupted by turning the lower device back off within the short circuit withstand time of the device.

A power device short circuit fault is shown in Figure 80 on a faster timebase than that used in Figure 79. Note that whereas the fault shown in Figure 79 was effectively a 'hard fault', the fault in Figure 80 is a 'fault under load'. The fault was imposed by forcing the upper, 'faulted', device to turn on whilst the lower, 'healthy' device is conducting.

Note the dip in the gate volts on the 'faulted' device at the time when the 'healthy device turns off. At the point immediately before the 'healthy' device turns off, the 'faulted' device has not had time to turn on fully so the fault current is limited by the 'faulted' device and the full DC link voltage is dropped across it. As the 'healthy' device turns off, the volts across the 'faulted' device fall rapidly and discharge the gate of the 'faulted' device through the Miller capacitance.

As the volts across the 'faulted' device fall, so the volts across the 'healthy' device rise, charging the gate of the 'healthy' device through the Miller capacitance and extending

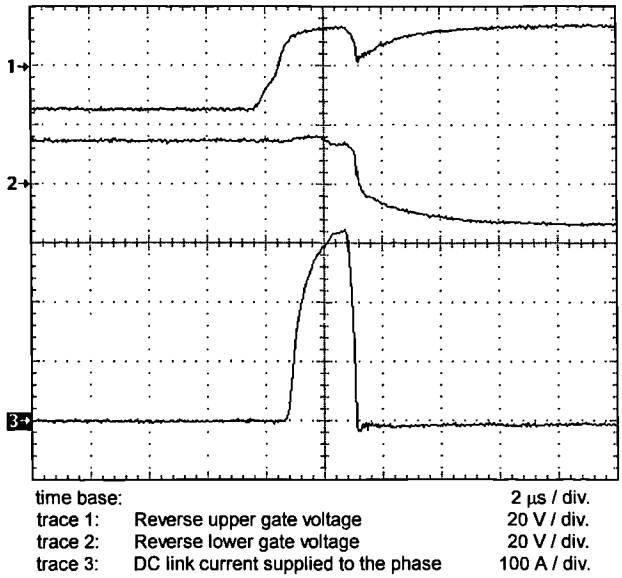

Figure 80 : Measured Gate Drives and Fault Current for Reverse Upper IGBT Short Circuit Fault

the turn off time. In the case of the 'fault under load', the Miller effect tends to oppose efforts to limit the fault current in the healthy device by controlling the gate voltage.

Sensor Operation during the Fault

The behaviour of the sensor during a power device short circuit fault is shown in Figure 81. The case shown in the oscillogram is a 'fault under load' for the healthy power switch in the faulted limb. This is apparent because the voltage measured across the device, trace 1 in the oscillogram, is near-zero when the fault current starts to flow,

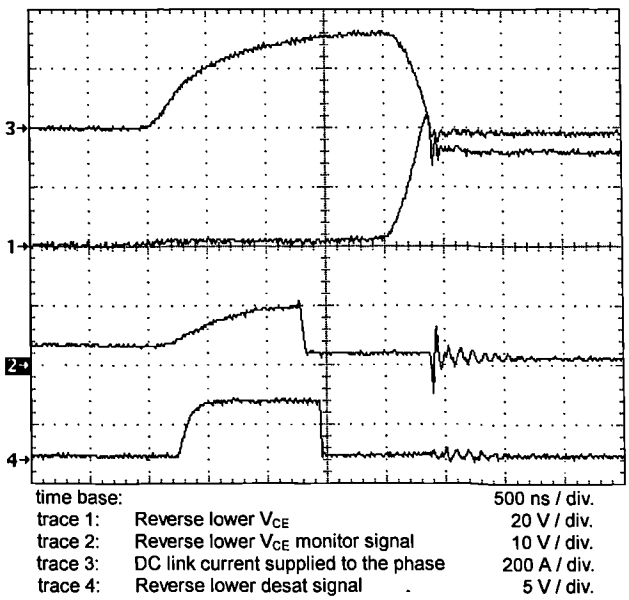

Figure 81 : Sensor Behaviour under Reverse Upper IGBT Short Circuit Fault

indicating that this device is already on when the fault is imposed.

At the point when the fault is imposed, the fault current supplied to this phase by the DC link, trace 3 in the oscillogram, rises towards 320 A. This causes the voltage across the healthy device carrying this current to rise. Whilst the healthy device remains on, the rise in V_{CE} is tracked by the sensor circuit 'V_{CE} monitor signal' and this signal, trace 2 in the oscillogram, is fed into a comparator. The comparator drives the 'desat signal' active whenever the V_{CE} monitor signal exceeds 6.8 V and this can be seen happening in traces 3 and 4 of the oscillogram.

The detection circuit takes approximately 0.5 μs (2 cycles of the detector circuit 4.0 MHz clock) to recognise the active desat signal and a further 0.5 μs to apply the new set of gate commands to the bridge. The instant that new gate commands are applied to the bridge is apparent on the oscillogram as the V_{CE} monitor signal falls to zero when the 'healthy' device receives a turn off command. From this point a further 1.1 μs is required for the 'healthy' device to turn off.

The fault detection scheme based on the on-state voltage monitor has been shown to successfully detect the fault and take appropriate action to interrupt the fault current within 3 μs. This is easily fast enough to protect the healthy device and to ensure that the DC link remains undisturbed.

In the results presented, the only evidence that the correct post fault action has been taken is that the DC link current falls to zero after the fault. Since the post fault control strategy is independent of the fault detection scheme, the effectiveness of the post fault operation will be discussed in the next chapter.

In the results presented, the fact that the DC link current falls to zero, after the fault, has been taken as sufficient evidence that the correct post fault action has been taken. The effect of the post fault control on the phase current and torque has not yet been adequately considered. Since the post fault control strategy is independent of the fault detection scheme, detailed discussion of post fault operation appears separately in Chapter 9.

8.4.6 On-State Voltage Monitor Failure Detection

The importance of detecting every failure in a fault tolerant system, to prevent the build up of undetected faults in a system, has been discussed. This must include detection of fault sensor failures. For the on-state voltage monitor, the desat output stuck active is automatically detected because it is seen as a indication that the corresponding device is carrying excessive current. The action taken for the latter fault is also acceptable for the former. However an on-state voltage monitor output stuck inactive is not automatically detectable. This is a serious problem which will be addressed in the next section.

8.4.7 Improving the Fault Detection Scheme

The operation of a simple fault detection scheme using the on-state voltage monitor has been demonstrated. The scheme presented is capable of detecting many faults including a shorted power device or a shorted phase with sufficient speed to allow fault propagation to be prevented. This section will examine how the scheme and sensor can be improved to eliminate each of the following deficiencies:

- Ground faults result in the same fault indication as other faults but require different action.
- The best course of action to deal with an open circuit gate drive has not been addressed
- The best course of action to deal with fault indications caused by externally induced noise has not been addressed
- It is not possible to detect when the sensor output is stuck inactive

Differentiating Ground Faults from Other Faults

In the section on phase current based detection schemes, it was shown that a ground fault within the machine is most likely to result from an undetected turn to turn short. By employing a turn to turn short circuit detection scheme, the chance of a ground fault within the machine should be considerably reduced. However, the discussion did not cover the power converter or the interconnections to the machine, where the possibility of a ground fault cannot be ruled out.

Consider a power converter configuration fed from a DC link whose centre point is connected to ground. Suppose the phase develops a ground fault whilst one of the lower devices is turned on. Fault current will flow from the earth fault to the DC link negative rail, through the lower device, as shown in Figure 82. In the simple fault detection

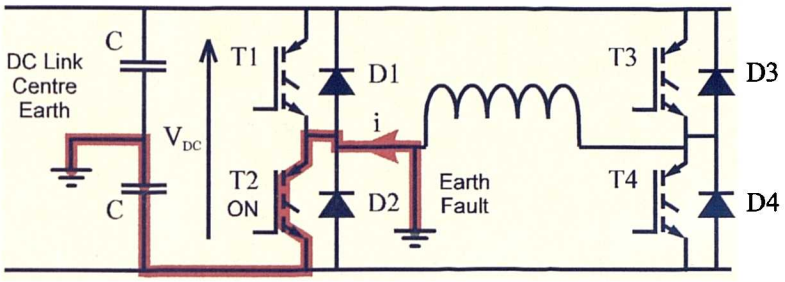

Figure 82 : Current Path Created by Turning a Lower Device on into an Earth Fault

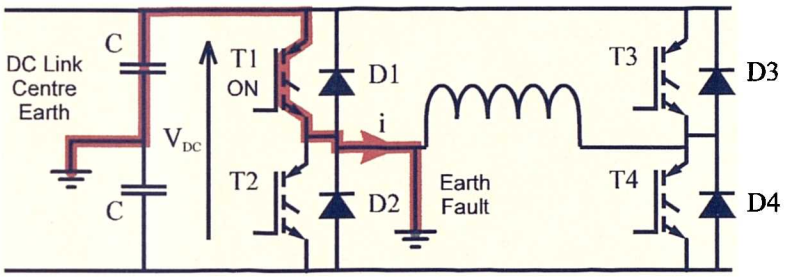

Figure 83 : Current Path Created by Turning an Upper Device on into an Earth Fault

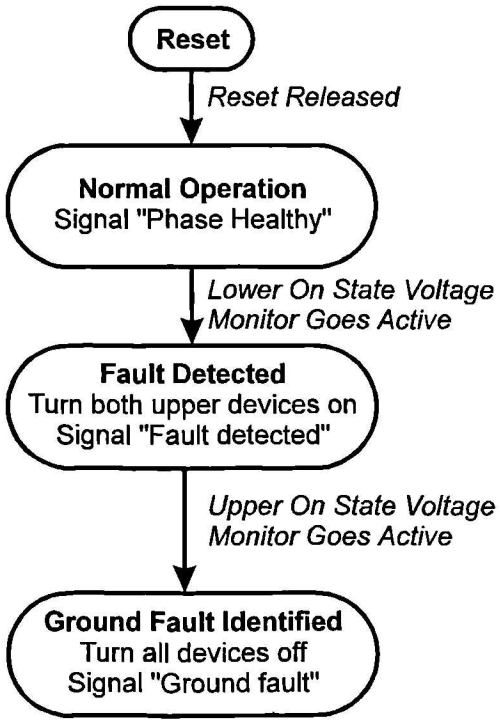

Figure 84 : Procedure for Handling Ground Faults

scheme, the resulting active desat signal from the lower switch will cause both lower switches to turn off and both upper switches to turn on. In the case of the ground fault, turning both upper devices on will result in another pulse of fault current of the opposite polarity to the first, Figure 83. Fortunately this situation will also be detected, as one of the upper device on-state voltage monitor outputs will go active. At this point all devices must be driven off and held off. This can be expressed in a flow diagram as shown in Figure 84.

Handling the Open Circuit Gate Drive Fault

It was established in Chapter 6, that the 'open circuit gate drive' fault could result in intermittent operation of the associated power switch. This is because, even whilst the gate is no longer being driven by the gate drive circuit, it is still coupled to the power device output by the (parasitic) Miller capacitance. As a result it is difficult to devise a simple post fault control scheme for the open circuit gate drive fault. One possible solution was hinted at in Table 15. This is illustrated in terms of the flow diagram in Figure 85.

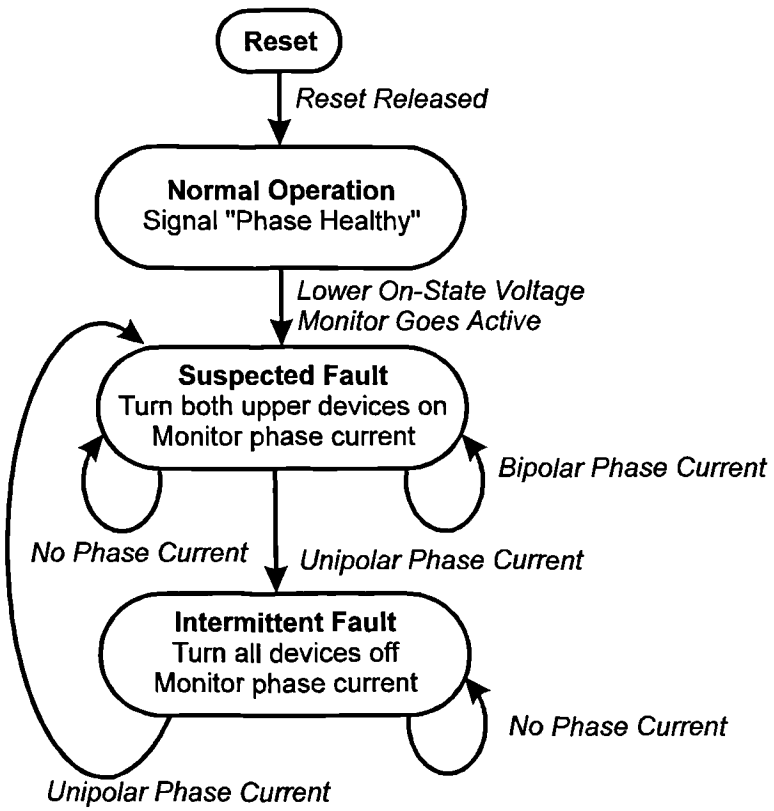

Figure 85 : Procedure for Handling Open Circuit Gate Drive Faults

Avoidance of Drive Shut Down Resulting from Induced Noise

Permanent shut down of a phase due to an erroneous fault indication could be avoided as follows. Upon any fault indication, the action appropriate to that indication would be applied for a short interval which would act as a 'cooling off period'. This would be followed by a short trial period of resumed operation as shown in Figure 86. If a further fault indication is generated during the trial period then the fault is confirmed. If the short trial period passes without further fault indication the first fault indication is disregarded and normal operation is resumed.

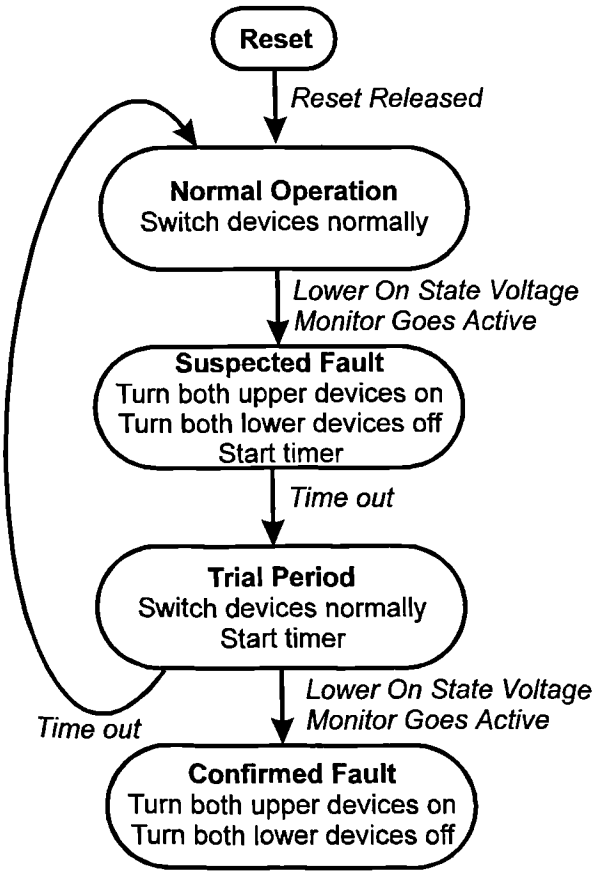

Figure 86 : Procedure for Confirming Faults

Self Checking On-State Voltage Monitor

How can the on-state voltage monitor be configured to enable faults in the sensor itself to be detected?

The solution to this problem is to configure the sensor so that it monitors the device voltage even when the device is off. The sensor would be monitored at all times, not just when the corresponding gate drive is active. The PWM cycle would be split into 4

periods as follows: (1) the first 2 μs after the device turns on during which time the sensor output is ignored; (2) from 2 μs after the device turns on, until the device turns off the sensor output is monitored and an active output indicates a fault; (3) the first 2 μs after the device turns off during which time the sensor output is ignored; (4) from 2 μs after the device turns off, until the device turns on during which time the sensor output is monitored. In period (4) a healthy sensor output will go active during those periods when the anti-parallel diode is not conducting. Taking the reverse lower switch as an example, the anti-parallel diode will only conduct when negative phase current flows. Therefore during periods of positive phase current an inactive reverse lower on-state voltage monitor output is an indication of the failure of that sensor.

8.5 Freewheel Current Sensing

8.5.1 Introduction to Freewheel Current Sensing

It is worth quickly summarising the essential features of the freewheel current sensor, developed in Chapter 6. Recall that the freewheel current sensor, shown in Figure 87, is placed so as to measure only freewheeling and shoot-through currents in the H-bridge. By correlating the measurement from this fault detection sensor with the phase current sensor output and the switching state of the H-bridge, the health of the bridge may be determined. Whilst this technique requires one additional sensor per phase, Craig claims that it can be used to identify the power device which has failed and the failure mode. If this is true then this simple additional sensor could prove to be invaluable in fault detection and identification.

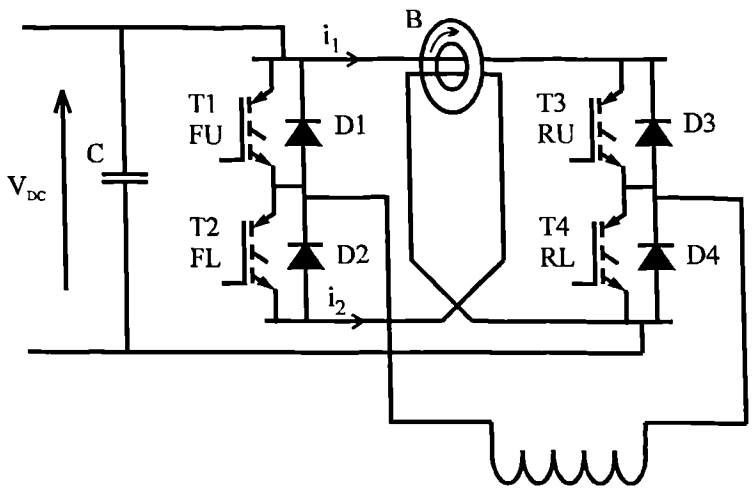

Figure 87 : H-Bridge Configured to Include a Freewheel Current Sensor

The freewheel current sensor is arranged to provide a positive output when positive phase current is freewheeling in the upper path. Positive phase current freewheeling in the lower path will result in a negative output from the freewheel current sensor.

In the following sections, it will often be convenient to refer to fault conditions using a set of abbreviations. The proposed set of abbreviations is defined in Table 16. Abbreviations will also be used to refer to the power device switching commands. These abbreviations, FL,FU,RL,RU refer to the gate switching commands for the forward lower, forward upper, reverse lower and reverse upper devices respectively. Occasionally the four switching commands will grouped together and referred to as the switching state. For example, the switching state: FU;FL;RU;RL = 1001 would indicate that the forward upper and reverse lower devices have been commanded on and the forward lower and reverse upper devices have been commanded off.

Fault Abbreviation	Fault Description
Ctrl Fault	Controller fault. i.e. the current controller has supplied an illegal set of power device switching commands.
FLOC	The Lower power device in the Forward leg has failed Open Circuit
FLSC	The Lower power device in the Forward leg has failed Short Circuit
FUOC	The Upper power device in the Forward leg has failed Open Circuit
FUSC	The Upper power device in the Forward leg has failed Short Circuit
ITSC	A Short Circuit fault is present at the Inverter Terminals
MTSC	A Short Circuit fault is present at the Motor Terminals
RLOC	The Lower power device in the Reverse leg has failed Open Circuit
RLSC	The Lower power device in the Reverse leg has failed Short Circuit
RUOC	The Upper power device in the Reverse leg has failed Open Circuit
RUSC	The Upper power device in the Reverse leg has failed Short Circuit

Table 16 : Fault Abbreviations

8.5.2 Information Required for the Fault Detection Scheme

Chapter 6 suggested that this sensor can detect a wide range of faults by dividing the measured current into a small number of discrete bands - but how many bands are required and where should the levels separating these bands be set? A few simple faults will be considered and it will become apparent that the answer to the question depends to some extent on the range of faults which is to be detected.

Machine Phase Terminal Short Circuits

A machine phase terminal short circuit will result in excessive current delivered into the

fault by the inverter as the power devices switch the DC link voltage into the short circuit. This current will flow through the phase current sensor, but will not be seen by the freewheel current sensor. This fault can be detected by using information from the phase current sensor alone, as shown in Table 17.

Iph in range	No Fault
|Iph| >>0	MTSC

Table 17 : Phase Terminal Short Circuit Detection Table

Power Device Short Circuits

In the case of a short circuit power device, when the other power device in the faulted limb turns on, excessive current will be delivered into the fault from the DC link. This shoot-through current will not flow through the phase current sensor, but will be seen by the freewheel current sensor. Detection of machine phase terminal short circuits is combined with detection of power device short circuits in Table 18.

	Ifw<<0	Ifw in range	Ifw>>0
Iph in range	RUSC or RLSC	No Fault	FUSC or FLSC
|Iph| >>0	Not possible	MTSC	Not possible

Table 18 : General Short Circuit Detection Table

Using the freewheel current and phase current sensors, it is not possible to separate a lower device short circuit from an upper device short circuit. For example the forward upper device short circuit and the forward lower device short circuit both result in fault current flowing through the freewheel current sensor in the positive sense. In both cases the fault current bypasses the phase current sensor so there is nothing to distinguish between these 2 faults. This is a serious failing as the remedial action required for each of these faults is very different and, in each case, application of the alternate action would be catastrophic. A solution to this problem is found through monitoring the switching commands.

For example, if the phase current and freewheel current measurements indicate that one of the two devices in the forward leg is short circuit and the switching commands indicate that the forward lower device should be on, then the forward upper device must be short circuit. Moreover the switching commands may be used to monitor the current controller as any occurrence of a switching command with both devices in a limb commanded on together is a clear indication of a fault in the current controller.

Power Device Open Circuits

If open circuit faults are to be detected then adding in the switching states alone still does not provide sufficient information and further current levels must be added. The phase current sensor output must be encoded into 5 bands - zero, positive, negative, excessive positive and excessive negative. The freewheel current sensor output must also be encoded into 5 bands, representing the same bands as the phase current sensor. In practice the zero band must be a 'near zero band' and the phase current sensor zero band cannot be the same size as the freewheel current sensor zero band. The selection of the levels used to distinguish between current bands is discussed in more detail in Appendix G.

In subsequent sections the 5 non-overlapping bands will be referred to by the symbols as follows:

'<<0' $\Rightarrow$ negative out of range

'<0' $\Rightarrow$ negative, in range

'=0' $\Rightarrow$ near zero

'>0' $\Rightarrow$ positive, in range

'>>0' $\Rightarrow$ positive, out of range

8.5.3 Fault Detection and Identification

Craig [50] approaches fault diagnosis by looking at each possible conduction situation and asking whether normal operation could cause it or what faults would be required to create the particular situation. For some situations Craig observes that a multiple failure is required but in the situations which could result from normal operation he disregards

the possibility of a fault and in the situations which could result from single point failures he disregards the possibility of a multiple failure. In contrast to Craig's approach, here each single point failure is considered and the resulting conduction situation is determined.

The fault diagnosis table for short circuit detection, Table 18, contained 2 phase current levels and 3 freewheel current levels - a total of 6 cells. The addition of 16 switch command combinations will increase the number of cells to 96. The additional current bands required for open circuit detection will result in a fault detection table with 5 phase current levels, 5 freewheel current levels and 16 switch command combinations - a total of 400 cells. Each of these cells could indicate one of 12 fault conditions. Handling this amount of information using intuition alone is prone to error. A systematic technique for fault diagnosis would be of great value in compiling the fault detection table. Such a method has been devised by the author, based on the Karnaugh map used for state minimisation in simple digital systems.

A 4 bit Karnaugh map allows each possible switching state to be represented by assigning one bit to represent each switch. Immediately it is possible to see that in a healthy phase the bridge cannot exist in all the states represented on the diagram. Any state where both devices in a limb are on together, shown as the red region in Figure 88, is an illegal or faulty state which will result in shoot-through.

The exact location on the map corresponding to healthy operation of the phase can be determined easily from the switching commands. For example the situation with both lower devices commanded on is illustrated in Figure 89. The expected location, given by

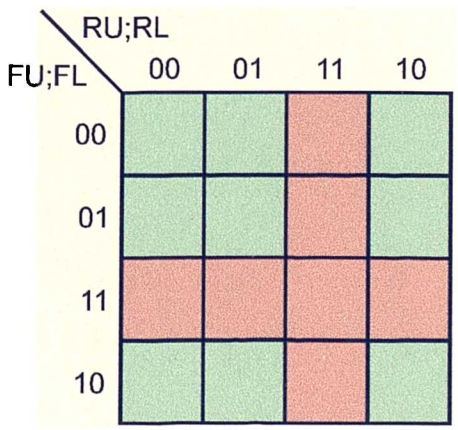

Figure 88 : Illegal Switching Commands on the Karnaugh Map

the switching state FU;FL;RU;RL = 0101, is shaded green and marked 'OK'. Any single point failure would result in a switch being on when commanded off or being off when commanded on. All these situations result in a move of one square horizontally or vertically. There are always 4 possible switching states resulting from a single point failure, as shown in red in Figure 89. The cells resulting from a single failure are marked with the corresponding fault. From comparison with Figure 88 it is clear that 2 of the fault conditions will result in shoot-through. Finally the possible switching states resulting from a double failure, reached by 2 horizontal or vertical moves on the map, are shaded in grey. Whilst this shows that the map could be used to analyse double failures it also suggests that the wide range of states possible would make the analysis very difficult.

For a given switching state and phase current, the freewheel current sensor output can be calculated. The freewheel current sensor output for positive phase current has been calculated and marked into the five cells of interest in Figure 89. Now, for the switching state FU;FL;RU;RL = 0101 and positive phase current, it is possible to infer the fault condition by matching the actual freewheel current sensor output to the predicted freewheel current sensor outputs. From Figure 89, for the given switching state and phase current, a freewheel current sensor output in the '>0' band would be expected. If the freewheel current sensor output were in the '=0' band, this would indicate a move from the 0101 square into the 0001 square, this would be a clear indication of a forward lower device open circuit fault.

The expected freewheel current level for a given switching state depends only on the

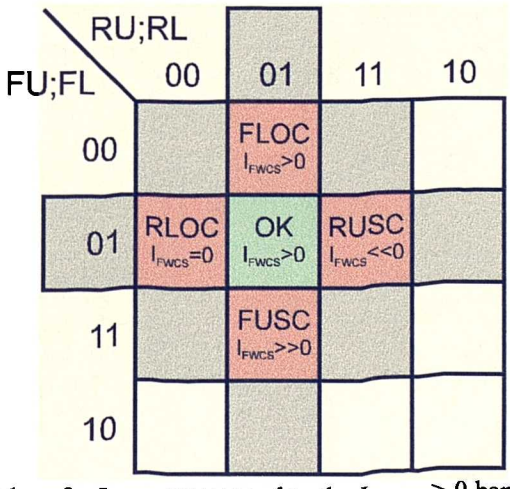

Values for I_{FWCS} correspond to the $I_{PHASE} > 0$ band

Figure 89 : Using a Karnaugh Map for Fault Diagnosis

phase current level. Figure 90(a) and Figure 90(b) show the currents seen by the freewheel current sensor for positive and negative phase current respectively. Now using these maps it is a simple matter to compile the table required for fault diagnosis.

Detectability

The map shows which faults are detectable at any point. For operation with positive phase current in the 0101 square the forward lower device open circuit is not detectable as this fault corresponds to moving to operation in the 0001 square - a move which does not result in a change in current seen by the freewheel current sensor. All other moves from the 0101 square do result in a change in current seen by the freewheel current sensor and are detectable. The undetectability of a fault in one square is not a serious problem as in the normal course of PWM switching an operating square (switching pattern) will soon occur for which the fault is detectable. In instances where the value of freewheel current for a faulted state is the same as for fault-free operation the phase must be assumed to be healthy.

Unique Identification

For operation with a given switching state and a given phase current it is possible to determine whether the 4 possible failures are uniquely identifiable. For operation with positive phase current in the 0101 square all 4 faults result in a different current seen by the freewheel current sensor, thus each fault is uniquely identifiable. However for

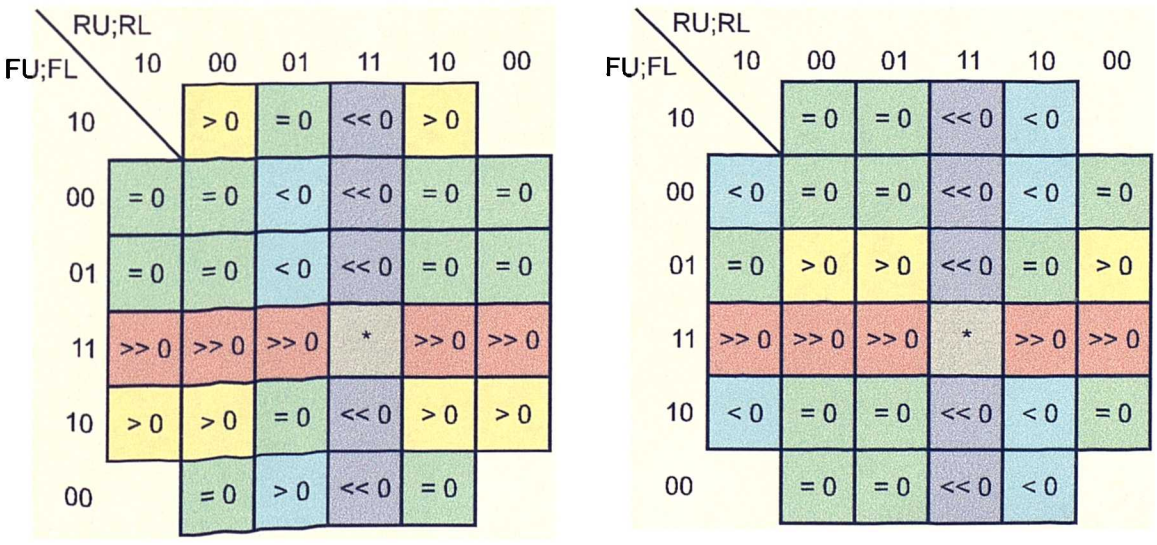

(a) $I_{PHASE} > 0$ (b) $I_{PHASE} < 0$

Figure 90 : Extended Karnaugh Maps for Fault Diagnosis

operation with positive phase current in the 0001 square 2 faults both result in a transition in the current seen by the freewheel current sensor from the '< 0' band to the '= 0' band. Thus these 2 faults (reverse lower open circuit and forward upper short circuit) are indistinguishable.

Fault Detection Lookup Table

Analysis of the possible switching states and current levels results in the fault detection and identification table below. This forms the basis of the detection scheme. Note that

Iphase	<0					~0					>0					<<0 or >>0
Ifwcs	<<0	<0	~0	>0	>>0	<<0	<0	~0	>0	>>0	<<0	<0	~0	>0	>>0	X
Switches (=FU;FL ;RU;RL) = 0000	-	RU SC	No fault	FL SC	-	-	No fault	No fault	No fault	-	-	RL SC	No fault	FU SC		-
0001	RU SC	-	No fault	FL SC	-	RU SC	No fault	No fault	No fault	-	RU SC	No fault	FU SC or RL OC	-	-	-
0010	RL SC	No fault	RU OC or FL SC	-	-	RL SC	No fault	No fault	No fault	-	RL SC	-	No fault	FU SC	-	-
0011	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault
0100	-	-	RU SC or FL OC	No fault	FU SC	-	No fault	No fault	No fault	FU SC	-	RL SC	No fault	-	FU SC	-
0101	RU SC	-	FL OC	No fault	FU SC	RU SC	No fault	No fault	No fault	FU SC	RU SC	No fault	RL OC	-	FU SC	-
0110	RL SC	FL OC	No fault	RU OC	FU SC	RL SC	No fault	No fault	No fault	FU SC	RL SC	-	No fault	-	FU SC	MT SC
0111	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault
1000	-	RU SC	No fault	-	FL SC	-	No fault	No fault	No fault	FL SC	-	-	RL SC or FU OC	No fault	FL SC	-
1001	RU SC	-	No fault	-	FL SC	RU SC	No fault	No fault	No fault	FL SC	RU SC	FU OC	No fault	RL OC	FL SC	MT SC
1010	RL SC	No fault	RU OC	-	FL SC	RL SC	No fault	No fault	No fault	FL SC	RL SC	-	FU OC	No fault	FL SC	-
1011	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault
1100	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault
1101	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault
1110	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault
1111	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault	Ctrl fault

Table 19 : Fault Lookup Table for Freewheel Current Sensing

some cells are empty. These represent situations which cannot be reached, either in normal operation or with any of the single point failures considered.

A few cells contain two possible faults. In these cases the presence of a fault has been detected but the fault has not been uniquely identified so there are 3 options open. Either continue operation until the fault is uniquely identified or cease normal operation and begin fault investigation or apply action which would be suitable for all possible faults, assuming such an action exists.

Every instance of a state where a fault has been detected but cannot be uniquely identified occurs during a transient switching state. These transient switching states only last for the dead time. None of the ambiguous states is associated with a large fault current and all ambiguity will disappear when the next switching state is applied. For these reasons the most appropriate response to an ambiguous detection is to continue normal operation until a definite fault identification is made.

The cells corresponding to gate commands which would result in both devices in a leg turned on simultaneously, are marked to indicate a controller fault. The correct response to a controller fault is to turn all the devices off.

8.5.4 Practical Implementation

The freewheel current sensor implementation, presented in Chapter 6, produced an output with a significant noise content. Even with this relatively poor freewheel current sensor output, it is still possible to implement a robust detection scheme. To avoid problems with noise on the freewheel current sensor output, for the first 16 μs after each change in bridge switching state only the '>>0' and '<<0' levels are monitored. Hence this delay does not affect the detection of those faults which result in excessive fault current.

To further minimise the possibility of inaccurate fault detection, any fault must persist for 6 cycles of the state machine before it is acknowledged and acted upon. This adds 1.5 μs to the detection time for all faults, but this is still only 15% of the short circuit withstand time of the power devices. The resulting state machine is detailed in Appendix G.

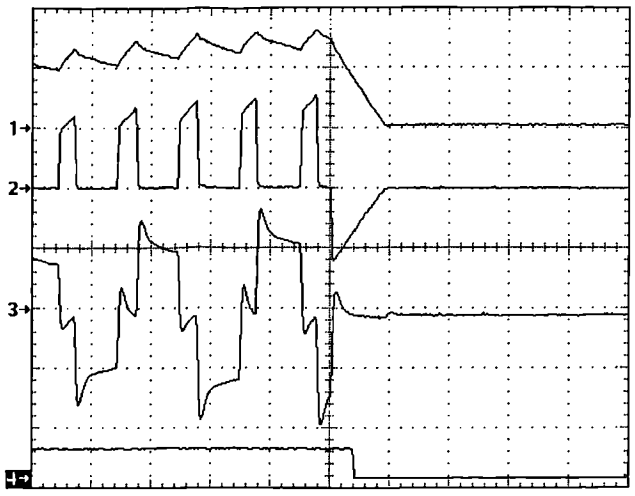

Figure 91 : The Introduction and Detection of a Reverse Lower IGBT Open Circuit

Power Device Open Circuit

The response of the phase current, the current delivered to the phase from the DC link and the freewheel current transducer output during the introduction of a reverse lower IGBT open circuit fault are shown in Figure 91. It is possible to tell that the fault was introduced during a lower loop freewheeling period because the freewheel current transducer indicates a current flowing in the freewheeling path with a polarity opposite to that of the phase current. At the instant that the fault is introduced, the current which was circulating in the freewheeling path is blocked by the open circuit device. The current must then commutate into a new path and this results in it flowing back to the DC link. It can be seen that the current returned to the DC link after the introduction of the fault decays rapidly, as its flow is opposed by the full DC link voltage.

This fault was detected and identified as a RLOC fault, as expected. The signal displayed as trace 4 in Figure 91 goes from low to high at the instant that the fault detection circuit replaces the gate commands from the current controller with gate commands determined by the post fault operating strategy. From the timing of the transition in this signal, it is apparent that the detection does not take place until 16 µs after the introduction of the fault. This delay corresponds to the 16 µs after each change in bridge switching state during which only the '>>0' and '<<0' levels from the freewheel current sensor are monitored.

Note that an open circuit IGBT can only be detected if it occurs at a time when there is sufficient phase current of the correct polarity. This is because the fault detection scheme relies on detecting current which is forced to return to the DC link rather than circulating around a freewheeling path. The device failure prevents the build up of current, making subsequent fault detection impossible.

Power Device Short Circuit

The response of the phase current, the current delivered to the phase from the DC link and the freewheel current transducer output during the introduction of a reverse lower IGBT short circuit fault are shown in Figure 92. The two oscillograms in the figure are of the same event but on different timebases. The particular case shown is a 'hard fault'. The command to force the reverse lower IGBT on is introduced at an instant when this device is already on. Thus it is subsequently unable to turn off and the reverse upper device duly turns on into the resulting fault. The fault current builds to 200 A in 1.0 μs. The freewheel current transducer output slews towards the negative rail at a rate limited by the filtering on the transducer output. The polarity of the freewheel current transducer output is opposite to the polarity of the DC link current because the shoot-through occurs in the reverse leg. Trace 4 in the oscillograms indicates that the fault is acknowledged and the post fault switching commands are applied by the detection circuit 3.0 μs after the introduction of the fault. After a further 0.5 μs the reverse upper gate

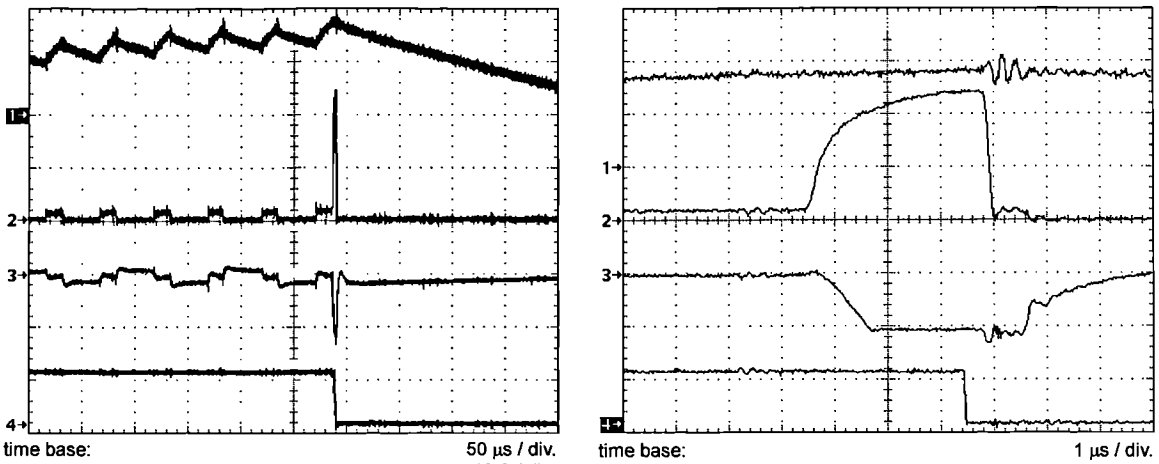

(a) Event in Context (b) Event in detail

Figure 92 : Introduction and Detection of a Reverse Lower IGBT Short Circuit Fault

voltage has been reduced to a level where the fault current is rapidly reduced from the peak value approaching 300 A to zero. The shoot-through current flowed for a total of 3.5 μs, considerably less than the 10 μs short circuit withstand time of the power device. This fault was reported as a RLSC immediately upon detection.

8.5.5 Freewheel Current Sensor Failure

The sensor outputs are exercised continuously, therefore any failure of the current transducer on which the sensor is based will quickly be seen by the detector. It will not be possible to differentiate between a sensor failure and the fault indicated by the erroneous sensor output but the phase will nevertheless be shut down in a controlled fashion in response to the reported failure.

The output of the transducer is digitised into 5 discrete levels. In normal drive operation the much greater than zero and the much less than zero current bands are not entered so a failure of one of the digital signals corresponding to these levels would not be detected. This is a serious problem.

Craig [50] observes that because the freewheel current sensor and associated health assessing circuitry are not in the main path, the addition of this fault sensor does not affect the reliability or fault tolerance of the drive system. This does not appear to be correct because if the sensor itself fails then its output will still be acted on with the likely loss of the torque contribution of the associated phase. The addition of the extra parts inevitably results in an increased failure rate.

8.6 Conclusion

Two fault detection schemes have been presented. The first scheme used on-state voltage monitors whereas the second scheme was based on the freewheel current sensor but also relied on information from the phase current sensor.

On-state voltage sensing is already used in many gate drive circuits to protect power devices from damage due to excessive current. Information from on-state voltage monitoring has been used here to detect a range of short circuit faults. It has been shown that it is a simple matter to adapt the sensor and detection scheme to detect open circuit IGBT faults and to detect failure of the on-state voltage monitor itself.

A fault detection scheme based on the freewheel current sensor has been developed and implemented using a freewheel current sensor integrated into a power hybrid. The ability of this scheme to detect open circuit and short circuit faults has been demonstrated. However the detection scheme does require an additional sensor, albeit a simple one and is not naturally capable of detecting all the failure modes of this extra sensor. Moreover freewheel current sensing is not capable of detecting any faults which cannot be detected with a detection scheme based on on-state voltage monitoring.

This work has shown that in practice each drive fault can be detected using either phase current sensing or on-state voltage monitoring. Many drives already include a phase current sensor and on-state voltage monitors, thus the original goal of detecting all possible drive faults without incorporating any additional sensors has been achieved.

If current sensing IGBTs are introduced successfully and widely adopted then it may be possible to eliminate the phase current sensor and the on-state voltage monitors from most drives. The current sensing IGBTs will give as many current signals as there are IGBTs carrying current thus providing redundant current sensing. This situation will require its own special detection and diagnostic scheme but will naturally borrow many of the techniques from the phase current and on-state voltage monitoring schemes presented here.

Chapter 9

FAULTED OPERATION

9.1 Objectives

This chapter will examine each of the faults introduced in Chapter 5. For each fault, appropriate post fault control will be developed. With fault detection strategies from Chapter 7 and Chapter 8 and an appropriate post fault control strategy it will be possible to operate the drive during and after the introduction of any fault. The final aim is to demonstrate that operation up to, during and beyond the introduction of even the most serious faults is possible. The level of disturbance to the overall drive function, resulting from the introduction of each fault and the application of appropriate post fault control, will be assessed.

9.2 Post Fault Control

9.2.1 Phase Winding Open Circuit

The incidence of this fault prevents any current from flowing in the faulty phase winding so, in theory, no specific post fault control is required for this phase. In practice all attempts to drive current into the faulty phase should cease. In the event of an intermittent fault, discontinued operation of the phase will prevent the intermittent nature of the fault being identified but will also prevent the fault causing erratic operation. More importantly the cessation of operation of the damaged phase will minimise the stress on the remaining healthy components and thus minimise the likelihood of fault propagation.

The action taken in response to an open circuit phase winding does not modify the effect of the fault on the phase current, current drawn from the DC link or on the torque output. Thus the effect of the fault together with the chosen post fault control, remains the same as shown in Figure 45 of Chapter 5.

9.2.2 Turn to Turn Fault

Once a turn to turn fault has been detected the action required to enable continued operation of the healthy phases will vary according to the machine type. For the fault tolerant synchronous PM machine studied here, the entire winding containing the fault must be short circuited at the machine terminals. The strong mutual coupling between individual turns will help to ensure that the resulting MMF is shared equally between the turns. The result will be that each turn in the phase carries the terminal short circuit current.

The requirement to apply a short circuit to the terminals of the faulted phase may seem onerous but in practice the requirement can be met without any change to the drive topology. A terminal short circuit is easily applied through the single phase bridge which normally drives this phase. The upper two transistors should be turned on and the lower two turned off as shown in Figure 93 (or vice versa). The resulting bi-directional current path is shown in bold.

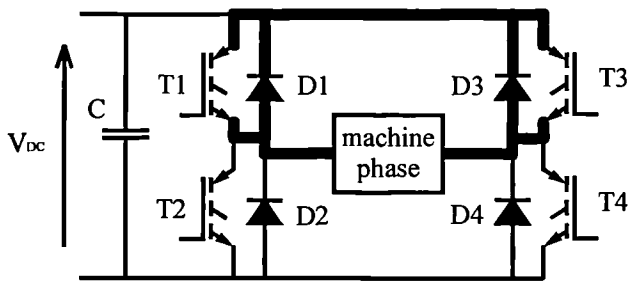

Figure 93 : Short Circuit Applied to the Faulted Phase by the Inverter

The plot in Figure 94 shows the currents measured in the first turn of the modified phase and in the remainder of the modified phase, with the demonstrator machine rotating at 3500 rpm, i.e. at 27% of full speed. Initially the additional turn is shorted but the remainder of the phase is left open circuit. The single turn carries a very large current which is driven by the proportion of the back-EMF appearing across the turn but limited to about 140 A rms by the impedance of the turn. During this time the remainder of the phase carries no current. After a short period a short circuit is applied to the terminals of the faulted phase. Thereafter both the shorted turn and the remainder of the phase each carry a current of 19 A rms, a little less than rated current for this machine.

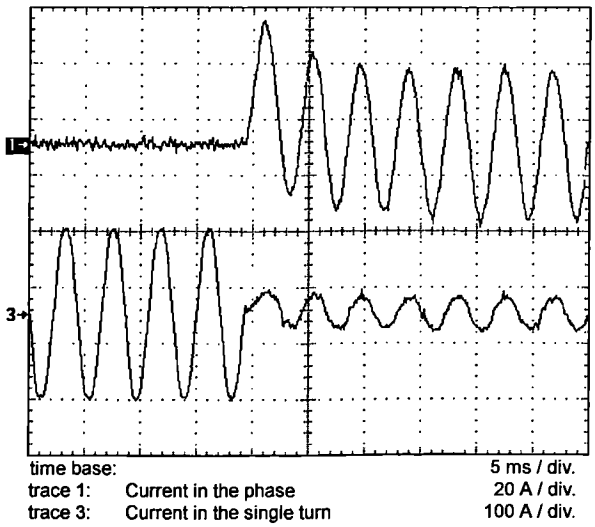

Figure 94 : The Effect of Applying a Terminal Short Circuit on the Shorted Turn Current

The oscillogram shows that with the phase shorted at the terminals by the power electronics the situation is returned to that of a terminal short circuit. The drive can operate indefinitely in this state with the consequences for torque output described in Section 5.6.

Continued Operation of an Inverter Fed Machine with a Shorted Turn

Many of the effects of a shorted turn on an inverter fed drive can be seen in Figure 95. The drive has been operated for a short period with a shorted turn in one phase. Then a short circuit was applied to the terminals of the faulted phase through the single phase bridge which was previously driving the current. Trace 3 (red) is the current in the adjacent healthy phase. This is unaffected by the fault and by the subsequent terminal short circuit applied.

Trace 1 (black) shows the current in the faulted phase. In the first half of the plot the current driven into the faulted phase clearly includes increased level of ripple at the PWM frequency. The oscillation seen in the sampled values shown in Figure 72 earlier is also clearly present here. Trace 2 (green) shows the current driven in the shorted turn as a result of voltages induced by the passing magnet flux and by the changing flux resulting from volts applied at the terminals of the faulted phase.

The oscillogram in Figure 95 shows the drive operating for tens of PWM cycles in order to illustrate the effect of a shorted turn on the current in the affected phase. The 'phase

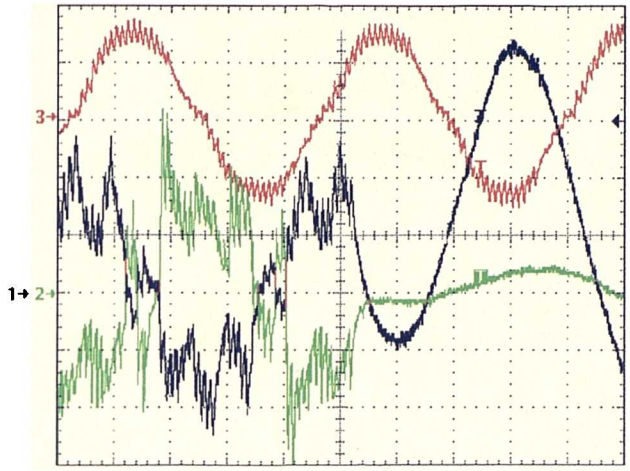

time base:		500 µs / div.
trace 1:	Current in faulted phase (black)	10 A / div.
trace 2:	Current in shorted turn (green)	100 A / div.
trace 3:	Current in adjacent healthy phase (red)	10 A / div.

Figure 95 : Effects of a Shorted Turn before and after the Faulted Phase is Shorted

to model comparison' based fault detection scheme, proposed in Chapter 7, is able to detect a shorted turn at the sampling instant that the phase current first exceeds the predicted value. This will typically lead to detection of the fault and reconfiguration of the power electronics to short the phase within one sampling interval.

The armature inductance of this machine is high at 1.0 per unit. However this is dominated by cross slot leakage flux, so that the armature driven air gap flux is small, both in normal operation and with a short circuit at the phase terminals. The result is that the air gap flux distribution hardly changes when a terminal short circuit is applied. Thus little unbalanced magnetic pull results.

After the phase is shorted through the power converter, both the shorted turn current and the current in the remainder of the faulted phase settle out to the short circuit value. The consequences for the DC link voltage and torque production of the resulting phase short circuit, applied through the power converter, were shown in Chapter 5.

9.2.3 Phase Terminal Short Circuit

The machine is capable of continued operation in the presence of a phase terminal short circuit. However, the fault presents itself to the power converter as an inverter output short circuit. If the inverter continues to switch into the fault, the occurrence of switching states with a diagonally opposite pair of IGBTs turned on will deliver shoot-

through current into the fault. This fault current is detected by either freewheel current sensing or by on-state voltage sensing and thereafter all the switching devices in the inverter supplying the faulted motor phase can be turned off. With all switching devices in the phase turned off, the DC link current supplied to this phase will be zero and the net torque resulting from the faulted phase will be small as described in Chapter 5.

Whilst this is one possible solution, a problem arises with a turn to turn fault which results in a large proportion of the phase winding becoming a shorted turn group. This will appear to the power converter as a phase terminal short circuit. A turn to turn fault requires that a short circuit be applied at the machine terminals by turning on the two upper IGBTs in the inverter supplying the faulted winding, as described in a previous section. This action is also appropriate for the phase terminal short circuit fault and is therefore the preferred action in the event of any phase terminal short circuit indication.

9.2.4 Open Circuit IGBT

In Chapter 5 the effect of an open circuit IGBT was portrayed as a benign failure which results in the reduction of current drive capability from bi-directional to unidirectional. Thus, through the continued operation of the healthy diagonal pair of IGBTs as a unipolar bridge, the phase is still capable of delivering half rated torque. The oscillogram in Figure 96 shows the phase current and the current supplied to the phase from the DC link through the introduction of an open circuit fault in the reverse lower IGBT.

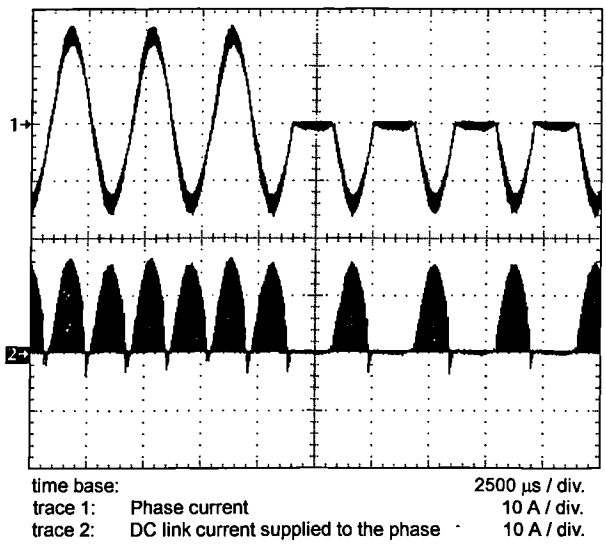

Figure 96 : Operation throughout the Introduction of an Undetected Open Circuit IGBT

After the introduction of the fault, control of the phase continues unaltered and negative half cycles of phase current are driven in the phase winding. Trace 2 shows the current supplied to the phase from the DC link. The current is actually pulse width modulated, with pulses spaced at 50 μs intervals, but this detail is lost on the time base used, so only the envelope of the current is seen. The sinusoidal current in a healthy phase requires that a series of half sine currents be drawn from the link. At the end of each half sine section a small amount of current is returned to the link. This current corresponds to the return of inductive energy. After the fault is introduced the current supplied by the link is unaltered except that alternate half-sine sections are absent.

When the fault was introduced, freewheel current sensing was the only fault detection scheme operating. It is interesting to note that because the fault was introduced during a negative half cycle of phase current, whilst the affected device was not conducting, it remained undetected by this detection scheme. Furthermore, the fault results in the loss of forward current capability making subsequent detection impossible.

Even without fault detection the resulting operation of the phase is acceptable. In Chapter 5 it was shown that a phase operated in this manner delivers half the torque of a healthy phase. However, it is worth considering what action would give the best possible performance of a drive with an open circuit IGBT fault. Whilst the continued operation of the phase with unipolar currents arguably results in the minimum reduction in torque output of the drive, there is also the possibility of ceasing operation of the faulted phase completely. The course of action taken will depend on the degree of redundancy in the drive and this will have been determined by the demands of the application.

The aircraft fuel pump application is regarded as safety critical. The degree of redundancy present is sufficient to compensate for the total loss of output from one phase. Therefore it is possible to cease operation of the faulted phase without an overall loss of drive output. This is the preferred option since it minimises stress on the remaining healthy devices in the faulted phase and thus minimises the likelihood of a further fault in the faulted phase.

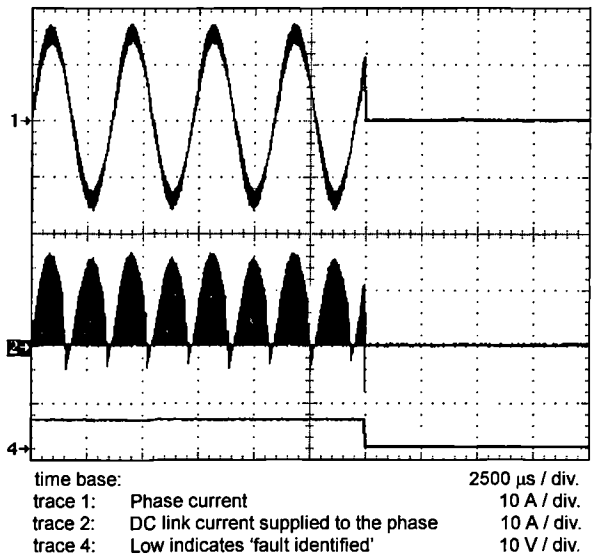

Figure 97 : Operation throughout the Introduction of a Detected Open Circuit IGBT

This preferred method of drive control, after the introduction of a reverse lower IGBT open circuit fault, is shown in Figure 97. This time the fault is introduced during a positive half cycle of phase current. The fault prevents the phase current from circulating within the bridge during alternate freewheeling periods in the PWM cycle. Thus the phase current is forced to flow back to the DC link during these periods. The current returned to the DC link is driven rapidly to zero by the full DC link voltage. The resulting pulse of current returned to the DC link is just visible in the oscillogram.

Freewheel current sensing is in operation and, because the fault is introduced whilst the faulted device is carrying current, the fault is detected by the sensing scheme in the brief period of current return to the DC link. Immediately upon detection of the fault, all the switching devices in the faulted bridge are driven off. The effect of this action on drive output, DC link and so on, is identical to the effect of a open circuit phase winding.

9.2.5 Power Device Short Circuit

In Chapter 5 it was shown that the healthy device in the faulted limb must be turned off to prevent uncontrolled shoot-through currents from flowing. It was also shown that if the two devices in the healthy limb are turned off then a pulsating unipolar current will flow with an RMS level in excess of rated current. Figure 98 shows the effect of turning off all the devices in the faulted bridge on both the phase current and the current supplied to the phase by the DC link.

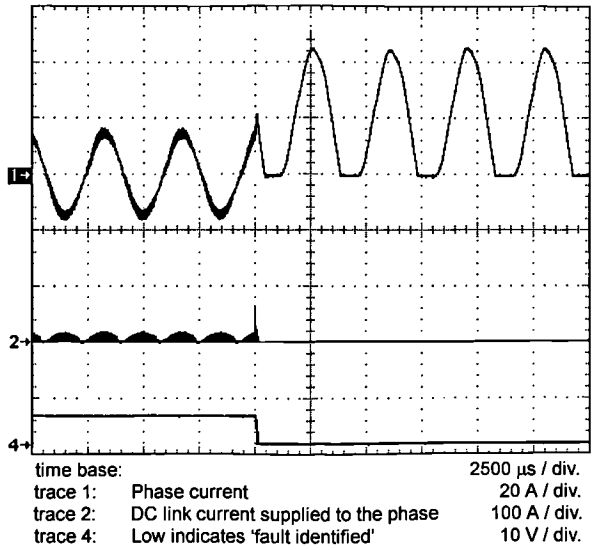

Figure 98 : The Effect of Turning Off all Healthy Devices after the Introduction of a Reverse Lower IGBT Short Circuit Fault

A solution to the problem of more than rated current flowing in the faulted phase is to turn on the device in the healthy leg adjacent to the shorted device in the faulted leg. This results in the machine winding being completely shorted out. This action is the same as the response to a shorted turn except that with the shorted turn it was of no importance whether the two upper devices or the two lower devices were turned on in response to the fault. Here it is essential to turn on the pair of devices which include the faulted device. For example, if one of the lower devices fails short circuit then the appropriate response is to turn on both lower devices. The situation is now equivalent to a terminal short circuit in the machine except that the phase short circuit current also flows through the inverter. Again the drive is capable of indefinite operation in this state.

Figure 99 shows phase current, DC link current and output from the freewheel current sensor during the introduction of a reverse lower IGBT short circuit. The figure comprises three oscillograms, each showing the same event on a different time base. The fault was introduced when the reverse lower device was conducting. When the reverse upper device subsequently turns on into the fault, the DC link is shorted through the reverse pair of IGBTs and shoot-through current flows. It is clear from Figure 99(b) and Figure 99(c) that the shoot-through current exceeds 200 A and that this current level is an order of magnitude higher than the current taken from the DC link in normal operation. It is only possible to prevent this fault from having a serious impact on the

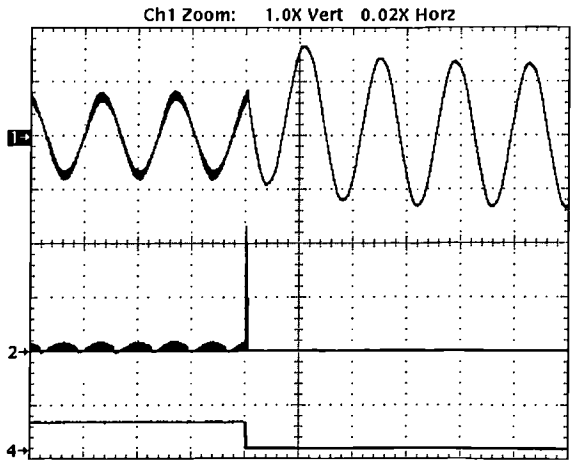

(a) Over a few machine cycles

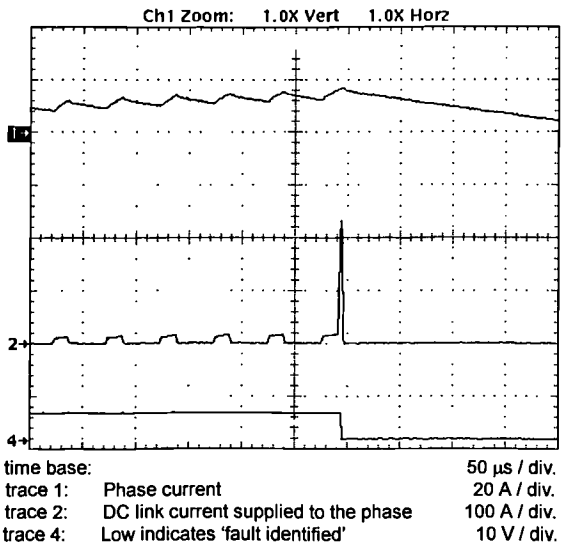

(b) Over a few PWM cycles

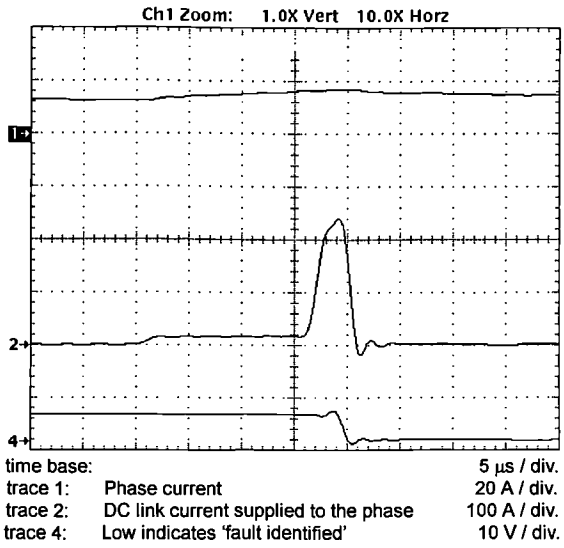

(c) The fault current in detail

Figure 99 : The Effect of Introducing a Reverse Lower IGBT Short Circuit Fault

operation of the drive by limiting the duration of the fault current to a few micro seconds. This fault is detected and identified as a reverse lower short circuit by the freewheel current based detection scheme. The fault detection circuitry drives all the healthy devices off, limiting the duration of shoot-through fault to 3.5 µs. After a further 16 µs the two lower devices are driven on, imposing a phase terminal short circuit. Thereafter short circuit current flows. The effects of this action on drive output, DC link and so on, were presented in the phase terminal short circuit section of Chapter 5.

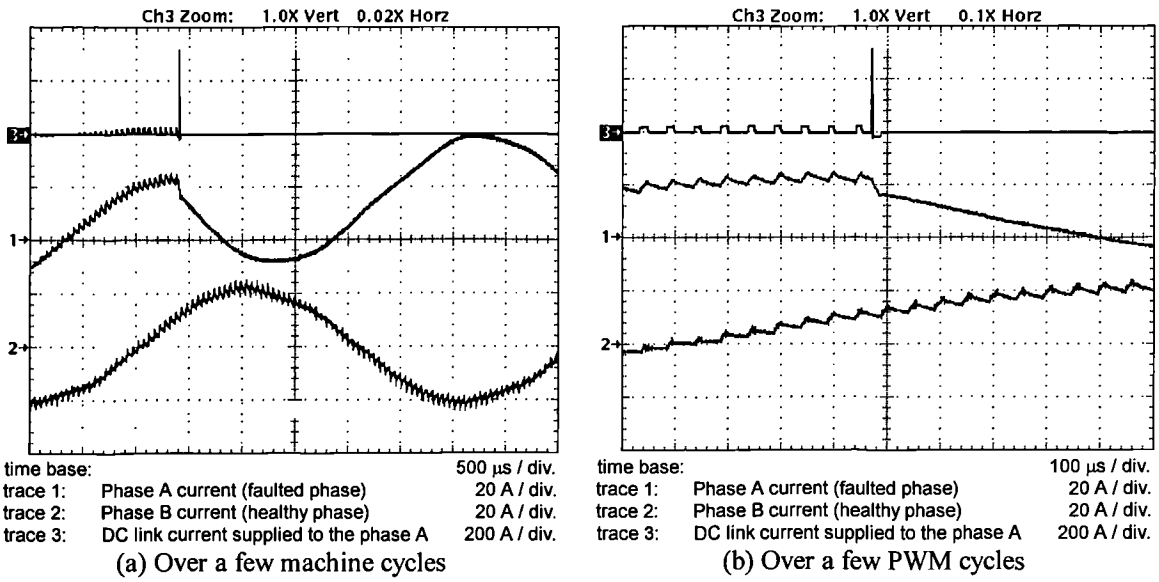

(a) Over a few machine cycles (b) Over a few PWM cycles

Figure 100 : Operation throughout the Introduction of a Reverse Lower IGBT Short Circuit Fault

Figure 100 shows both the phase current and the current supplied to the faulted phase from the DC link during the introduction of a reverse upper IGBT short circuit. The figure comprises two oscillograms, each showing the same event on a different time base. The fault was introduced when the reverse lower device was conducting. In this case the fault detection was performed using on-state voltage monitoring but the post fault control strategy is identical whatever the detection method. The oscillograms demonstrate that the current in the adjacent phase is completely unaffected, even at the instant that the faulted phase draws the pulse of shoot-through current from the DC link, in spite of the phases sharing a common DC link. The fault is detected quickly and the shoot-through current flows for only 3 µs. The short duration of the fault current allows the DC link capacitance to effectively de-couple the phases.

9.2.6 Phase Current Sensor Failure

In the event of phase current sensor failure, the preferred action would be to discontinue operation of the faulted phase. However the actions taken in the event of each of the faults described in this chapter will also be acceptable in the event of a phase current sensor failure. Thus it is not important to be able to distinguish phase current sensor failure from other faults.

Consider the fault detection scheme based on 'phase to model' comparison, a current sensor failure cannot go undetected. If the sensor output remains at zero, for example, this will appear to the fault detection algorithm as an open circuit phase. Whilst the fault will be erroneously reported as an open circuit phase, the action taken, to turn off all power devices in the phase, will prevent the controller from driving excessive current into the phase.

In general the action to be taken in the event of a suspected fault will be tolerable even if that phase were actually fault-free. Thus a sensor failure erroneously detected as another fault will not be a problem, in fact it is much better for a sensor failure to be detected as another fault than for it to go undetected.

9.3 The Thermal Situation

The preceding discussion has shown that appropriate post fault control is available to limit fault current to rated value for all faults. Thus, once the post fault control is applied, the drive is capable of operating indefinitely without excessive temperature rise.

During the fault transient, excessive currents will flow which will lead to an increase in the temperature of the drive. The thermal time constant of the machine is approximately 2 minutes. This short time constant is a result of intensive cooling. The time constant of the power converter would be similar if it were also liquid cooled. All faults can be detected within one machine cycle, which is several orders of magnitude faster than the thermal time constant for any reasonable operating speed. Post fault control is applied immediately after fault detection and the subsequent transient has a time constant of approximately 5 ms, as shown in Figure 94. The speed of fault detection and the rapid decay of the fault current prevent the drive temperature from rising significantly before the post fault control becomes effective.

Certain faults, notably shoot-through faults and shorted turn faults, result in fault currents which exceed rated current by more than an order of magnitude. The situation is made even more serious because in each case the power is dissipated in a very small part of the drive. For these faults it is not sufficient to say that the overall drive temperature will not increase significantly. The issue of localised heating must be considered. The rate of temperature rise for the shorted turn fault was calculated in Chapter 5 to be 1545°C/s.

'Phase to model comparison' is capable of detecting this fault in one PWM cycle, i.e. in 100 μs, which should result in a localised temperature rise of less than 1.0 °C. The rate of temperature rise for an IGBT die carrying shoot-through current is not known, but these devices are designed to withstand short circuit conditions for 10 μs. The fault detection schemes based on the freewheel current sensor and on the on-state voltage monitor can detect and turn off the shoot-through current in 3.5 μs and 3.0 μs respectively, comfortably less than the short circuit withstand time.

In the case of each fault, the temperature rise has been limited by detecting the fault in a timely manner and adopting an appropriate post fault control strategy immediately after fault identification.

9.4 Torque Recovery

Until now much of the emphasis has been on the faulted phase. All the actions taken after detecting a fault have been designed to prevent fault propagation. Once a fault has been detected and the appropriate post fault control applied the drive can be operated in the faulted state indefinitely. However it is not enough for the drive to continue to operate in the presence of a fault. To be of any use the drive must continue to produce useful torque.

The fault tolerant demonstrator drive controls the level of current and thus torque in response to a current reference. There is no speed loop, so drive speed is dependent only on the current reference and the load torque. In the event of a fault, the power output of the drive for a given current reference will reduce as a result of lost torque contribution from the faulted phase. In the demonstrator drive, the reduction in power output will result in a reduction in shaft speed. The oscillograms in Figure 101 show a visible reduction in shaft speed resulting from the response of the drive to the introduction of a power device short circuit fault.

There are two ways in which the control can be designed to make up for the loss of torque which results from a fault. If the fault is detected and diagnosed quickly then knowledge of the fault can be used to calculate the torque loss. Then the torque demand applied to the remaining phases can be factored up by an appropriate amount.

ORIGINAL IN COLOUR

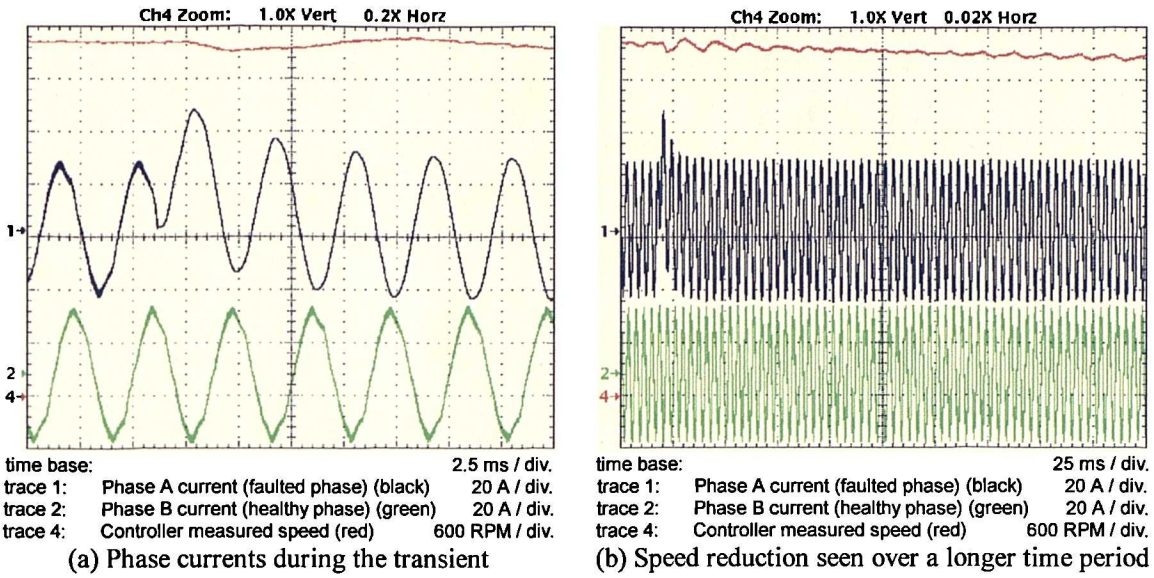

(a) Phase currents during the transient (b) Speed reduction seen over a longer time period

Figure 101 : Plot of Speed before and after Short Circuit Power Device Fault

Even in a drive where precise fault detection is not available, the use of a speed control loop will naturally compensate for any loss of torque resulting from a fault. Even a simple PI speed controller will respond to a fault which results in a loss of torque and therefore speed by increasing the current/torque demand to all phases until the actual speed is returned to the speed set point. The drawback of relying on the speed loop is that the speed must fall before the torque demand can be increased to make up for the torque loss due to the fault. However, it has the advantage that, in the long term, precise drive power control is achieved without requiring knowledge of the fault or corresponding torque loss. Torque recovery of this type is demonstrated by Lui [102].

The best performance is likely to be achieved with a combination of torque demand adjustment based on fault diagnosis and speed loop regulation. The first (feed forward) adjustment will adjust the post fault torque to approximately the pre-fault level within a PWM cycle. After this the (feed back) speed controller will make the fine adjustments to bring the post fault torque to precisely the pre-fault level.

To demonstrate the effect of torque recovery, Figure 102 shows the drive continuing to operate during and after the introduction of a power device short circuit. When the fault is identified appropriate post fault control is applied to the faulted phase. At the same time the current demand to each of the remaining healthy phases is increased by a factor

ORIGINAL IN COLOUR

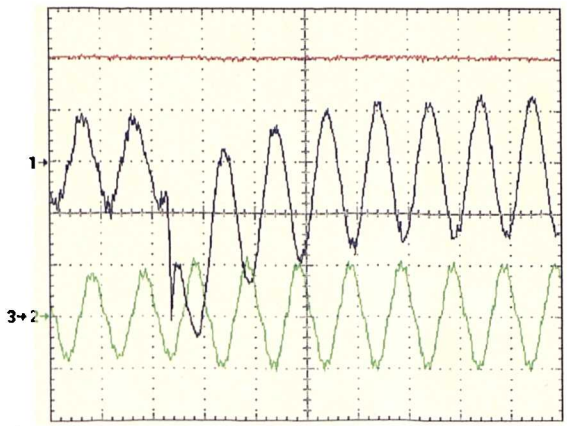

(a) Phase currents during the transient

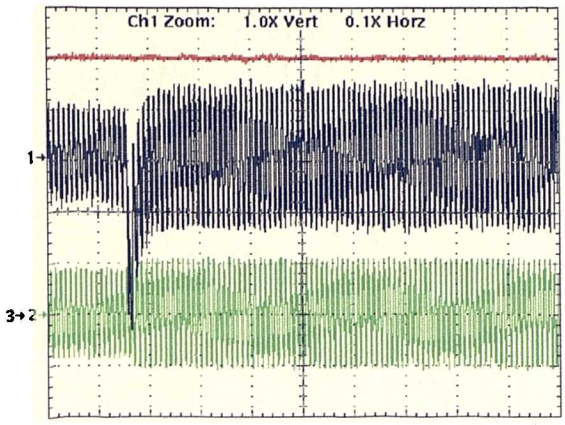

(b) Speed reduction seen over a longer time period

Figure 102 : Speed before and after a Fault with Lost Torque Compensation

of 6/5. It can be seen that the current in the healthy phase is increased by 20% after the introduction of the fault. The increased torque from the healthy phases makes up for the loss of torque from the faulted phase. The increase in current in the healthy phases is achieved within 200 μs of the introduction of this fault, which is easily fast enough to prevent any significant disturbance to the machine speed. At this operating point, simply factoring the current demand ensures that there is very little loss of drive speed in spite of the drive having no speed loop. Note that the inertia of the machine coupled to the test rig is larger than the inertia that would result if the machine were coupled to a fuel pump. Thus the degree of speed disturbance seen in these tests is not necessarily representative of the speed disturbance that would be seen in the fuel pump application.

9.5 Torque Smoothing

By designing the drive with n+1 redundancy it is possible to deliver rated torque from the drive even when the torque producing capability of one phase is lost completely. However, whilst the faulted drive continues to deliver the required amount of torque the quality of the torque delivered is reduced, as shown in Chapter 5. For the aircraft fuel pump application, the increased level of torque ripple resulting from even the most severe faults is not thought to be a problem. This is especially true, given that the drive will be required to operate for only a few hours in the faulted state.

For applications where smooth torque production is essential, a different approach is required. It has been shown that the time averaged torque loss due to a fault can be recovered by factoring up the current demand in remaining healthy phases. If smooth torque were required then this could be achieved by factoring the instantaneous current demand in each of the healthy phases to make up for the faulted phase. This approach was presented by Jack et al [35] who show that the current in the healthy phases of a 3 phase PM synchronous drive can be shaped to partially compensate for the torque ripple introduced by a terminal short circuit. The paper implies that there is a trade off, in the faulted state, between smooth torque and total torque capability.

The requirement for smooth torque has serious consequences for the level of redundancy required in the drive. Consider a drive which is required to deliver rated torque at stall. The peak current driven in the faulted drive must not exceed rated current for the drive. If a phase fails open circuit then the loss to mean torque in this drive is 16.7%. The instantaneous loss of torque is greatest during that part of the electrical cycle when the faulted phase should have been delivering peak torque. At this point the instantaneous torque producing capability of the drive is reduced by 33.3%.

In the event of a short circuit fault the situation is worse, as the faulted phase contributes a pulsating torque. The loss of healthy torque contribution for this phase, combined with the pulsating torque contributed by the fault, results in a peak reduction in instantaneous torque of approximately 50%. Furthermore, the current waveforms required to produce smooth torque from the faulted machine demand faster changes in current than those occurring in normal operation and may only be achievable with a disproportionate increase in power converter VA rating.

9.6 Conclusion

A post fault control strategy has been proposed for each fault. In each case post fault control has been developed to minimise the impact of the fault on the DC link and the torque output of the drive. The thermal situation has been considered, both in terms of the transient and steady state temperature rises. The issue of localised heating has also been addressed for the most serious faults. In the case of each fault, the temperature rise

has been limited by detecting the fault in a timely manner and by applying an appropriate post fault control strategy immediately after fault identification.

The flux in an operating machine will normally drive many times rated current in a shorted turn fault. This chapter has demonstrated that it is possible to limit the current in the shorted turn to rated value, simply by using the power converter to impose a short circuit on the phase terminals. A combination of rapid fault detection, described in Chapter 7, and the technique for limiting current in the shorted turn ensure that the shorted turn fault does not result in an excessive temperature rise. The technique for limiting current in the shorted turn allows the drive to continue to operate indefinitely with a shorted turn fault.

Shorted power devices and phase terminal short circuit faults result in shoot-through currents of the order of 10 times rated current. The fault current can be detected and turned off within the 10 μs short circuit withstand time of the power devices. The fault current does not flow in the machine phase and this can be seen from the phase current in the faulted phase. During the transient, the phase current can reach a peak value of twice rated current, but this settles within 10 ms - so no significant over heating will occur. Operation of the adjacent healthy phases remains undisturbed throughout the periods of fault introduction, detection and post fault control.

Some consideration has been given to the control of the remaining healthy phases after a fault is detected. Operation of the healthy phases for torque recovery and torque smoothing has been discussed. A technique for adjusting the torque produced by the healthy phases, to make up for the torque loss resulting from a fault, has been demonstrated. This technique accepts that the torque ripple increases as a result of the fault and simply ensures that the average torque delivered after a fault remains the same as the average torque produced before the fault. It would be possible to continue to produce smooth torque, even in the presence of a fault by applying instantaneous torque control. However the production of smooth torque in the faulted state demands a higher level of redundancy. Torque recovery without torque smoothing is well suited to the aircraft fuel pump application as the torque ripple introduced by the most severe fault is still within acceptable limits.

With fault detection strategies from Chapter 7 and Chapter 8 and an appropriate post fault control strategy it is possible to operate the drive during and after the introduction of any fault. The introduction of a fault, and the application of the chosen post fault control, will result in very little disturbance to the overall drive function. The proposed approach to recover the lost torque ensures that the torque output of the drive is not altered significantly either by the fault transient or in subsequent faulted operation. The post fault control strategies ensure that the drive will operate indefinitely, in the presence of any fault, without resulting in an excessive temperature rise.

Chapter 10

CONCLUSION

10.1 Fault Tolerant Permanent Magnet Drives

The principle aim of the work presented here has been to design a drive capable of surviving any single point failure, detecting the failure and continuing to deliver full rated torque. Every effort has been made to achieve this goal with the minimum penalty in terms of increased complexity, parts count, cost, weight and size. Now it is time to assess the extent to which these objectives have been achieved and to review the trade-offs which have been made.

The scope of the work has been limited to electrical failures in the electric machine and power converter. Failure of the rotor position sensor and failures in the control electronics were not considered.

The requirements of fault tolerance have been adopted from the power systems field and applied to a brushless permanent magnet drive. First the drive function must be partitioned correctly. Then it is possible to achieve the necessary level of redundancy without seriously increasing the cost or complexity of the drive. The drive design presented has 5+1 redundancy, which only results in a 20% increase in drive mass.

A novel design of permanent magnet machine has been developed for use in this fault tolerant drive. The machine topology provides effective isolation between machine phases. The short circuit current in the machine is limited to 1.0 per unit by the phase inductance. The machine design presented shows that the SRM is not the only machine suitable for applications demanding fault tolerance.

Each phase of the machine is driven by an independent single phase bridge, so the isolation between phases is maintained in the power electronic converter. The resulting drive is capable of continuous operation in the absence of any one phase. These qualities

are shared by the ship propulsion drive reported on by Olivera et al, [18]. However the drive here has several important advantages:

- The machine has a 1.0 per unit inductance, limiting fault current to rated value.
- The machine has effective physical, electrical, magnetic and thermal isolation between phases.
- The power converter uses switching devices which are capable of turning off the short circuit current.
- No fuses are required - faults are controlled by the remaining healthy devices in the faulted phase rather than simply being isolated.

In spite of the unusual design, the machine delivers a high specific power, only about 15% less than a conventional design of PM machine but still providing around twice the specific power of a similarly sized SRM. The power converter VA rating required to supply the machine must be increased by approximately 34% to handle the increased levels of reactive power resulting from the high machine inductance. The resulting power converter requirement of 10.4 kVA/kW compares favourably with a power converter requirement of approximately 10.5 kVA/kW for a typical switched reluctance drive, Miller, [65].

10.2 The Fault Tolerant Demonstrator Drive

A demonstrator drive has been constructed. The performance of the resulting drive has been compared to the original pump specification. The drive performance is thermally limited due to higher than expected thermal resistance from the winding to the case and the higher than expected iron losses. These shortcomings are associated with the efforts to achieve a high power density and do not result from those features in the machine essential to fault tolerance. The electro-magnetic performance of the machine is as expected and the drive proved capable of delivering the rated torque. The electro-magnetic coupling between phases has been measured as 2.2%, demonstrating that the aim of effective magnetic isolation between phases has been achieved.

The drive has been operated as a generator with all phases short circuit. The resulting temperature rise does not exceed that which results from fault-free motoring operation,

demonstrating that machine terminal short circuits do not present a thermal threat. Moreover, the machine has good thermal isolation between phases, with only 5.3% thermal coupling between phases, so even if a fault were to result in excessive power dissipation this would not have a large impact on the adjoining phases.

The power converter VA requirement for a given drive power, can be reduced by phase advancing the current with respect to the back-EMF. This is in contrast to a PM machine with low phase inductance and no saliency where the best operating mode is with the current in phase with the back-EMF. The pump drive requires full torque at maximum speed, however flux weakening is still useful in reducing the VA requirement of the power converter. The key attribute of this drive which makes flux weakening possible, and indeed advantageous, is the high per unit inductance.

The power output for a power converter operating from a given DC link is traditionally maximised by switching from current control to voltage control. Here the same effect has been achieved by continuing operation in current control but using a current profile designed to minimise the VA requirement. A technique has been presented for designing optimum current profiles from the ψ-i diagram. The technique gives more torque and a lower kVA/kW requirement at a given peak current. For a given RMS current, the resulting current profile gave 4.0% more torque than sinusoidal excitation with 30 degrees phase advance, and required 10.0% less kVA.

The result is a high performance permanent magnet drive with a similar degree of inherent fault tolerance to that provided by a switched reluctance drive. The drive is capable of delivering full rated torque in the presence of any one of the single point failures considered.

10.3 Model Based Current Control

A current control scheme has been presented, based on the measured flux linkage characteristic of the machine. This current controller was compared to a PI current controller on the basis of step responses and high speed tests.

With a simple PI current controller it was not possible to achieve the gain necessary to drive the current waveform desired with the sampling rate used. Modifying the PI controller to add a back-EMF feed-forward term gave a large improvement in the current

control performance at high speed. However, even with this improvement, the PI controller gave very poor current reference tracking as the machine phase approached saturation.

The model based current controller replaces the PI controller with a simple model of the plant. Thus the model based current controller requires no tuning and there is no trade off between stability and transient response. It has been demonstrated that the transient response of the model based current controller is exceptional. Some steady state error may be present due to voltage drops in the power devices and phase winding and due to the deadtime introduced by the PWM unit. The size of the steady state error can easily be estimated. It has been demonstrated that, by modelling the sources of this error in the controller, the steady state error can be reduced towards zero. The resulting controller will always drive the current to the level demanded, subject to there being sufficient volts available to do so. This is true even when the phase is driven into saturation.

The actual current which will be achieved at each sampling instant can be accurately predicted. This is true even in the situation where there are insufficient volts available to achieve the demanded level of current. This has important implications for fault detection. Any difference between the current measured and the current which the controller had expected to achieve indicates that the phase behaviour has deviated from the model. This can only indicate some fault in the phase. The way in which the measured current deviates from the expected current may even be used to diagnose certain faults.

10.4 Operating a Faulted Drive

A short list of the most likely faults has been presented. It has been demonstrated that the drive is capable of operating indefinitely with a power device or a phase open circuit or even with a phase terminal short circuit.

A single shorted turn in the machine or an undetected power device short circuit failure results in very large fault currents and correspondingly serious consequences. Without rapid detection and appropriate post fault control, these faults will propagate rapidly and/or disturb the operation of healthy phases. It is often argued that these are also the most likely faults.

It has been shown that, for a synchronous PM machine with a reactance of 1.0 per unit, a terminal short circuit should be applied to the faulted phase as soon as a turn to turn fault is detected. Since the demonstrator machine has little mutual coupling between phases it can continue to operate in this state indefinitely.

Shorted power devices and phase terminal short circuit faults result in shoot-through currents of the order of 10 times rated current. The fault current can be detected within the short circuit withstand time of the power devices and turned off. Whilst the shoot-through current does not flow in the machine phase, a shorted power device allows machine phase currents in excess of rated current to circulate, driven by the back-EMF. Turning on the adjacent power device in the healthy limb reduces the currents which circulate through the faulted power device to rated value thus permitting continued operation of the drive. The transient phase current can reach a peak of twice the rated value but this settles within 10 ms - so no significant overheating will occur. Operation of the adjacent healthy phases remains undisturbed throughout the periods of fault introduction, detection and post fault control.

The level of redundancy required in the drive was calculated at the design stage, assuming that the faulted phase would not contribute any significant braking torque. Some of the faults do result in current flowing in the phase and thus in losses. Whilst these losses are small, experimental measurements show that they are significant. The losses present in the faulted phase are highest for a phase terminal short circuit applied through the inverter. This fault results in a torque reduction of 21.3% at approximately one third of full speed, rather than 16.7% for a loss-less drive. Thus the drive must be designed with 27 % redundancy (rather than the 20 % redundancy required in a loss-less drive) if the drive is to continue to deliver rated power in the event of this condition.

It is important to remember that the torque reduction was measured at one third of rated speed and that the lost torque resulting from the short circuit faults will reduce for higher speeds. Thus 27 % redundancy is a conservative figure.

It has been demonstrated that, if appropriate action is taken, the drive can continue to operate indefinitely in the presence of any one of the faults considered. The introduction of a fault, and the subsequent application of post fault control, has no significant effect on the remaining healthy phases. It has been demonstrated that, with fault detection and

appropriate post fault control, the introduction of a fault will result in very little disturbance to the overall drive function.

If smooth torque were required from the faulted drive this would also be possible but at the penalty of reduced total torque or increased redundancy.

10.5 Fault Detection

Fault detection is a vital element in any fault tolerant system, as without fault detection a build up of failures will eventually result in a loss of system function. In an electric drive, shorted turns and power device short circuit failures result in fault currents of more than 10 times rated current, making rapid fault detection especially important.

To enable fault detection it will always be necessary to measure some quantities in the drive. Thus the development of fault sensors is a key area of interest. A range of sensors for fault detection was examined. These included sensors already present on many drives, such as the phase current sensor, as well as other sensors developed especially for fault detection. It should be noted that the ideal situation is to be able to detect all possible drive faults using only those sensors normally present in a drive. As more sensors are added, the level of redundancy in the sensing scheme increases, but integration of the extra sensor information becomes more difficult.

Phase current monitoring can be used to detect a number of faults including open circuit faults in the switching devices or in the motor phase and turn to turn short circuits in the motor phase. It has been shown that the detection scheme is sufficiently sensitive to detect a single shorted turn in the machine phase. This is particularly important as the turn to turn fault is widely regarded as a major cause of electric machine failure. The detection scheme is able to detect the failure of the phase current sensor itself. Thus this scheme is self checking.

On-state voltage sensing is already used in many gate drive circuits to protect power devices from damage resulting from excessive current. Information from on-state voltage monitoring has been used here to detect a range of short circuit faults. It has been shown that it is a simple matter to adapt the sensor and detection scheme to detect open circuit IGBT faults and to detect failure of the on-state voltage monitor itself.

If current sensing IGBTs are introduced successfully and widely adopted then it may be possible to eliminate the phase current sensor and the on-state voltage monitors from most drives. The current sensing IGBTs would give as many current signals as there were IGBTs carrying current, thus providing redundant current sensing. This situation would require its own special detection and diagnostic scheme but would naturally borrow many of the techniques from the phase current and on-state voltage monitoring schemes presented here.

A freewheel current sensor has been integrated into a power hybrid. Operation of the sensor has been analysed. The ability of the sensor to detect open circuit and short circuit faults has been demonstrated. However the detection scheme does require an additional sensor, albeit an unsophisticated one, and is not inherently capable of detecting all the failure modes of the sensor itself.

In practice all drive faults can be detected using either phase current sensing or on-state voltage monitoring. Given that many drives already include a phase current sensor and on-state voltage sensors, the original goal of detecting all possible drive faults without incorporating any additional sensors has been achieved.

One of the questions posed was 'Is it possible to detect these faults fast enough and can appropriate action be taken in the presence of these faults to allow continued operation of the drive, without significant disturbance to the supply or to the load?'. It is now clear that all the faults considered can be detected sufficiently quickly to prevent fault propagation and to ensure that, with appropriate action, the disturbance to the load and supply remains within acceptable limits.

10.6 Other Applications

Whilst the work presented has been strongly focused on a novel permanent magnet drive it is possible to see how some of the same principles could be applied to drives using different machines, or to drives for applications with a different balance between the requirements for cost and availability. For example, it is easy to see that the principles of partitioning and redundancy could also be applied to an SRM. Furthermore, it seems likely that all the possible electrical faults in a switched reluctance drive could be detected by a combination of phase current sensing and on-state voltage monitoring. In

particular the method of shorted turn detection presented for a fault tolerant synchronous PM machine is expected to work equally well for switched reluctance machines, for which a shorted turn is an equally serious fault [25].

10.7 Further Information

Whilst there are several ongoing developments such as direct torque control and current sensing IGBTs which could change the details of the approach, the overall goal of a fault tolerant drive has been achieved, using available and proven technology. The author believes that this is the first time that all of the elements required to realise a fault tolerant drive have been brought together.

More detailed information on certain aspects of this work may be found in other publications by the author. These publications are detailed in the List of References and include: [36], [37], [38], [39], [40], [41] and [103].

10.8 Future Work

This research covered a very wide area and certain issues have been discussed only superficially, or omitted completely, in an effort to limit the scope. Two such issues are:

1. Sensorless shaft position detection for fault tolerant drives.
2. Fault tolerant AC to DC converters for drives connected to an AC supply.

The area of sensorless shaft position detection is a popular topic of research. In the field of fault tolerance, Hofer [62] has used it as a backup in case of sensor failure. In the application examined in this thesis the situation differs considerably, since the sensorless algorithm would have to operate even in the event of a phase failure. A sensible structure might have each individual phase sensing its own position. In any case, there would be strong links between the position sensing scheme, model based current control and fault detection using 'phase to model comparison'.

Even within those areas covered by this research, some issues have not been fully investigated. In particular, the following require further attention:

1. ***Fault mechanisms in the power electronics.*** These may include voltage stress, current stress, thermal stress and thermal cycling, [104].

2. ***Transient behaviour of real faults.*** Do any of the failures appear intermittent for a short time? Detection times are in the region of 3 μs for detection schemes based on the on-state voltage monitor and 100 μs for 'phase to model comparison'. Any intermittence in time frames longer than these could cause problems.

3. ***Threats posed by inductive voltages.*** Some of the possible faults could result in the generation of inductive voltages in the machine. It appears that these are clamped to the DC link by the freewheeling diodes, so do not represent a threat to the healthy power devices. A more detailed examination of the inductive voltages generated by faults in the machine and power converter is desirable to verify whether or not these are a problem.

4. ***Power converters layout requirements.*** More work is required to establish design rules which will minimise the risk of fault propagation in the power converter through thermal and/or physical mechanisms. This work must establish the amount of thermal and physical separation required between power devices.

Intermittent faults have been only briefly considered. For instance, a power device with an open circuit gate drive will sometimes turn on when it should be off and will sometimes turn off when it should be on. This suggests that the monitoring of this condition should be continuous, with the action to be taken decided on an instant by instant basis. Further experimental work is required to validate this approach.

The DC link capacitor has received relatively little attention. A fault tolerant capacitor network has been suggested. Further analysis is required to compare this approach to possible alternatives. If this analysis indicates that the suggested fault tolerant capacitor topology is attractive then a fault tolerant capacitor network should be constructed and tested. This will provide a clear idea of the space and performance penalties of the scheme.

Appendix A

MACHINE DETAILS

This appendix gives mechanical details of the permanent magnet machine, designed for use on the fault tolerant drive demonstrator rig. The information here is intended to add to the description in Chapter 2. Some additional information on the smaller prototype machines is also included.

Note that all drawings have been reduced in size, so where scales are given, these should be disregarded.

Machine Drawings

The machine used on the demonstrator featured a 6 phase stator. The stator core was assembled from a stack of laminations, each with a thickness of 0.35 mm and a cross section as shown in Figure 103.

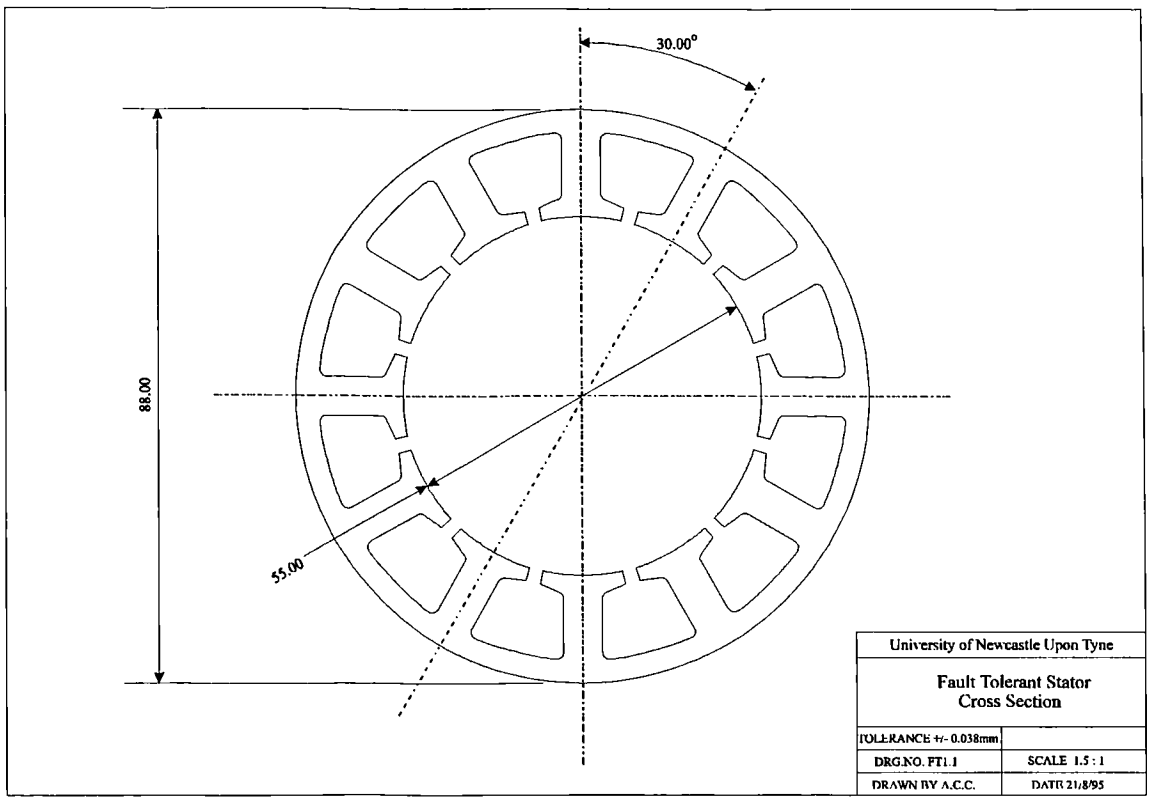

Figure 103 : Stator Lamination Drawing

The machine rotor comprises a mild steel shaft, with an octagonal section onto which eight magnet segments have been bonded. The eight pole magnet assembly is protected by an inconel sleeve with a wall thickness of 1.5 mm. The rotor assembly is shown in Figure 104

The rotor and stator are supported in an aluminium frame as shown in Figure 105. The resulting air gap is 0.5 mm. Note the following features in the assembly:

- The bearings are pre-loaded by thrust washers at the drive end of the machine.
- Shaft position information is provided by a two part brushless resolver, mounted at the non-drive end of the machine.
- The resolver is protected by an aluminium cover.
- A channel exists between the main part of the frame and an additional sleeve. This channel carries the liquid coolant, used to cool the stator.

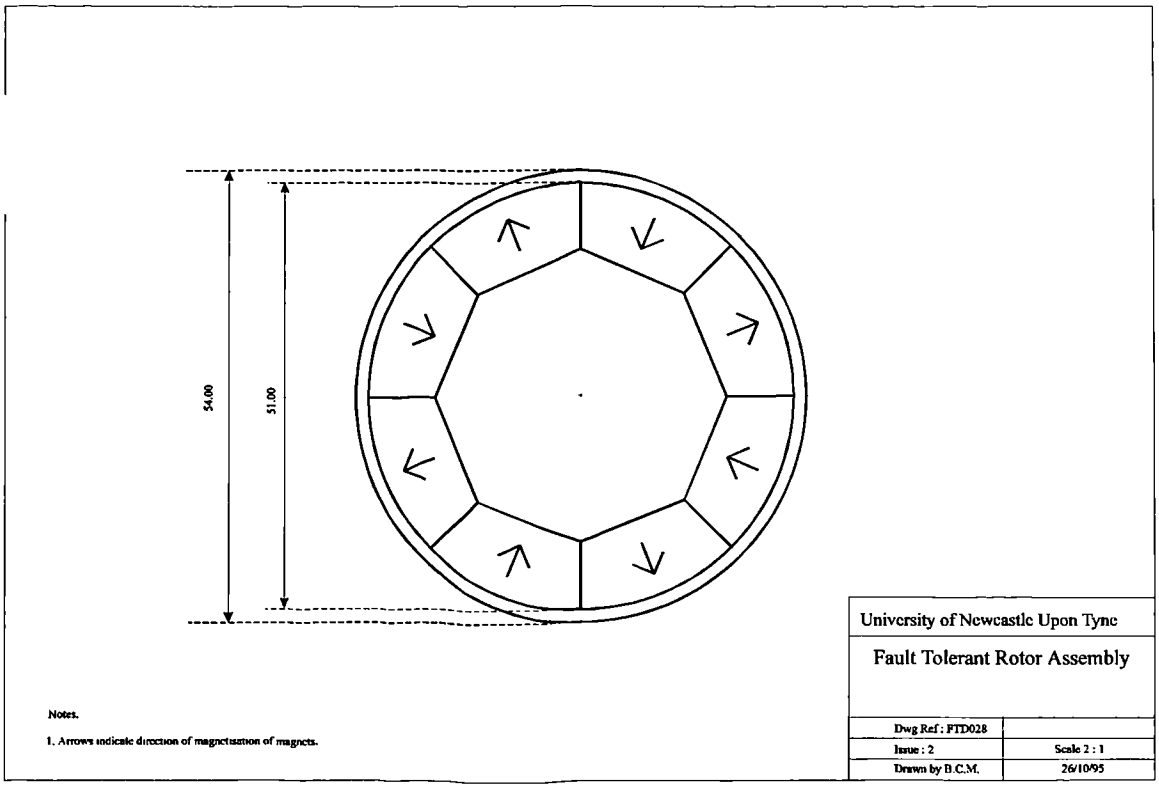

Figure 104 : Rotor Arrangement

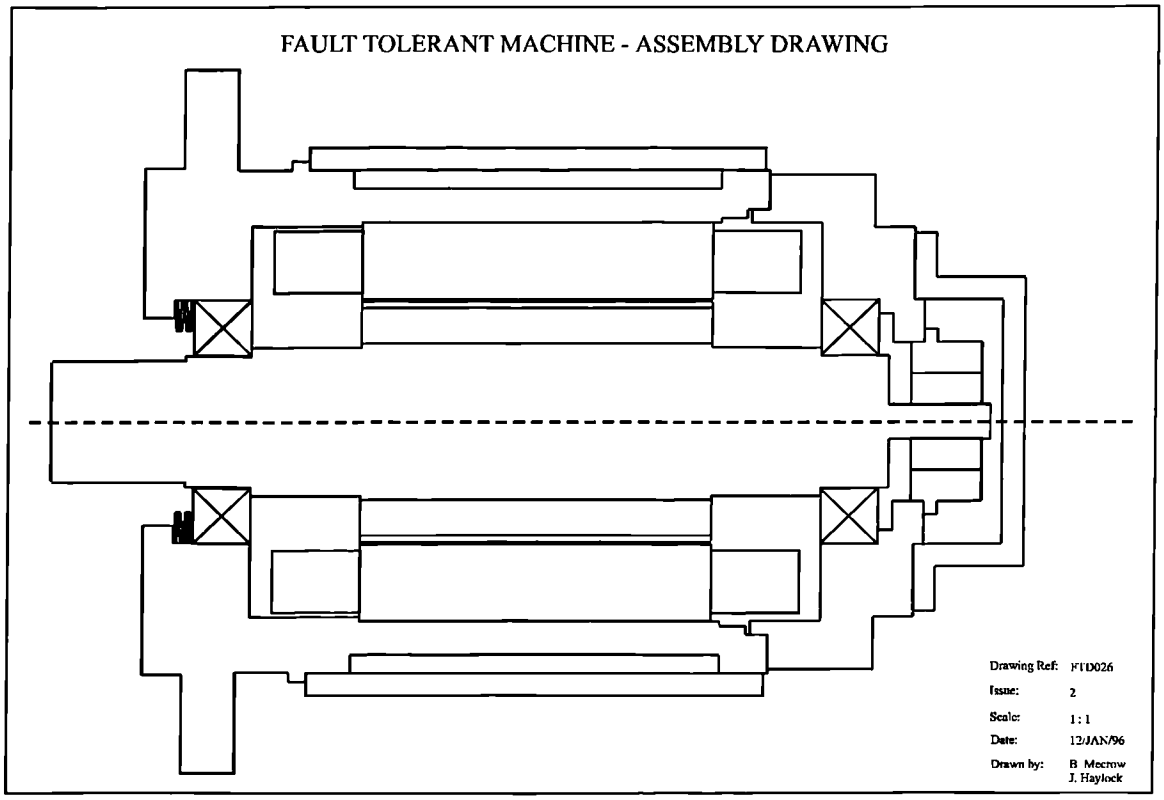

Figure 105 : Machine Assembly

Magnetic Flux Distribution in the Machine

Finite element analysis was used to obtain the magnetic flux distribution in the machine under open circuit conditions. This is shown in Figure 106.

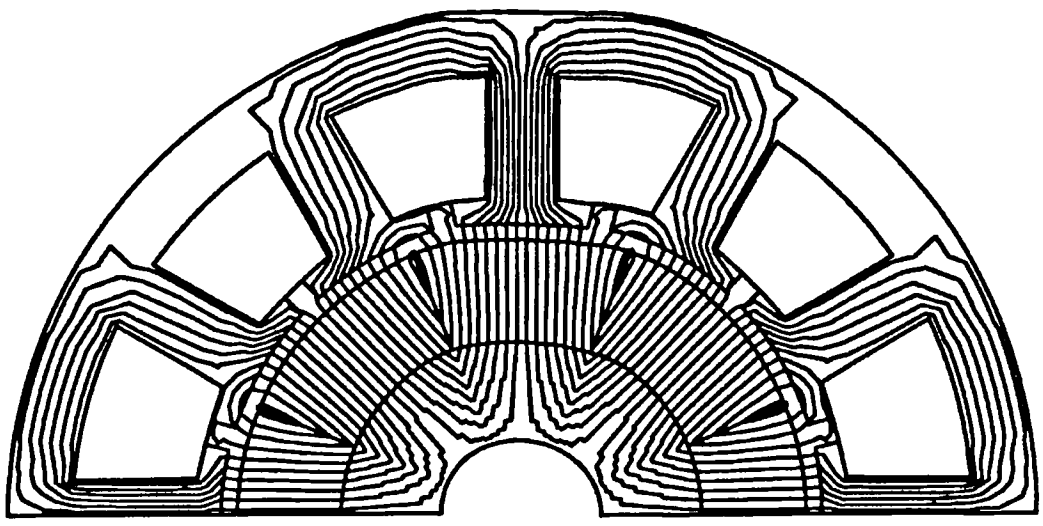

Figure 106 : Magnetic Flux Distribution in the Demonstrator Machine on No-Load

In the event of a phase terminal short circuit, the shorted phase carries current, producing flux which opposes the magnet flux. The result is that the magnet flux is shunted across the reactance slot, as shown in Figure 107

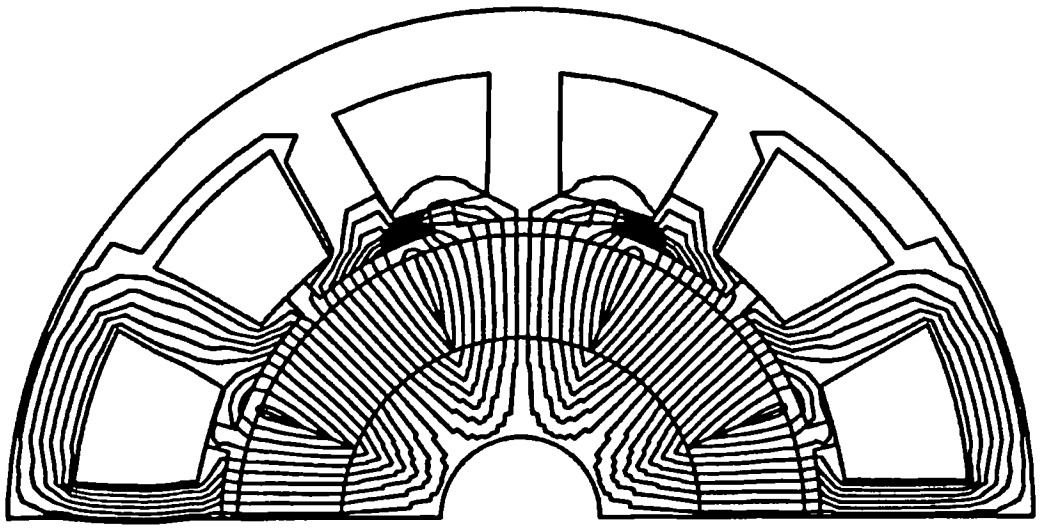

Figure 107 : Magnetic Flux Distribution in the Demonstrator Machine on No-Load, with One Phase Short Circuit.

Appendix B

THE POWER CONVERTER

The inverter design comprises six independent single phase bridges, supplied from a common DC link. Each bridge is rated as detailed in Table 20. Figure 108 shows the H-bridge assembly, comprising two IGBT modules, two electrolytic reservoir capacitors, three polypropylene decoupling capacitors and the DC link lamina.

Parameter	Value	Units
Proposed switching frequency	10	kHz
Reservoir capacitance per phase	2000	μF
Heatsink type	Forced air cooled	
Heatsink rating	< 0.20	°C/W per phase
Nominal DC link voltage	270	V
Maximum continuous DC current per phase	41.3	A

Table 20 : H-Bridge Parameters

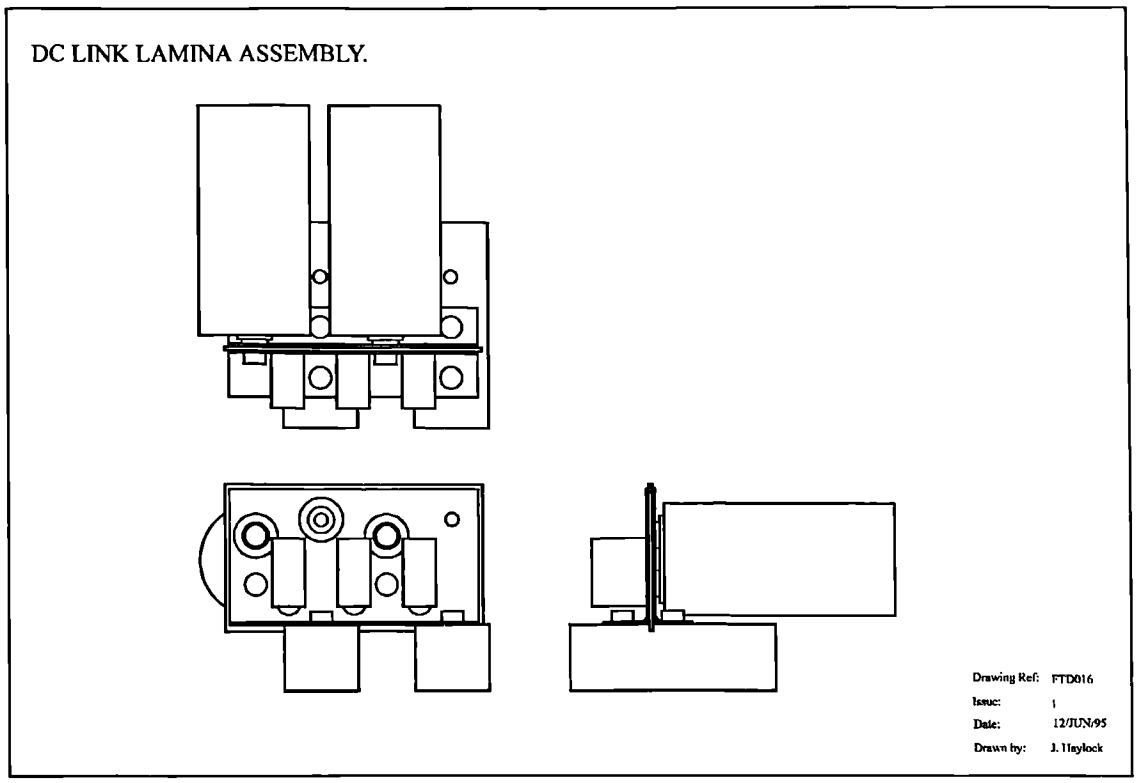

Figure 108 : H-Bridge Assembly Drawing

Appendix C

CONTROL CIRCUIT DETAILS

The power converter and electric machine are controlled by a set of electronics based on a TMS320C31 digital signal processor (DSP). These electronics ultimately control the amount of torque produced by the machine. The controlling DSP is supported by a field programmable gate array (FPGA) in which the PWM unit is implemented. Controllers configured in this way, with a DSP supported by PWM and protection functions implemented in gate arrays, are gaining gradual acceptance, [105].

The functional blocks which make up this controller are shown in Figure 109. The controller is physically split into 2 units. These units are referred to as the controller and the ADC pod. They are described here in terms of their component parts.

The control hardware samples each phase current at fixed points in each PWM interval. The rotor position and DC link voltage are sampled at the same instants.

Controller

The controller comprises a set of 6 cards mounted in a 19 inch 6U rack. The cards are as follows:

- DSP Card. The controller is based around a Texas Instruments TMS320C31 digital signal processor (DSP). The software to control the drive will run on this processor and this card controls all the data flow within the controller via a backplane based on the Loughborough Sound Images DBEX interface standard.
- ADC Interface Card. This card provides an interface between all analogue inputs and the DBEX backplane. The card includes 4 channels of 12 bit ADC circuitry so analogue signals may be input directly. The card also forms a link between up to 2 sets of 8 channels of external 12 bit ADC circuits and the DBEX backplane.

- Over-Current and Over-Voltage Monitor Card. This card takes inputs from the over-current and over-voltage detection circuits. These inputs are used to protect the system by forcing the gate commands inactive, and thus forcing some or all of the power devices to turn off. An over-voltage indication is used to turn off all power devices. An over-current indication results in the power devices in the affected phase turning off. The affected gate drive signals cannot be re-enabled until after the DSP card has been reset.

- Digital I/O Card. Any digital inputs or outputs for control or monitoring purposes will be connected to the controller via this card. It provides a pair of 24 bit inputs and a pair of 24 bit outputs. Also included as part of this card are a pair of 4.5 digit displays.

- PWM Card. The card is based around a three phase digital PWM generator circuit implemented in FPGA technology. Two of these devices are used to produce the 24 pulse width modulated (PWM) gate commands required to control the power converter. The card also produces a synchronisation pulse once per PWM cycle to allow the phase currents to be sampled at the correct instants.

- DAC / RDC Card. Two separate functions are implemented on this card. The first is to provide 6 analogue output channels from 12 bit DAC circuits. The second function is to derive a 12 bit rotor position measurement from a brushless resolver in the machine. The resolver output signals are converted to rotor position measurement using an AD2S80A, RDC from Analog Devices.

ADC Pod

The ADC pod is a small aluminium box, housing two cards associated with the measurement and monitoring of currents and voltages in the electric machine and power converter. Separating these functions from the controller and placing them in a casing adjacent to the power converter provides two advantages. These are the improved safety and accuracy of performing analogue to digital conversions at the point of measurement. The two cards within the ADC pod are:

- ADC Circuit. This card provides 8 channels of 12 bit ADC circuitry. Each channel can be configured to measure currents in the range ±50 amps or ±100 amps or voltages in the range ±500V. The outputs are in serial form and may be connected directly to the ADC Interface Card in the controller. Each voltage / current measurement channel is calibrated by adjusting the gain and zero offset of the analogue front end, feeding the ADC.

- Over-Current / Over-Voltage Detection Circuit. To protect power devices from faults resulting in excessive current very fast action is required. The response speed required is achieved by monitoring the analogue output of each current transducer and continuously comparing it to a threshold value. This comparison is performed by the Over-Current / Over-Voltage Detection Circuit and the result fed back to the controller. The over-current signals are latched in the controller and used to implement hardware over-current protection by immediately forcing all IGBT gate commands for the affected bridge to the off state. A similar comparison is made on the DC link voltage to check for over-voltage conditions caused by excessive regenerative braking of the electric machine.

Parameter	Value	Units
ADC pod input filter bandwidth	80	kHz
ADC pod input filter time constant	12.5	μs
ADC conversion time	8	μs
Voltage channel resolution (± 500 V / 4096)	0.244	V
Current channel resolution (± 100 A / 4096)	0.049	A
Rotor position resolution (360 ° mechanical / 4096)	0.088	mechanical degrees

Table 21 : ADC Channel Specification

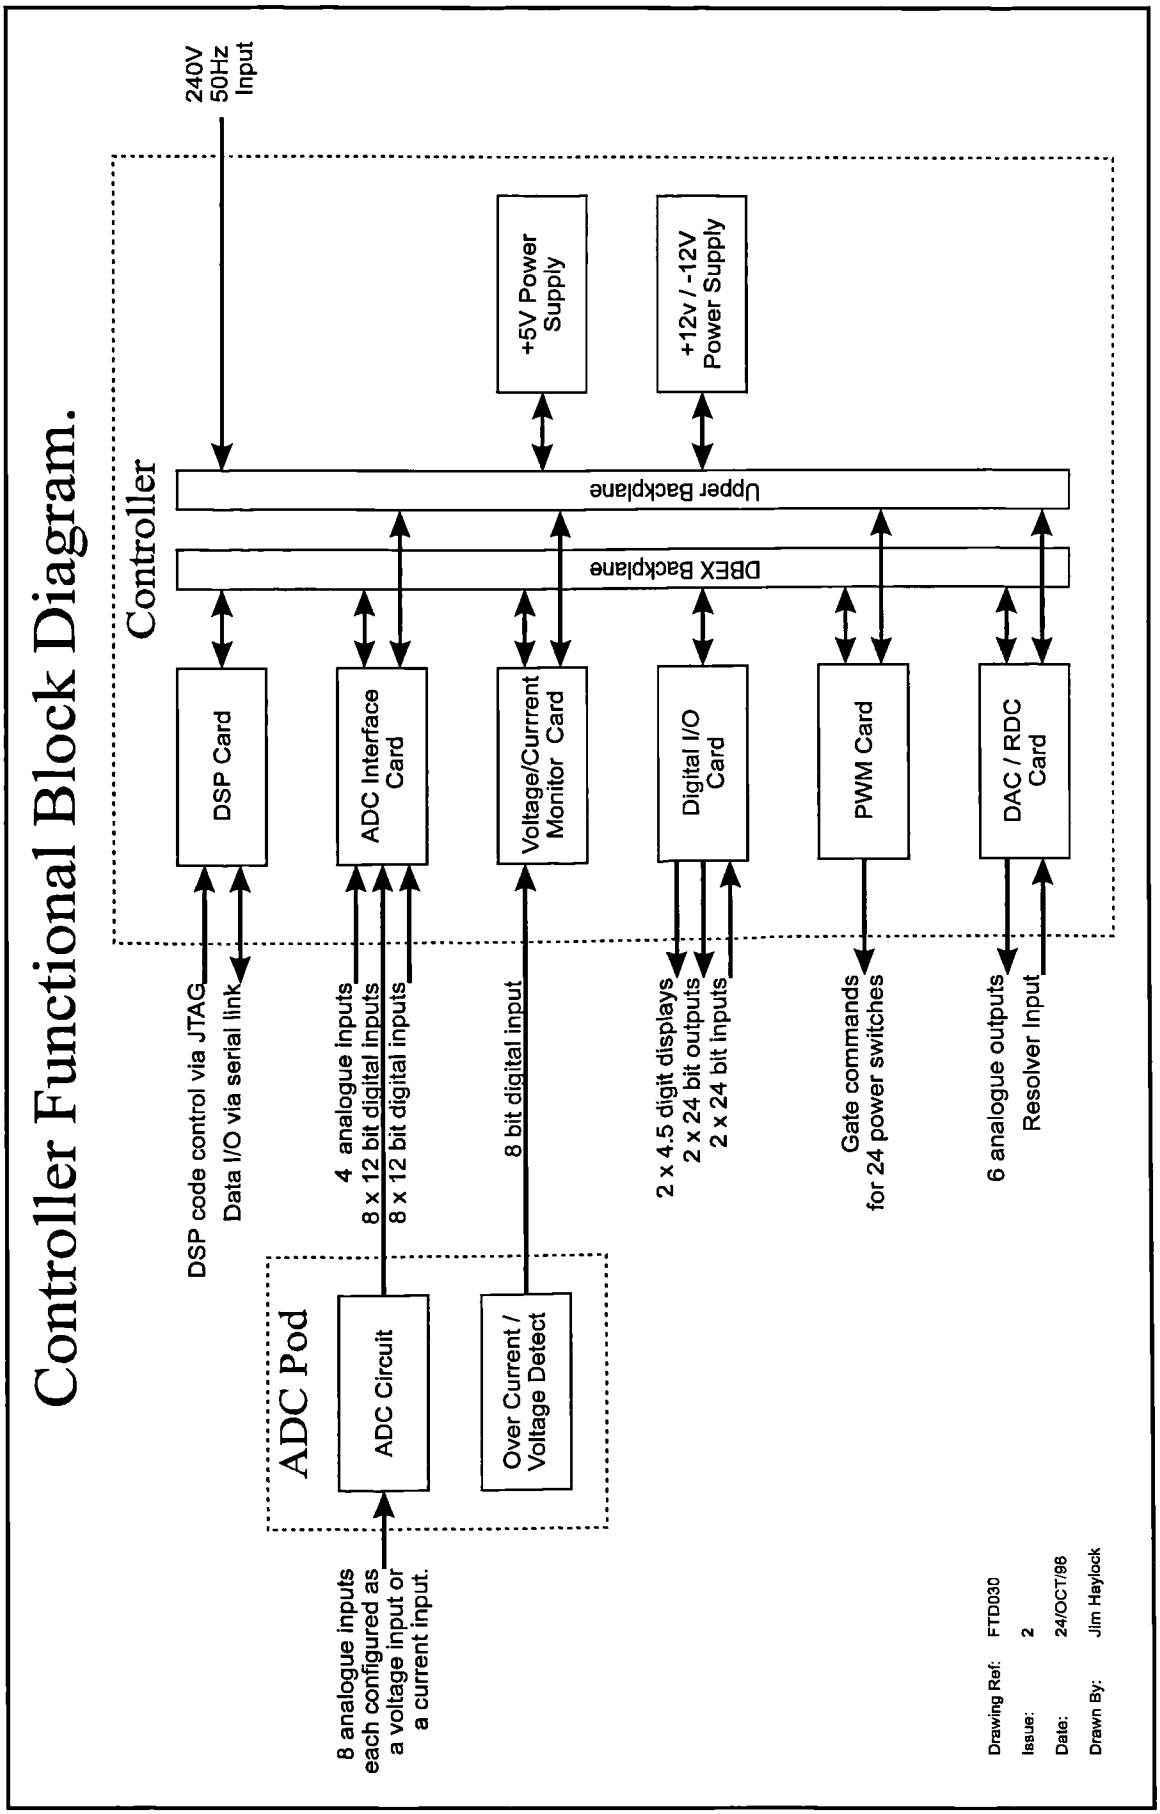

Figure 109 : Control Electronics Functional Block Diagram

Appendix D

THE MACHINE TEST RIG

A test rig has been constructed to allow the fault tolerant demonstrator drive to be tested on load. This comprises a DC power supply and a high speed motor loading rig. Figure 110 shows the functional blocks in the demonstrator drive and test rig.

A simple DC power supply, not shown in Figure 110, has been built to supply the DC link required by the inverter. The DC voltage is derived by rectifying a 3 phase 50 Hz supply. The power supply is housed in the same enclosure as the inverter and further details are given in Table 22.

The demonstrator drive comprises a set of control electronics, a six phase inverter and the demonstrator machine, shown on the left of Figure 110. The demonstrator machine is coupled to a DC load machine through a torque meter and a toothed belt. The Toothed belt provides the gearing required to allow the DC machine to operate at 3000 rpm whilst the demonstrator machine is running at 13000 rpm. The torque meter provides shaft torque measurement from speeds of 13000 rpm down to standstill. Power generated by the load machine is dissipated in a load bank.

Figure 111 shows the demonstrator machine mounted on the loading rig. The demonstrator machine is clearly visible in the top left corner of the photo. The toothed belt drive is situated in the guarded area, whilst the torque meter and the load machine

Parameter	Value	Units
Power supply rating	20	kW
Voltage input (rms. line voltage)	236	V
Current drawn from ac supply (rms. per phase)	60.4	A
Unregulated voltage output	315	V DC
Maximum continuous output current	74	A DC

Table 22 : DC Link Power Supply Specification.

are clearly visible on either side of this.

In addition to torque and speed measurements from the torque meter, measurements are made of the DC link voltage and current and the load machine armature voltage and current. This allows power flow through the rig to be calculated.

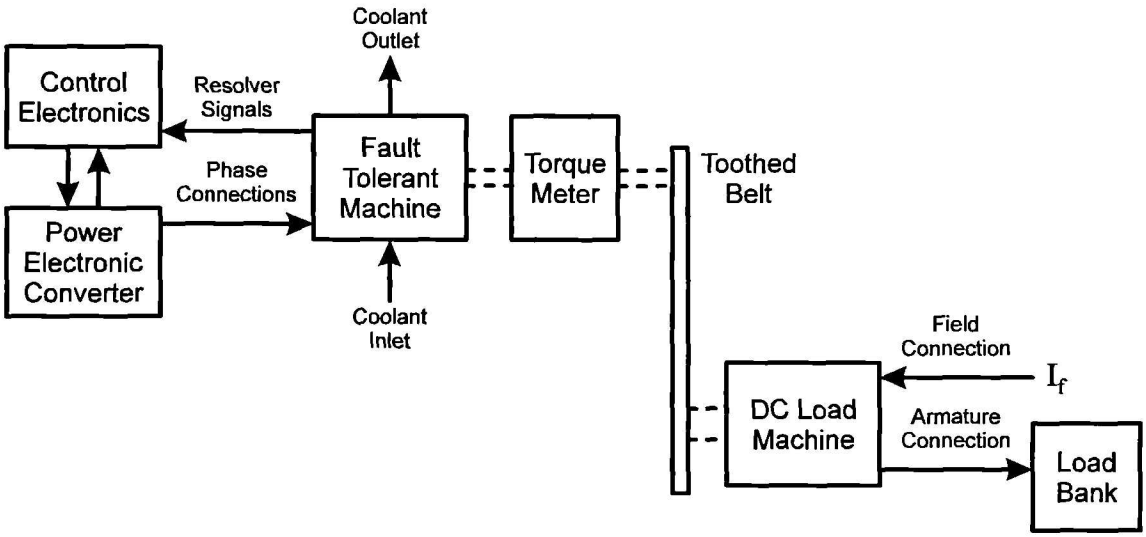

Figure 110 : Drive Test Rig Functional Block Diagram

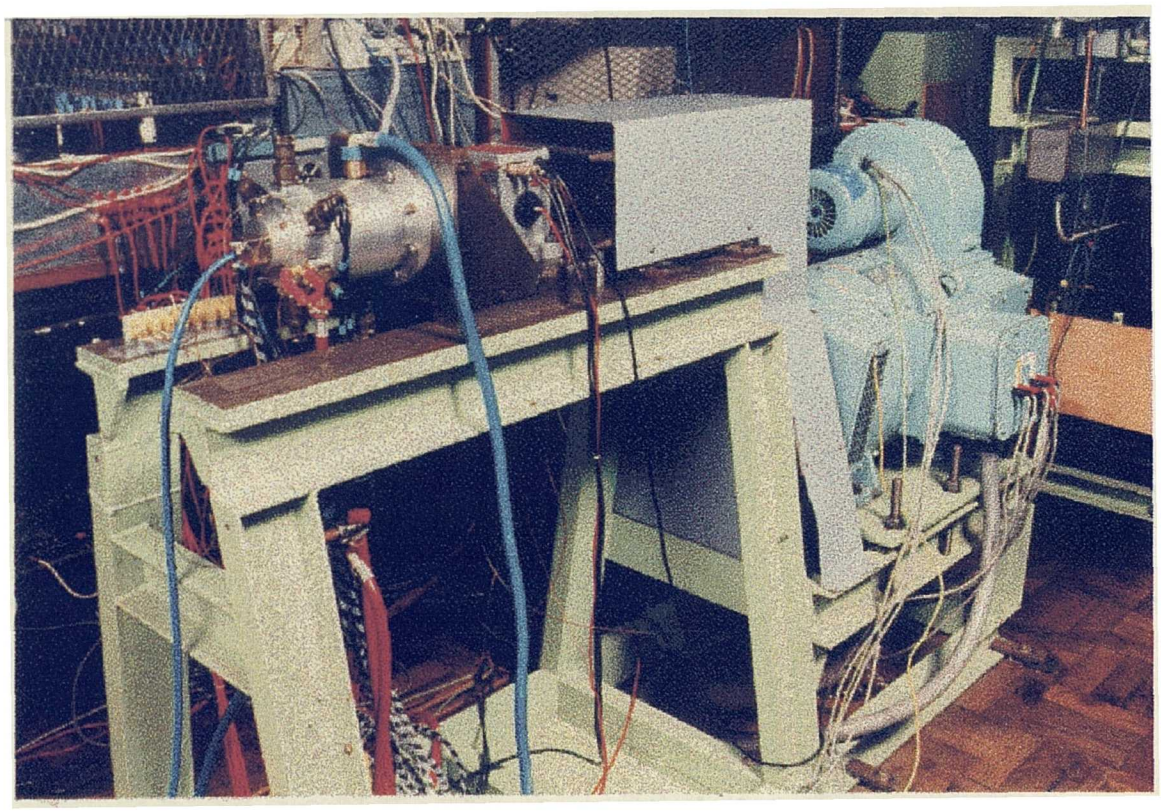

Figure 111 : The Drive Test Rig

Appendix E

COLLECTING FLUX LINKAGE DATA

Introduction

In Chapter 3, a current controller based on a model of the machine was presented. In Chapter 6 and Chapter 7 the same machine model was used in a fault detection scheme. A look-up table, relating flux-linkage in the machine to current and rotor angle, was central to the machine model. The method used for gathering the flux linkage-data was briefly outlined in Chapter 3. A more detailed description of the flux-linkage data gathering process is given here.

Recall from Chapter 3 that, in general, the flux linkage, ψ, will be a function of the rotor position and the current in each phase [75]. However, in this machine each phase is magnetically isolated from all other phases so the flux linking a phase is a function of rotor position and the current in that phase alone. Thus expression below relates the terminal volts for a phase to the current in that phase and the rotor position.

$$V = i \cdot R + \frac{d\psi}{dt} \tag{31}$$

The flux linkage in a permanent magnet machine has two components. These are the magnet flux linking the phase and the flux resulting from current in the phase which links that phase. It is simplest to measure the flux linkage in two stages. First the magnet flux linking the phase is measured. Then the flux linkage resulting from phase current is measured and added to the magnet flux linkage.

Measurement of Magnet Flux Linking a Phase

The change in magnet flux linking the phase gives the back-EMF for the phase. It is possible to work backwards, from the back-EMF profile for a phase, to the magnet flux

linking that phase. This is achieved by integrating the back-EMF, so the magnet flux linking the phase, $\psi_f(\theta)$, is given by:

$$\psi_f(\theta) = \int_0^{\theta} E(\theta)d\theta + \psi_f(0) \tag{32}$$

The value of magnet flux at the starting position, $\psi_f(0)$, is not known. However it is known that the magnet flux linking a phase averages to zero over an electrical cycle:

$$\int_0^{2\pi} \psi_f(\theta)d\theta = 0 \tag{33}$$

The value of $\psi_f(0)$ can be obtained (32) from and (33). First rearrange (33) to give:

$$\psi_f(\theta) - \psi_f(0) = \int_0^{\theta} E(\theta)d\theta \tag{34}$$

Integrating $\psi_f(\theta) - \psi_f(0)$ over an electrical cycle gives:

$$\int_0^{2\pi} \psi_f(\theta) - \psi_f(0)d\theta = \int_0^{2\pi} \psi_f(\theta) - \int_0^{2\pi} \psi_f(0)d\theta \tag{35}$$

Substituting (32) into (35):

$$\int_0^{2\pi} \psi_f(\theta) - \psi_f(0)d\theta = 0 + \int_0^{2\pi} \psi_f(0)d\theta \tag{36}$$

And $\psi_f(0)$ is a constant:

$$\int_0^{2\pi} \psi_f(\theta) - \psi_f(0)d\theta = 2\pi.\psi_f(0) \tag{37}$$

Rearranging gives $\psi_f(0)$:

$$\psi_f(0) = \frac{1}{2\pi} \int_0^{2\pi} \psi_f(\theta) - \psi_f(0)d\theta \tag{38}$$

Thus obtaining $\psi_f(\theta)$ is a four stage process:

1. Measure the back-EMF variation with rotor position, $E(\theta)$.
2. Integrate the back-EMF data between 0 and θ to obtain '$\psi_f(\theta)$ - $\psi_f(0)$', (34).
3. Integrate '$\psi_f(\theta)$ - $\psi_f(0)$' over a full electrical cycle to obtain, $\psi_f(0)$, (38).
4. Add $\psi_f(0)$ to '$\psi_f(\theta)$ - $\psi_f(0)$' to obtain, $\psi_f(\theta)$.

The flux linkage contributed by the magnets is shown in Figure 112. Note that, because the magnet flux linkage is obtained by integrating the back-EMF, the third harmonic content of the magnet flux linkage is much lower than for the back-EMF - see Figure 11 in Chapter 2.

Measurement of Flux Linking a Phase due to Current in that Phase

The flux linkage associated with a given phase current can be calculated from the time integral of the voltage required to drive the current from zero to the level of interest (with the rotor stationary).

The machine rotor is locked at the angle of interest, θ_0. With the phase current initially at zero, a voltage step is applied to the phase, at time t_0. The voltage at the machine terminals, v(t), and the current in the phase, i(t), are measured. When the phase current reaches the maximum value of interest the voltage source is disconnected. The data gathered during this experiment is then processed to calculate the flux linkage corresponding to each current measurement. The flux linkage corresponding to the current at time t_1 is given by:

$$\psi_a(i(t_1),\theta_0) = \int_0^{t_1} v(t) - i(t)\cdot R \, dt \tag{39}$$

The test is repeated for a series of rotor angles. The data gathered in this test is shown in

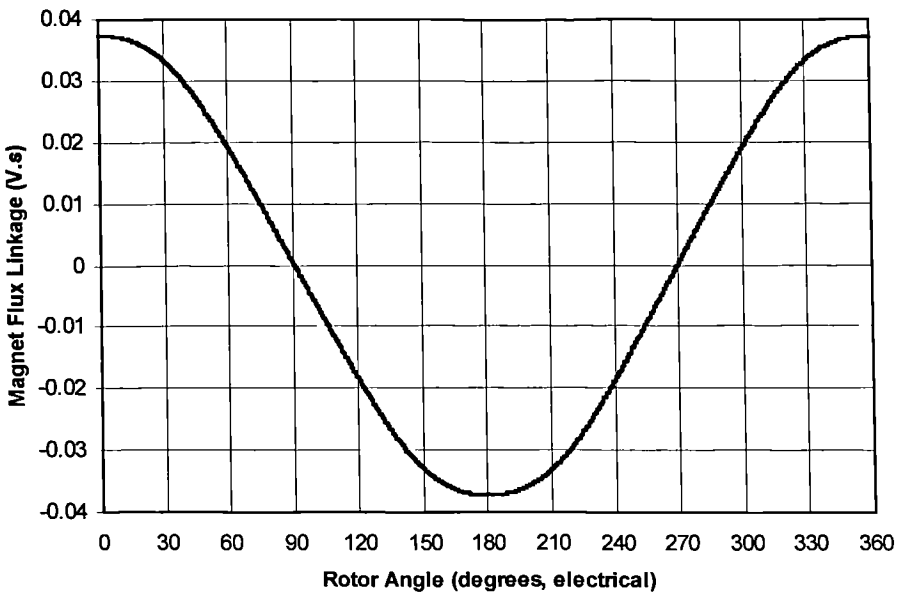

Figure 112 : Measured Magnet Flux Linking Phase A of the Machine

Figure 113. The effect of saturation is clearly visible at the point on each curve where the flux linkage per unit current begins to increase more slowly. This appears to occur at a different value of flux linkage at each rotor position. This is because the flux linkage shown does not include the flux contributed by the magnets.

Total Flux Linkage

The total flux linking a phase is the sum of the magnet flux linkage and the flux linkage associated with the phase current:

$$\psi(i,\theta) = \psi_a(i,\theta) + \psi_f(\theta) \tag{40}$$

The total flux linkage characteristic, measured on the demonstrator machine, is shown in Figure 114. The effect of saturation is clearly visible at the point on each curve where the flux linkage per unit current begins to increase more slowly.

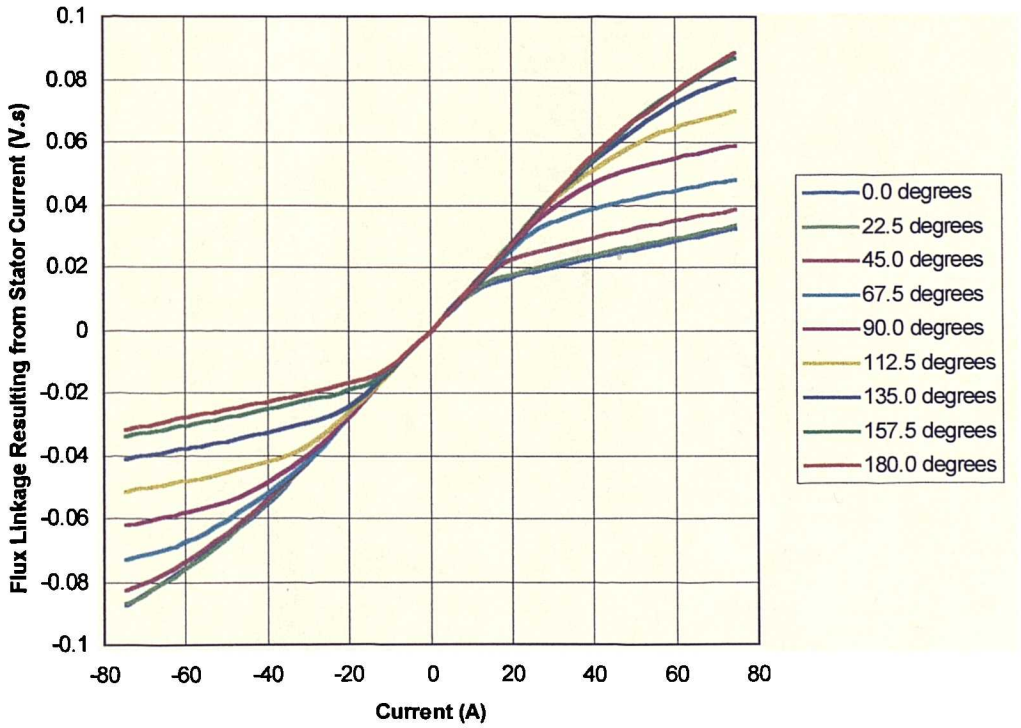

Each curve represents one rotor position from 0 degrees (bottom curve) to 180 degrees (top curve) in steps of 22.5 degrees.

Figure 113 : Measured Stator Current Driven Flux Linking Phase A of the Machine

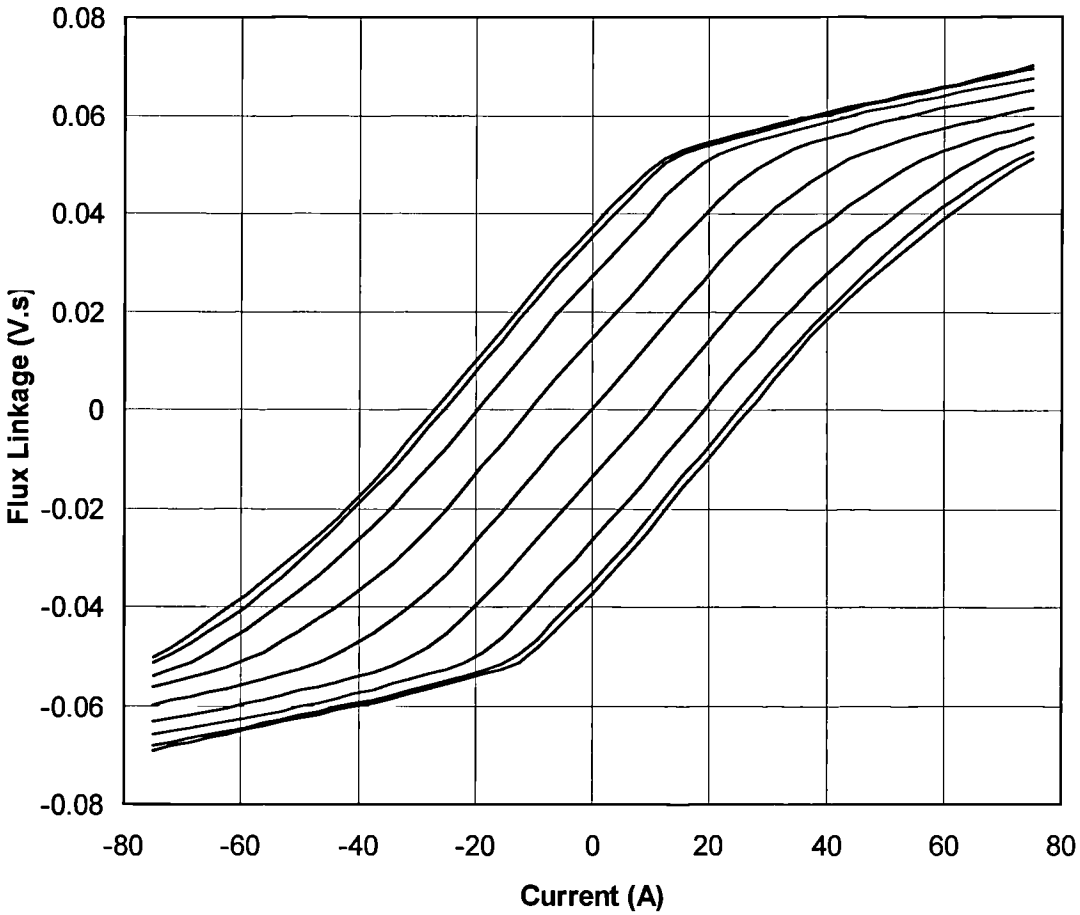

Each curve represents one rotor position from 0 degrees (top curve) to 180 degrees (bottom curve) in steps of 22.5 degrees.

Figure 114 : Flux Linkage Data for Demonstrator PM Machine

Appendix F

CALCULATING STATIC TORQUE FROM FLUX LINKAGE

When investigating operation in various fault conditions, it is important to consider the resulting torque waveform. As discussed in Chapter 4, the instantaneous torque delivered by a phase of the demonstrator machine is a function of the phase current and rotor angle. The function which relates torque to phase current and rotor angle may be obtained by static torque measurement. Alternatively the torque function can be calculated from the data relating flux linkage to phase current and rotor angle. The second alternative has the advantages that the resulting torque function is known for a wider range of currents and that the ψ-i-θ data is already available as a lookup table.

The technique for deriving static torque from the Flux-MMF diagram is described by Staton et al [75]. The instantaneous electro-magnetic torque, T_e, is calculated from the rate of change of co-energy, W', with respect to angle as shown in (41).

$$T_e(\theta_n, i_n) = \left.\frac{\partial W'}{\partial \theta}\right|_{i=\text{constant}} \tag{41}$$

The change in co-energy, for an incremental change in rotor angle, is the energy converted from electrical energy to rotational energy. The co-energy associated with a particular locus on the ψ-i plane is the line integral around that locus. Thus the rate of change of co-energy with respect to θ may be calculated from the small change in co-energy that results as the rotor moves between an angle just less than θ, referred to as θ_{n-1}, and an angle just above θ, referred to as θ_{n+1}. Then the static torque is given directly by (42).

$$T(\theta_n, i_n) = \frac{W'(\theta_{n+1}, i_n) - W'(\theta_{n-1}, i_n)}{\theta_{n+1} - \theta_{n-1}} \tag{42}$$

This change in co-energy can be obtained graphically from the ψ-i diagram. The change in co-energy required to give the torque at point θ_n , i_n (point E in Figure 115) is the area enclosed by the path:

- starting from i=0, $\theta = \theta_{n-1}$ (point A)
- holding $\theta = \theta_{n-1}$ while increasing i to i_n (point B)
- holding $i = i_n$ while increasing θ to θ_{n+1} (point C)
- holding $\theta = \theta_{n+1}$ while reducing i to 0 (point D)
- holding i = 0 while reducing θ to θ_{n-1} (point A)

Whilst adopting the convention that clockwise encirclements give negative area and counter clockwise encirclements give positive area. Note that the encirclement on the ψ-i diagram shown is only for illustration. In practice a θ_{n+1} and θ_{n-1} would be much closer to θ, giving a very narrow loop. This is possible because the diagram only shows ψ-i curves at angle increments of 22.5 degrees where as the lookup table ψ-i data at angle increments of 5.625 degrees.

The end result is a new lookup table containing T_e-i-θ data. Plotting data from the look up table yields the set of curves presented in Figure 116.

Calculation of instantaneous torque from the ψ-i-θ data assumes that there is no significant electromagnetic coupling between phases. This is a reasonable assumption

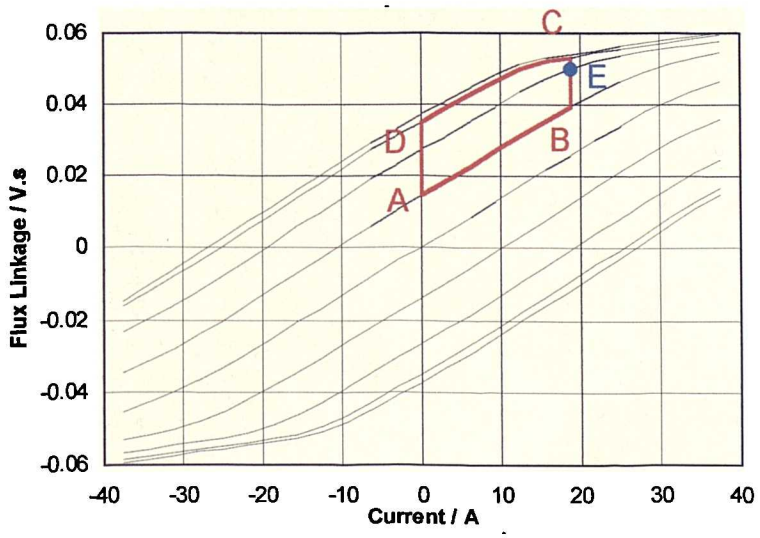

Figure 115 : Derivation of Static Torque Data from ψ-i-θ Data.

since the measured parameters of the machine show that the design aim of effective isolation between machine phases has been largely achieved, (Table 4 in Chapter 2). Note that the torque calculated using this technique accounts only for torque production resulting from flux which links the stator winding. The phenomenon of cogging in a PM machine results from interaction between the magnets and stator slotting and does not cause any change in the flux linking the stator coils. Thus the electromagnetic torque calculated from ψ-i-θ data does not take account of cogging torque, Deodhar et al [78]. This is not a serious concern here, since the cogging torque over a machine cycle averages to zero and, as shown in Chapter 2, the peak value of the ripple resulting from cogging is only 2.3% of rated torque.

Disregarding the cogging torque, the static torque curves calculated from ψ-i-θ data are in good agreement with those obtained from direct measurement of static torque in Chapter 2. The discrepancy between calculated and measured static torque is always less than 5%.

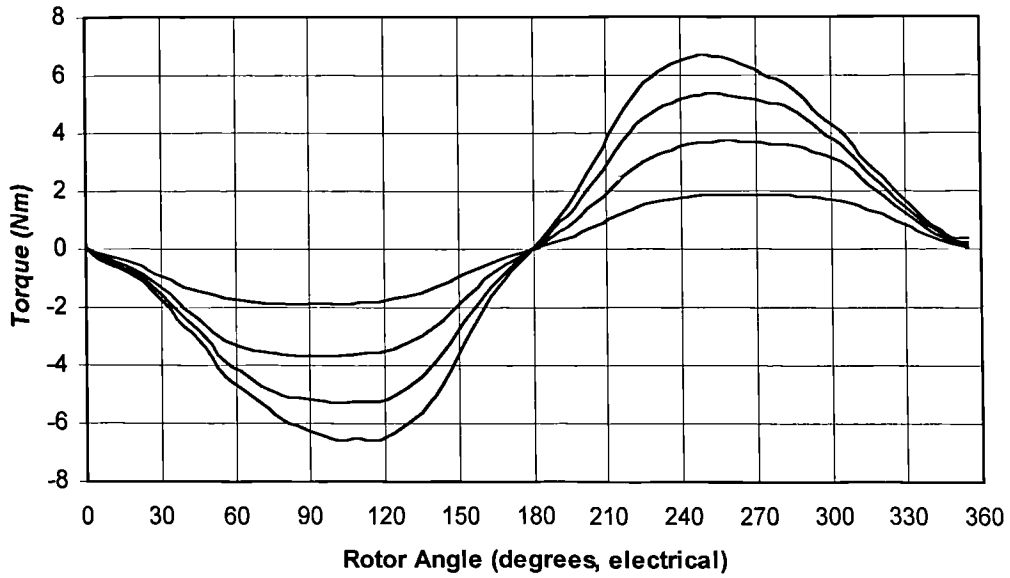

The curves show torque against angle for currents of 12.5 A, 25.0 A, 37.5 A and 50.0 A

Figure 116 : Static Torque Data Derived from ψ-i-θ Data.

Appendix G

FURTHER DETAILS ON FREEWHEEL CURRENT SENSING

A fault detection scheme based on the freewheel current sensor was discussed in Chapter 8. Further information on the practical considerations required to make this fault detection scheme work are presented here.

Current Bands Required for the Freewheel Current Sensor

Recall that a fault detection scheme based on the freewheel current sensor can detect a range of faults using only a few specific current bands. The discussion in Chapter 8 concluded that only three current bands are required to detect short circuit faults, but that open circuit fault detection required a total of five bands. This section will discuss where the levels separating these bands should be set.

The three bands required for short circuit fault detection are:

'<<0' $\Rightarrow$ negative out of range

'in range' $\Rightarrow$ in range

'>>0' $\Rightarrow$ positive, out of range

These bands are applied to the phase current and the freewheel current measurements as shown in Figure 117. Note that the break points for the freewheel current 'out of range' levels are set beyond the phase current 'out of range' levels. This ensures that, even accounting for offset and gain errors, any fault resulting in both excessive phase current and excessive freewheel current will always be reported as 'phase current out of range' in advance of 'freewheel current out of range'.

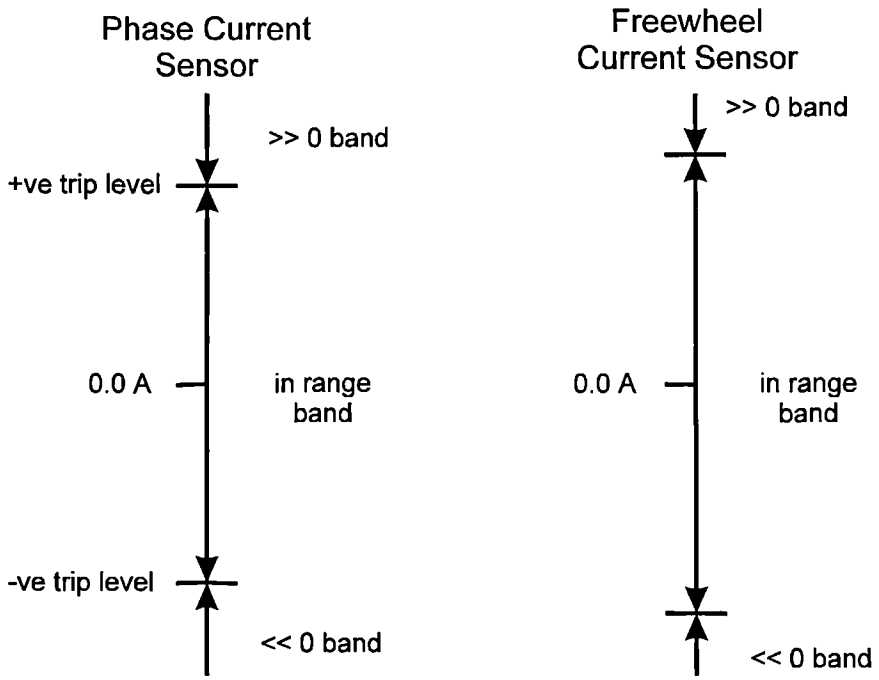

Figure 117 : Bands for Phase Current and Freewheel Current to Enable Short Circuit Power Device Detection

If open circuit faults are to be detected then the phase current 'in range' band and the freewheel current 'in range' band must each be subdivided into three bands. The resulting total of five non-overlapping bands is as follows:

'<<0' ⇒ negative out of range

'<0' ⇒ negative, in range

'=0' ⇒ near zero

'>0' ⇒ positive, in range

'>>0' ⇒ positive, out of range

In practice the zero band must be a 'near zero band' as shown in Figure 118. The width of the freewheel current sensor '=0' band is set to ensure that the largest zero offset in the current transducer output falls within the '=0' band. The phase current '=0' band extends beyond the freewheel current '=0' band, by an amount that ensures that no current falling inside the freewheel current '=0' band falls outside the phase current '=0' band. To achieve this the phase current '=0' band limit is set to the freewheel current

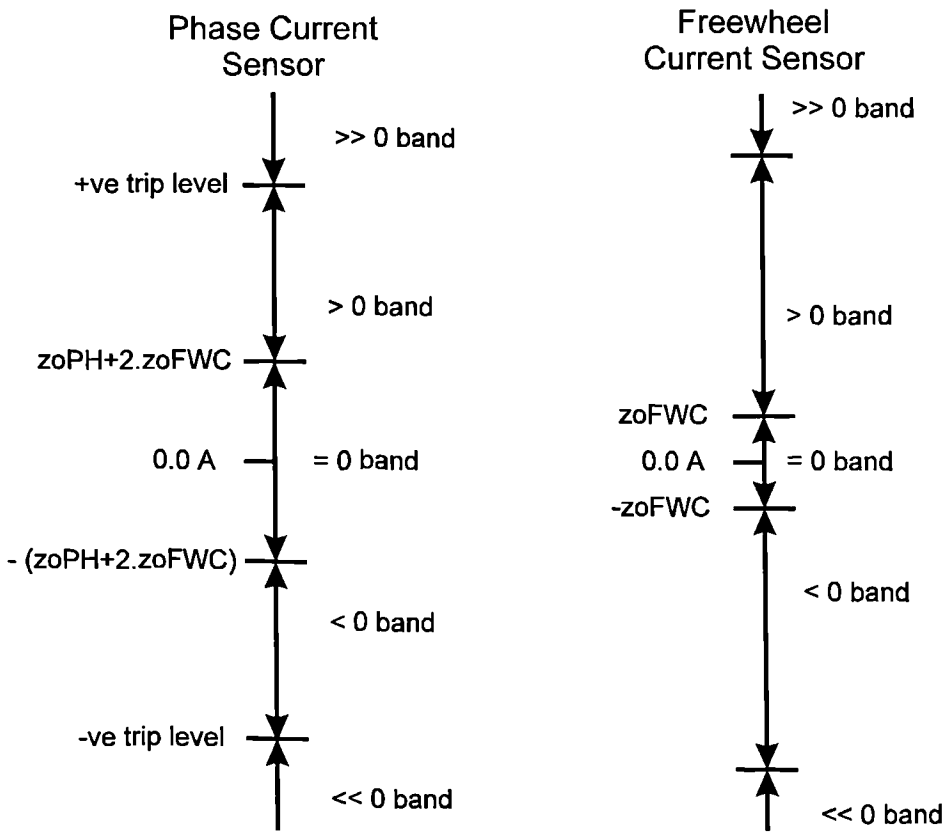

Figure 118 : Bands for Phase Current and Freewheel Current to Enable Open Circuit Power Device Detection

'=0' band limit plus the maximum possible zero offset in the freewheel current transducer output plus the maximum possible zero offset in the phase current transducer output.

State Machine Implementation of the Fault Detection Scheme

The fault detection scheme based on the freewheel current sensor, described in Chapter 8, was implemented as a finite state machine. The state machine was designed to implement the fault detection decision making described in Table 19 of Chapter 8 and to discriminate between noise on the freewheel current sensor output and real fault conditions. To avoid problems with noise on the freewheel current sensor output, only the '>>0' and '<<0' levels are monitored for the first 16 μs after each change in bridge switching state. To further minimise the possibility of inaccurate fault detection, any fault must persist for 6 cycles of the state machine before it is acknowledged and acted upon. This results in a delay of 1.5 μs in fault detection, but this is still only 15% of the short circuit withstand time of the power devices. A full list of the states and state assignments in the resulting state machine is shown in Table 23.

Name	ID	State Assignment (Binary)	State Assignment (Decimal)
Reset	RESET	000 0000 0	0
All OK, SW pattern = 0000 (FU;FL;RU;RL)	OK0000	000 0000 1	1
All OK, SW pattern = 0000 (FU;FL;RU;RL)	OK0001	000 0001 1	3
All OK, SW pattern = 0000 (FU;FL;RU;RL)	OK0010	000 0010 1	5
All OK, SW pattern = 0000 (FU;FL;RU;RL)	OK0100	000 0100 1	9
All OK, SW pattern = 0000 (FU;FL;RU;RL)	OK1000	000 1000 1	17
All OK, SW pattern = 0000 (FU;FL;RU;RL)	OK0101	000 0101 1	11
All OK, SW pattern = 0000 (FU;FL;RU;RL)	OK0110	000 0110 1	12
All OK, SW pattern = 0000 (FU;FL;RU;RL)	OK1001	000 1001 1	19
All OK, SW pattern = 0000 (FU;FL;RU;RL)	OK1010	000 1010 1	21
Control Failure	CF	010 0000 0	64
Suspected FUOC	SFUOC	111 0000 1	225
Suspected FLOC	SFLOC	110 0000 1	193
Suspected RUOC	SRUOC	101 0000 1	161
Suspected RLOC	SRLOC	100 0000 1	129
Suspected FUSC	SFUSC	111 1010 1	245
Suspected FLSC	SFLSC	110 0101 1	203
Suspected RUSC	SRUSC	101 1010 1	181
Suspected RLSC	SRLSC	100 0101 1	139
Suspected MTSC	SMTSC	011 1010 1	117
Dead time for FUSC	DFUSC	111 1000 0	240
Dead time for FLSC	DFLSC	110 0100 0	200
Dead time for RUSC	DRUSC	101 0010 0	164
Dead time for RLSC	DRLSC	100 0001 0	130
Dead time for MTSC	DMTSC	011 0000 0	96
Confirmed FUOC	CFUOC	111 0000 0	224
Confirmed FLOC	CFLOC	110 0000 0	192
Confirmed RUOC	CRUOC	101 0000 0	160
Confirmed RLOC	CRLOC	100 0000 0	128
Confirmed FUSC	CFUSC	111 1010 0	244
Confirmed FLSC	CFLSC	110 0101 0	202
Confirmed RUSC	CRUSC	101 1010 0	180
Confirmed RLSC	CRLSC	100 0101 0	138
Confirmed MTSC	CMTSC	011 1010 0	116
Multi-point Fault	MPF	111 1111 1	255

Table 23 : State Assignments for the Freewheel Current Sensing State Machine

List of References

[1] G. Cosulich, M. Fracchia, A. Marriscotti, S. Savio, 'Comparative Dependability Analysis of Electric Ship Propulsion Systems.', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 3.328-3.333.

[2] R.H. Caplen, 'A Practical Approach to Reliability', Business Books Limited, 1972. ISBN 0-220-66809-4.

[3] D.E.A. Clarke, G.F. Cornwell, 'A Fault Tolerant Power Converter System', INTELEC'85, Seventh International Telecommunications Energy Conference (proceedings), Munich, West Germany, 14-17 October 1985, pp. 499-504.

[4] D. Kastha, B.K. Bose, 'Investigation of Fault Modes of Voltage-Fed Inverter System for Induction Motor Drive', IEEE Transactions on Industry Applications, Vol. 30, No. 4, July/August 1994, pp. 1028-1037.

[5] J. Liddiard, 'An Introduction to Safety Critical Systems', Embedded System Engineering, February/March 1994, pp. 38-41.

[6] R.V. White, 'Fault Tolerance In Distributed Power Systems', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 2.851-2.857.

[7] D. Kastha, B.K. Bose, 'On-Line Search Based Pulsating Torque Compensation of a Fault Mode Single-Phase Variable Frequency Induction Motor Drive', IEEE Transactions on Industry Applications, Vol. 31, No. 4, July/August 1995, pp. 802-811.

[8] D. Kastha, 'Fault Mode Single-Phase Operation of a Variable Frequency Induction Motor Drive and Improvement of Pulsating Torque Characteristics', IEEE Transactions on Industry Applications, Vol. 41, No. 4, August 1994, pp. 426-432.

[9] G. Sinha, C. Hochgraf, R.H. Lasseter, D.M. Divan, T.A. Lipo, 'Fault Protection in a Multilevel Inverter Implementation of a Static Condenser', IEEE 30th Industry Applications Society Annual Meeting, IAS'95, Orlando, Florida, USA, 8-12 October 1995, Part 3 (of 3), pp. 2557-2564.

[10] M. Karyagina 'Designing For Fault-tolerance in the Commercial Environment', Proceedings IEEE Annual Reliability and Maintainability Symposium, Las Vegas, NV, USA, 22-25 January 1996, pp. 258-262.

[11] R.K. Prasad, J.H. Aylor, B.W. Johnson, 'Fault Tolerant Motor Drives for Powered Wheelchairs', IEEE Proceedings of 22nd Annual Southeastern Symposium on System Theory, Cookeville, TN, USA, 11-13 March 1990, pp. 420-425.

[12] P.K. Steimer, 'Redundant, Fault Tolerant Control System for a 13 MW High Speed Drive', IEE Third International Conference on Power Electronics and Variable Speed Drives, London, England, 13-15 July 1988, pp. 343-346.

[13] S.Corrias, K. Metzger, K.Schweizer, P. Steimer, 'Adjustable Speed, Converter-Fed Synchronous Motor Used to Drive a 13-MW Turbocompresser', ABB Review, Switzerland, 1990, ISSN 1013-3119, No. 7, pp. 19-26.

[14] T.M. Jahns, 'Improved Reliability in Solid State AC Drives by Means of Multiple Independent Phase Drive Units', IEEE Transactions on Industry Applications, Vol. 16, No. 3, May/June 1980, pp. 321-331.

[15] E. Wu, J. Hwang, J. Chladek, 'Fault Tolerant Joint Drive System for the Space Shuttle Remote Manipulator System: Analysis and Experiment', IEEE Transactions on Robotics and Automation, Vol. 9, No. 5, October 1993, pp. 675-684.

[16] K.R. Boyer, 'Redundant, Thin Wing, High Pressure, Direct Drive Valve Actuation System', Aerospace Technology Conference and Exposition, Long Beach, CA, USA, 1-4 October 1990, p.15p.

[17] G. Nerowski, K. Plankner, B. Piepenbreier, H.J. Tolle, 'New Permanent Field Synchronous Motor With Integrated Inverters', International Conference on Electric Machines, ICEM'90, Cambridge, MA, USA, 13-15 August 1990, Vol. 1, pp. 124-131.

[18] A.M.Oliveira, A.G. Balan Palhares, A.H. Kumakura, G. Winnischofer, A. Hoshino, 'Analysis of Brushless DC Motor Performance when Faults Occurrence', IEE, Proceedings of the 4th European Conference on Power Electronics and Applications, EPE'91, Florence, Italy, 3-6 September 1991, Vol. 3, pp. 445-450.

[19] E. Richter, 'High Temperature, Lightweight, Switched Reluctance Motors and Generators for Future Aircraft Engine Applications', Proceedings of 1988 American Control Conference, Atlanta, GA, USA, 15-17 June 1988, pp. 1846-1851.

[20] E. Richter, 'Switched Reluctance Machines for High Performance Operations in a Harsh Environment - A Review Paper', International Conference on Electric Machines, ICEM'90, Cambridge, MA, USA, 13-15 August 1990, Vol. 1, pp. 18-24.

[21] E. Richter, A.V. Radun, C. Ferreira, E. Ruckstadter, 'An Integrated Electrical Starter/Generator System for Gas Turbine Application, Design and Test Results', International Conference on Electric Machines, ICEM'94, Paris, France, 5-8 September 1994, Vol. 3, pp. 286-291.

[22] E. Richter, C. Ferreira, 'Performance Evaluation of a 250 kW Switched Reluctance Starter Generator', IEEE 30th Industry Applications Society Annual

Meeting, IAS'95, Orlando, Florida, USA, 8-12 October 1995, Part 1 (of 3), pp. 434-440.

[23] A.V. Radun, 'High Power Density Switched Reluctance Motor Drive for Aerospace Applications', IEEE Transactions On Industry Applications, Vol. 28, No. 1, January/February 1992, pp. 113-119.

[24] C.M. Stephens, 'Fault Detection and Management System for Fault-Tolerant Switched Reluctance Motor Drives', IEEE Transactions on Industry Applications, Vol. 27, No. 6, November/December 1991, pp. 1098-1102.

[25] T.J.E. Miller, 'Faults and Unbalanced Forces in the Switched Reluctance Machine', IEEE Transactions on Industry Applications, Vol. 31, No. 2, March/April 1995, pp. 319-328.

[26] B.C. Mecrow, 'New Winding Arrangements for Doubly Salient Reluctance Machines', IEEE 27th Industry Applications Society Annual Meeting, IAS'92, Houston, Texas, USA, 8-12 October 1992, Vol. 1, pp. 249.

[27] P.R. Palmer, 'A Unipolar Inverter Drive for a Cage Induction Motor', PhD Thesis, London University, UK, 1985.

[28] D. Kastha, B.K. Bose, 'Investigation of Fault Modes of a Voltage Fed Inverter System for Induction Motor Drive', IEEE 27th Industry Applications Society Annual Meeting, IAS'92, Houston, Texas, USA, 8-12 October 1992, Vol. 1, pp. 858-866.

[29] T. Liu, J. Fu, T.A. Lipo, 'A Strategy for Improving Reliability of Field Oriented Controlled Induction Motor Drives', IEEE Transactions On Industry Applications, Vol. 29, No. 5, September/October 1993, pp. 910-917.

[30] T. Elch-Heb, J.P. Hautier, 'Remedial Strategy for Inverter-Induction Machine System Faults Using Two-Phase Operation', IEE, Proceedings of the 5th European Conference on Power Electronics and Applications, EPE'93, Brighton, U.K., 13-16 September 1993, Vol. 5, pp. 151-156.

[31] T. Elch-Heb, E. Fan, J.P. Hautier, 'Reliability Improvement of Field-Oriented Controlled Three-Phase AC Drives By Means of Two-Phase Remedial Operation', International Conference on Electric Machines, ICEM'94, Paris, France, 5-8 September 1994, Vol. 2, pp. 194-198.

[32] T.A. Nyamusa, N.A. Demerdash, 'Transient Analysis of a Partial Armature Short Circuit in an Electronically Commutated Permanent Magnet Motor System Using an Integrated Nonlinear Magnetic Field-Network Model', IEEE Transactions on Energy Conversion, Vol. 2, No. 1, March 1987, pp. 86-92.

[33] A.K. Wallace, R. Spee, 'The Simulation of Brushless DC Drive Failures', IEEE, Proceedings of the 19th Annual Power Electronics Specialists Conference, PESC'88, Kyoto, Japan, 11-14 April 1988, Vol. 2, pp. 199-206.

[34] R. Spee, A.K. Wallace, 'Remedial Strategies for Brushless DC Drive Failures', IEEE Transactions on Industry Applications, vol. 26, no. 2, July/August 1996, pp. 259-266.

[35] A.G. Jack, B.C. Mecrow, 'Safety Critical Drives for Aerospace Applications', International Conference on Electric Machines, ICEM'94, Paris, France, 5-8 September 1994, Vol. 1, pp. 91-96.

[36] B.C. Mecrow, A.G. Jack, J.A. Haylock, J. Coles, 'Fault Tolerant Permanent Magnet Machine Drives', IEE, Proceedings of the 7th International Conference on Electric Machines and Drives, EMD'95, Durham, UK, 11-13 September 1995, pp. 433-437.

[37] A.G. Jack, B.C. Mecrow, and J. Haylock, ' A Comparative Study of Permanent Magnet and Switched Reluctance Motors for High Performance Fault Tolerant Applications', IEEE Transactions on Industry Applications, Vol. 32, No. 4, July/August 1996, pp. 889-895.

[38] A.G. Jack, B.C. Mecrow, and J. Haylock, ' A Comparative Study of Permanent Magnet and Switched Reluctance Motors for High Performance Fault Tolerant Applications', IEEE 30th Industry Applications Society Annual Meeting, IAS'95, Orlando, Florida, USA, 8-12 October 1995, Part 1 (of 3), pp. 734-740.

[39] B.C. Mecrow, A.G. Jack, J.A. Haylock, J. Coles, 'Fault Tolerant Permanent Magnet Machine Drives', IEE Proceedings, Part B : Electrical Power Applications, Vol. 143, No. 6, November 1996, pp. 437-442.

[40] J. A. Haylock, B. C. Mecrow, A. G. Jack and D. J. Atkinson 'Operation of Fault Tolerant Machines With Winding Failures', IEEE International Electric Machines and Drives Conference, , IEMDC'97, Milwaukee, Wisconsin, USA. May 18-21 1997, pp. MC3-10.1 - MC3-10.3.

[41] J.A. Haylock, B.C. Mecrow, A.G. Jack, D.J. Atkinson, 'Operation of a Fault Tolerant PM Drive for an Aerospace Fuel Pump Application', IEE, Eighth International Conference on Electrical Machines and Drives, Robinson College, Cambridge, UK, 1-3 September 1997.

[42] M. Ronkowski, K. Iwan, R. Szczesny, 'Simulation of Converter - Synchronous Machine Drives in Normal and Fault Operation', IEE, Proceedings of the 5th European Conference on Power Electronics and Applications, EPE'93, Brighton, U.K., 13-16 September 1993, Vol. 5, pp. 174-178.

[43] RE. Szczesny, P.J. Chrzan, K. Iwan, 'Modelling and Simulation of Electric Drives at Normal and Fault Conditions', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 3.580-3.585.

[44] N. Bianchi, S. Bolognani, M. Zigliotto, 'Analysis of PM Synchronous Motor Drive Failures During Flux Weakening Operation', IEEE, Proceedings of the

27th Annual Power Electronics Specialists Conference, PESC'96, Maggorie, Italy, 20 January 1996, Vol. 2, pp. 1542-1548.

[45] M.F. Rahman, M.A. Rahman, L. Zhong, W.Y. Hu, K.W. Lim, 'A Direct Torque Controller for Permanent Magnet Synchronous Motor Drives', IEEE International Conference on Electric Machines and Drives, IEMDC'97, Milwaukee, Wisconsin, USA. May 18-21 1997, pp. TD1-2.1 - TD1-2.3.

[46] A.C. Renfrew, J.X. Tian, 'The Use of a Knowledge-Based System in Power Electronics Circuit Fault Diagnosis', IEE, Proceedings of the 5th European Conference on Power Electronics and Applications, EPE'93, Brighton, U.K., 13-16 September 1993, Vol. 7, pp. 57-62.

[47] K. Debebe, V. Rajagopalan, T.S. Sankar, 'Diagnosis and Monitoring for AC Drives', IEEE 27th Industry Applications Society Annual Meeting, IAS'92, Houston, Texas, USA, 8-12 October 1992, Vol. 1, pp. 370-377.

[48] I.B. Aris, L. Zhang, L.N. Hulley, 'Fault Detection of an Inverter Circuit By Digital Signal Processing and Knowledge Based Approach', Proceedings of the 29th Universities Power Engineering Conference, UPEC'94, Galway, Ireland, 14 - 16 September 1994, Part 1 (of 2), pp. 489-492.

[49] L. Zhang, I.B. Aris, L.N. Hulley, 'A Knowledge Based System for On-Line Fault Diagnosis of Power Inverter Circuits for AC Machine Drive', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 3.334-3.339.

[50] E. Craig, B.C. Mecrow, D.J. Atkinson, A.G. Jack, 'Fault Detection Procedure For Single Phase Bridge Detectors', IEE, Proceedings of the 5th European Conference on Power Electronics and Applications, EPE'93, Brighton, U.K., 13-16 September 1993, Vol. 4, pp. 466-471.

[51] M.J. Cronin, 'The All Electric Aircraft as an Energy Efficient Transport', SAE, Aerospace Congress & Exposition, Los Angeles, 13-16 October 1980, Paper 801113.

[52] J.A. Weimer, 'Power Management and Distribution for the More Electric Aircraft', Proceedings of the 30th Intersociety Energy Conference, IECEC, Orlando, FL, USA, 30 July - 4 August 1995, Part 1 (of 3), pp. 273-277.

[53] S.J. Przybylko, 'Future Advanced Control Technology Study (FACTS) -- A Look at Emerging Technologies', AIAA/SAE/ASME/ASEE 23rd Joint Propulsion Conference, San Diego, 1987, Paper AIAA-87-1930.

[54] M.J. Cronin, 'The All Electric Airplane Revisited', Aerospace Power System Technology, 1988, SAE Publication SP758, Paper 881407, pp. 9-27.

[55] R.E. Quingley, 'More Electric Aircraft', 8th Annual Applied Power Electronics Conference and Exposition - APEC'93, San Diego, CA, USA, 7-11 March 1993, pp. 906-911.

[56] W.D. Jones, M.S. Jarvis, 'Electrically Driven Engine Controls and Accessories for Future Aircraft', American Helicopter Society, Rotary Wing Propulsion Specialists' Meeting.

[57] T.M. Jahns, R.C. Van Nocker, 'High Performance EHA Controls Using an Interior Permanent Magnet Motor', IEEE Transactions on Aerospace and Electronic Systems, Vol. 26, No. 3, May 1990, pp. 534-542.

[58] R. M. Crowder, 'Electrically Powered Actuation for Civil Aircraft Applications', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 3.793-3.798.

[59] 'Hydraulics Lift Design of Lightweight Aircraft', Drives and Controls magazine, September 1994, pp. 40.

[60] S.R. Jones, B.T. Drager, 'Performance of High Speed Switched Reluctance Starter/Generator System Using Electronic Position Sensing', IEEE 30th Industry Applications Society Annual Meeting, IAS'95, Orlando, Florida, USA, 8-12 October 1995, Part 1 (of 3), pp. 294-253.

[61] H.H. Belmont, 'Digital Fault Tolerant Flight Actuation Systems', Proceedings of the IEEE 1985 National Aerospace Electronics Conference, NAECON 1985, Dayton, OH, USA, 20-24 May 1985, pp. 618-627.

[62] K. Hofer, 'Redundant Drive Control By Observers', IEE, Proceedings of the 5th European Conference on Power Electronics and Applications, EPE'93, Brighton, U.K., 13-16 September 1993, vol. 4, pp. 330-335.

[63] R. Chokhawala, J. Catt, L. Kiraly, 'A Discussion on IGBT Short Circuit Behaviour and Fault Protection Schemes', IR IGBT Designers Manual [IGBT3] 1994, pp. E-117-E-126.

[64] B.K. Bose, D. Kastha, 'Electrolytic Capacitor Elimination in Power Electronic System by High Frequency Active Filter', IEEE 26th Industry Applications Society Annual Meeting, IAS'91, Dearborn, MI, USA, 28 September-04 October 1991, Part 1 (of 2), pp. 869-878.

[65] T.J.E. Miller, 'Converter Volt-Ampere Requirements of the Switched Reluctance Motor Drive', IEEE Transactions on Industry Applications, Vol. 21, No. 5, September/October 1985, pp. 1136-1144.

[66] T.J.E. Miller, 'Switched Reluctance Motors and their Control', Monographs in Electrical and Electronic Engineering, Magna Physics Publishing and Oxford Science Publications, 1993. ISBN 0-19-859387-2.

[67] C.R. Elliott, J.M. Stephenson, M.J. Turner, 'High Performance Control of Phase Current in the Switched Reluctance Motor', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 3.223-3.228.

[68] P.G. Barrass, 'High Performance Switched Reluctance Drives', PhD Thesis, University of Newcastle upon Tyne, UK, 1995.

[69] J. Lee, B. Kang, 'Digital Current Controller with Delay Compensator for PWM Converters', IEEE International Conference on Electric Machines and Drives, Milwaukee, Wisconsin, USA. 18-21 May 1997, pp. TB2-6.1 -TB2-6.3.

[70] E. Monmasson, J.C. Hapiot, M. Grandpierre, 'Analysis of a Current Controller for AC Drives Entirely Based on FPGAs', International Conference on Electric Machines, ICEM'94, Paris, France, 5-8 September 1994, Vol. 3, pp. 1-5.

[71] L.A. Belfore, A.A. Arkadan, 'Neurogenetic Models for the Characterisation of Fault Tolerant Switched Reluctance Motors', IEEE International Conference on Electric Machines and Drives, Milwaukee, Wisconsin, USA, 18-21 May 1997, pp. TB1-8.1 -TB1-8.3.

[72] C. French, 'Real Time Control of Electric Drives', PhD Thesis, University of Newcastle upon Tyne, UK, 1995.

[73] G.E. Franklin, J.D. Powell, A. Emami-Naeini, 'Feedback Control of Dynamic Systems', 3rd Edition, Addison-Wesley Publishing Company, 1994. ISBN 0-201-52747-2.

[74] G. Champenois, 'Effective Digital Torque Control Systems for Permanent Magnets Converter Fed Sinusoidal Synchronous Machines', IEEE 23rd Industry Applications Society Annual Meeting, IAS'88, Pittsburgh, PA, USA, 02-07 October 1988, Vol. 35, No. 6, pp. 259-265.

[75] D.A. Staton, R.P. Deodhar, W.L. Soong, T.J.E. Miller, 'Torque Prediction Using the Flux-MMF Diagram in AC, DC, and Reluctance Motors', IEEE Transactions on Industry Applications, Vol. 32, No. 1, January/February 1996, pp. 180-188.

[76] P.G. Barrass, B.C. Mecrow, 'Torque Control of Switched Reluctance Drives', International Conference on Electric Machines, ICEM'96, Vigo, Spain, 10-12 September 1996, pp. 254-259.

[77] R.L. Burden, J.D. Faires, 'Numerical Analysis', Fourth Edition, PWS-KENT Publishing Company, 1989, ISBN 0-534-91585-X.

[78] R.P. Deodhar, D.A. Staton, T.M. Jahns, T.J.E. Miller, 'Prediction of Cogging Torque using the Flux-MMF Diagram Technique', IEEE Transactions on Industry Applications, Vol. 32, No. 3, May/June 1996, pp. 569-575.

[79] T.M. Jahns, 'Flux-Weakening Regime Operation of an Interior Permanent-Magnet Synchronous Motor Drive', IEEE Transactions On Industry Applications, Vol. 23, No. 4, July/August 1987, pp. 681-689.

[80] A. Ludbrook, M. Ehsani, 'Burndown Prevention in Static Power Converter Equipment', IEEE Industry Applications Magazine, March/April 1995, pp. 46-55.

[81] P. Calonnec, T. Derrey, E. Destobbeleer, L. Protin, 'Induction Motors Reliability : Use of a Fault Tree', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 3.281-3.285.

[82] C.S. Moo, S.Y. Chan, Y.C. Chunag, Y.N. Chang, 'Investigation on Fault Operation of Full Converter Fed DC Drives', IEEE 30th Industry Applications Society Annual Meeting, IAS'95, Orlando, Florida, USA, 8-12 October 1995, Part 1 (of 3), pp. 3-8.

[83] J. Penman, H.G. Sedding, B.A. Lloyd, W.T. Fink, 'Detection and Location of Interturn Faults in the Stator Windings of Operating Motors', IEEE Transactions on Energy Conversion, Vol. 9, No. 4, December 1994, pp. 652-658.

[84] J.A. Wagner, 'The Shorted Turn in the Linear Actuator of a High Performance Disk Drive', IEEE Transactions On Magnetics, Vol. 18, No. 6, November 1982, pp. 1770-1772.

[85] S. Duong, S. Rael, C. Schaeffer, J. De Palma, 'Short Circuit Behaviour for PT and NPT IGBT Devices - Protection Against Explosion of the Case by Fuses', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 1.249-1.254.

[86] R.R. Schoen, T.G. Gabetler, 'A New Method of Current-Based Condition Monitoring in Induction Machines Operating Under Arbitrary Load Conditions', International Conference on Electric Machines, ICEM'94, Paris, France, 5-8 September 1994, Vol. 2, pp. 282-287.

[87] P.D. Evans, B.M. Saied, 'Protection Methods For Power Transistor Circuits', IEE Proceedings, Part B : Electric Power Applications, Vol. 129, No. 6, November 1982, pp. 359-363.

[88] A.S. Carvalho, A.P. Martins, A.S. Araujo, 'Improving IGBT Behaviour on Short Circuit Condition', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 1.831-1.835.

[89] G. Castino, A. Dubashi, S. Clemente, B. Pelly, 'Protecting IGBTs Against Short Circuit', IR Publication IGBT DESIGNERS MANUAL (IGBT2) 1991.

[90] H.-G. Eckel, L. Sack, 'Optimisation of the Short-Circuit Behaviour of NPT-IGBT By Gate Drive', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 2.213-2.218.

[91] R. Chokhawala, S. Sobhani, 'Switching Voltage Transient Protection Schemes for High Current IGBT Modules', IR IGBT Designers Manual [IGBT3] 1994, pp. E-145-E-154.

[92] S. Gediga, R. Marquawrdt, R. Sommer, 'High Power IGBT Converters with New Gate Drive and Protection Circuit', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 1.066-1.070.

[93] R.J. Valentine, 'Power Module Control Design', IEEE 30th Industry Applications Society Annual Meeting, IAS'95, Orlando, Florida, USA, 8-12 October 1995, Part 1 (of 3), pp. 904-910.

[94] Semikron International, 'SKHI-Drivers, The Right Way to Control IGBT/MOS Power', Semikron International Dr Fritz Martin GmbH & Co. D-90253 Nurnberg.

[95] Y.C. Liang, V. Hor, 'Design Of Lateral IGBT Power Devices with Current Sensor', IEEE 30th Industry Applications Society Annual Meeting, IAS'95, Orlando, Florida, USA, 8-12 October 1995, Part 1 (of 3), pp. 1010-1015.

[96] Z.J. Shen, S.P. Robb, A. Taomoto, 'Current Sensing Characteristics of IGBTs under Short Circuit Conditions', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 2.202-2.206.

[97] Concept Technology Ltd., 'CONCEPT IHD 215/280/680 Datemblatt, Intelligenter Halbbrucken-Treiber fur IGBT und Power-MOSFET', Data Sheet Ref. IHDDAT11 12/1994 © 1992, 1993, 1994 CT-Concept Technology Ltd, Switzerland.

[98] M. Nicolaidis, S. Noraz, B. Courtois, 'A Generalised Theory of Fail Safe Systems', IEEE 19th International Symposium on Fault-Tolerant Computing, Chicago, IL, USA, 21-23 June 1989, pp. 398-406.

[99] E. Woschnagg, 'Turbogenerator Field Winding Shorted Turn Detection by AC Flux Measurement', IEEE Transactions on Energy Conversion, Vol. 9, No. 4, June 1984, pp. 427-431.

[100] K. Weinreb, P. Drozdowski, 'Detection of Winding Faults of a Salient Pole Synchronous Machine by a Spectral Analysis of Currents', International Conference on Electric Machines, ICEM'94, Paris, France, 5-8 September 1994, Vol. 2, pp. 56-61.

[101] P.L. Walstrom, 'Shorted Turn Detection in the Winding and Operation of Super Conducting Magnets', IEEE Transactions on Magnetics, Vol. 19, No. 3, May 1983, pp. 1094-1096.

[102] T.H. Lui, 'A New Control Method and Reliability Improvement for a Switched Reluctance Motor Drive', Electric Power Systems Research, Vol. 37, No. 1, April 1996, pp. 7-18.

[103] J. A. Haylock, B. C. Mecrow, A. G. Jack and D. J. Atkinson, 'Operation of Fault Tolerant Machines with Winding Failures', IEEE Transactions on Energy Conversion, To be published, 1998.

[104] S. Januszewski, M. Kociszewska-Szczerbik, H. Swaitek, G. Swaitek, 'Causes and Mechanisms of Semiconductor Device Failures in Power Converter Service Conditions', 6th European Conference on Power Electronics and Applications, EPE'95, Sevilla, Spain, 19-21 September 1995, pp. 1.625-1.630.

[105] L. Chen, J.C. Balda, K.J. Olejniczak, 'Implementation of a High Performance Induction Motor Drive Using a Three Phase Inverter Module, a PWM ASIC and a DSP', IEEE 30th Industry Applications Society Annual Meeting, IAS'95, Orlando, Florida, USA, 8-12 October 1995, Part 1 (of 3), pp. 217-223.