

High Performance Switched Reluctance Drives

Peter Guy Barrass

NEWCASTLE UNIVERSITY LIBRARY

095 51172 3

Thesis L5580

A thesis submitted for the degree of
Doctor of Philosophy

© September, 1995

**Department of Electrical Engineering
University of Newcastle upon Tyne**

Abstract

The fully-pitched winding arrangement is one of the most radical changes in the design of doubly-salient reluctance motors in recent times. By replacing conventional short-pitched windings with fully-pitched windings, the resulting machine has a strong and position dependant mutual coupling between phases. The major torque producing mechanism is due to changes in mutual inductance with rotor position. This enables the windings to be better utilised, and with correct selection of excitation all phases can contribute useful torque all of the time. The increased winding utilisation requires a lower MMF per phase in comparison with a short-pitch wound machine with a single phase excited. Given a suitable winding configuration and machine dimensions, the copper losses for a given torque can be significantly lower than an equivalent conventional switched reluctance machine.

Operation of a three phase fully-pitched winding switched reluctance machine has been studied theoretically, in simulation and experimentally. The experimental drive comprises of a D132 frame 12:8 machine, IGBT power converter and DSP controller.

Operation with unipolar phase currents has been investigated over a wide speed range and performance compared with a conventional switched reluctance machine. Bipolar operation with several different excitation patterns has been investigated. Unipolar operation gives the largest torque/speed envelope with a simple controller, although bipolar modes can equal this with a more complex controller. Results show that for equal RMS phase current the average torque produced by four different modes of excitation are approximately equal. However, there is a large difference in the torque ripple and acoustic noise performance of each mode.

Current control in switched reluctance machines is complicated by the non-linear nature of the load. By controlling flux-linkage rather than current a linear load model can be used. A discrete time 'dead-beat' flux-linkage controller has been implemented which gives superior phase current control performance to other types of controller with the same sample interval.

A new method of constant torque operation based on 'flux ramps' has been proposed. This method gives predictable performance and enables constant torque operation over a wide speed range. A Genetic Algorithm has been shown to be very effective when applied to the problem of optimising the 'flux ramps' for minimum torque ripple. A speed controller has been implemented which makes use of the Genetic Algorithm optimised flux ramps to give smooth torque over a wide speed range.

Acknowledgements

Well I think the biggest thanks should go to Barrie because if he hadn't come up with the fully-pitched winding idea, or pointed me in the right direction now and then, then all this would never have happened.

Life would be very dull in the UG Lab if it wasn't for the people in it. There have been many passing through but the most important to me are Chris (I'm never drinking again), Andy (You did remember to tell Mandy didn't you?), Howard (He's done a Howard) and Jim H. (Food! - What an excellent idea). Other notable inmates are Jim K., Brian, Alun, Gav, Ken, Bernard, Volker, Roger, Husain, Adil, Jawad, Ibrahim and Anthony. And those who have mounted a successful escape bid, Nesimi, Peter, Teresa, Eugene and Sab.

Outside the lab thanks go to Hilary for sticking with the long distance travelling, and the people I have been fortunate (or unfortunate) enough to share a house with - Chris, Mark and Julian.

Down in the UG Lab Dave always managed to find enough meters for me and would always go out of his way to be helpful - like the time we brought a large transformer down several floors without the aid of the lift. In the electronics workshop thanks go to Tom for ordering the bits and Brian and Steve for making the boards for them to go on. In the mechanical workshop biggest thanks must go to Steve and Charlie who spent many an hour putting the machine rig together and straightening out the bend. Thanks also to Allen and Les who have been a great help.

On the Thesis writing front thanks go to Barrie for general guidance, my Mum for teaching me where and when to use punctuation (its only taken me 30 years to find out), Chris for going through the pain of proof reading the thing, and lastly to my trusty PC for not crashing the hard-disk half way through.

Like most things in life these days money comes into the equation at some point, and nothing would happen without it. Those who enabled it to happen for me are my Mum and Dad for the first 18 years, and the EPSRC for the last three years. Thanks also go to the Royal Society who came up with some money for a few bits of kit.

In moving on to the next stage of my career thanks go to Alan and Paul for putting in a good word for me at a couple of well known drives companies.

Table of Contents

Abstract	ii
Acknowledgements	iii
Table of Contents.....	iv
List of Figures	x
List of Tables.....	xviii
Chapter 1 - Introduction	
1.1 Objectives.....	1.2
1.2 The Contribution to Knowledge	1.2
1.3 Overview of the Thesis	1.3
1.4 Electrical Machine Drives.....	1.4
1.4.1 Single Speed Drives.....	1.4
1.4.2 Variable Speed Drives	1.4
1.4.3 Servo Drives.....	1.5
1.5 The state of the art in Switched Reluctance Drives	1.6
1.5.1 Fundamental Machine Parameters	1.6
1.5.2 Applications	1.9
1.5.3 Machine Winding Configurations.....	1.11
1.5.4 Power Converters.....	1.12
1.5.5 Control of Switched Reluctance Drives.....	1.17
1.5.6 Improving Drive Performance and Overcoming Prejudices.....	1.19
1.5.7 Switched Reluctance Drives in the Market Place	1.23
1.6 Conclusions.....	1.25
Chapter 2 - Fully-Pitched Winding Switched Reluctance Machines	
2.1 What is a Fully-pitched Winding Switched Reluctance Machine	2.1
2.1.1 The Mechanical and Winding Arrangement.....	2.1
2.1.2 Self and Mutual Inductances.....	2.3
2.1.3 The Method of Torque Production	2.6
2.1.4 Phase Conduction Sequences.....	2.7
2.2 The Equivalent Machine Transformations.....	2.12

2.3	Comparison of Fully-Pitched Winding and Single-Tooth Winding Switched Reluctance Machines	2.17
2.3.1	Phase Winding Lengths	2.17
2.3.2	Theoretical Copper and Torque Density	2.20
2.3.3	The Experimental Machine Copper and Torque Density	2.22
2.3.4	Power Converter Rating.....	2.24
2.4	Summary and Conclusions.....	2.28

Chapter 3 - Drive Implementation

3.1	Implementation of the Experimental Drive.....	3.1
3.1.1	Experimental machines and loading arrangement	3.1
3.1.2	Control Systems and Power Electronics	3.3
3.1.3	Drive Control Software.....	3.7
3.2	Summary and Conclusions.....	3.11

Chapter 4 - Current Control in Switched Reluctance Drives

4.1	Current Control of Switched Reluctance Machines.....	4.1
4.1.1	Current Control Methods	4.1
4.1.2	Current Controller System Model.....	4.4
4.1.3	PID Current Controllers.....	4.5
4.1.4	Transformed Current Controllers.....	4.8
4.2	Summary and Conclusions.....	4.10

Chapter 5 - Unipolar Operation

5.1	Theory of Operation.....	5.1
5.1.1	The Method of Torque Production	5.1
5.1.2	Phase Current Reference	5.3
5.2	Simulation Studies	5.4
5.2.1	Current Control at Low Speed	5.5
5.2.2	Voltage Control at High Speed	5.7
5.3	Experimental Operation of the Fully-Pitched Winding Switched Reluctance Drive.....	5.12
5.3.1	Low Speed Operation	5.13
5.3.2	High Speed Operation.....	5.17
5.3.3	Power Converter Rating.....	5.21
5.4	Torque/Speed Characteristics.....	5.22
5.4.1	Operation with Series Connected Phase Coils.....	5.23

5.4.2	Operation with Parallel Connected Phase Coils	5.26
5.5	Operation of the Single-tooth Winding Switched Reluctance Drive	5.26
5.5.1	Experimental waveforms	5.26
5.5.2	Torque/speed Characteristics	5.28
5.5.3	Comparison with Fully-Pitched Winding SR Machine	5.29
5.6	Summary and Conclusions.....	5.30

Chapter 6 - Bipolar Operation

6.1	Theory of Operation.....	6.1
6.1.1	The Method of Torque Production	6.1
6.1.2	Phase Current Reference.....	6.4
6.2	Simulation Studies	6.5
6.2.1	Trapezoidal Operation	6.5
6.2.2	Squarewave Operation.....	6.6
6.2.3	Sinusoidal Operation.....	6.8
6.2.4	Comparison of Simulated Low Speed Performance	6.9
6.3	Drive Operation with Separately Connected Phases.....	6.10
6.3.1	Low Speed Operation	6.10
6.3.2	High Speed Operation.....	6.14
6.3.3	Operation with Novel Current Waveforms.....	6.17
6.4	Torque/Speed Characteristics.....	6.18
6.5	Star Connected Operation	6.20
6.5.1	Experimental Drive Operation.....	6.20
6.5.2	Torque/speed Characteristics	6.22
6.6	Delta Connected Operation	6.23
6.7	Summary and Conclusions.....	6.24

Chapter 7 - Flux Control of Switched Reluctance Machines

7.1	Theory of operation.....	7.1
7.1.1	Phase Current Control Methods.....	7.1
7.1.2	A Dead Beat Flux Controller	7.3
7.1.3	Implementation of a Flux-Linkage Controller for Single-Tooth Winding Machines	7.7
7.1.4	Implementation of a Flux-Linkage Controller for Fully-Pitched Winding Machines.....	7.9
7.1.5	The Effect of Position Errors on Flux-Linkage Control	7.10

7.2	Experimental Results	7.12
7.2.1	Comparison of Flux-Linkage and Current Controller Responses.....	7.12
7.2.2	Flux-Linkage Control at Low Speed.....	7.14
7.2.3	Flux-Linkage Control at High Speed.....	7.17
7.3	Advanced Switched Reluctance Drives Using Flux-Linkage Control.....	7.18
7.3.1	Calculation of Flux References for High Speed Operation	7.18
7.3.2	Improving High Speed Torque Capability.....	7.21
7.4	Summary and Conclusions.....	7.23

Chapter 8 - Constant Torque Operation

8.1	Theory of Operation.....	8.1
8.2	Review of Constant Torque Methods	8.2
8.2.1	Basis of Methods	8.3
8.2.2	Simulation.....	8.5
8.2.3	Experimental Comparison	8.9
8.3	Simulation	8.14
8.3.1	Current Control for Constant Torque.....	8.14
8.3.2	Voltage Control for Constant Torque	8.15
8.4	Flux Ramps for Constant Torque.....	8.17
8.4.1	Theory of Operation.....	8.18
8.4.2	Flux-linkage Profiles for Constant Torque	8.18
8.4.3	Simulation.....	8.21
8.4.4	Systematic Search	8.23
8.4.5	Genetic Algorithm Search.....	8.23
8.4.6	An Array of Optimised Flux Ramps.....	8.28
8.4.7	Experimental Results	8.29
8.5	A Speed Controller.....	8.33
8.5.1	System Model and Controller	8.34
8.5.2	Experimental Results	8.37
8.6	Summary and Conclusions.....	8.38

Chapter 9 - Comparison of Operating Modes

9.1	Torque Speed Characteristics.....	9.1
9.2	Sensitivity to Control Angles.....	9.2
9.3	Torque Ripple and Torque per Unit Loss	9.4

9.4	Acoustic Noise	9.6
9.4.1	Measurement of Acoustic Noise	9.6
9.4.2	Variation of Noise with Speed	9.7
9.4.3	Variation of Noise with Supply Voltage	9.8
9.4.4	Variation of Noise with Current and Torque Waveform	9.9
9.5	Summary and Conclusions.....	9.10

Chapter 10 - Conclusions

10.1	Switched Reluctance Machines.....	10.1
10.2	Operation with Current and Voltage Control.....	10.2
10.3	Advanced Drive Control	10.3
10.4	Further Work.....	10.5

Appendix A - Switched Reluctance Machine Characteristics

A.1	The Experimental Machines	A.1
A.2	Physical Dimensions and Windings.....	A.1
A.3	Flux-Linkage Characteristics	A.2
A.3.1	Measurement of Flux-Linkage Characteristics	A.2
A.3.2	Fully-Pitched Winding Switched Reluctance Machine Characteristics.....	A.5
A.3.3	Single-Tooth Winding Switched Reluctance Machine Characteristics.....	A.7
A.4	Static Torque Characteristics	A.8
A.5	Calculation of Flux-Linkage from Torque.....	A.9

Appendix B - Machine Loading Rig Characteristics

B.1	Physical Arrangement	B.1
B.2	DC Machine Characteristics	B.2
B.3	Calculation of Shaft Torque	B.5
B.4	Speed Controller Load Model	B.8

Appendix C - Electrical and Electronic Circuits

C.1	Load and Measurement Circuit.....	C.1
C.2	Power Converter.....	C.3
C.2.1	Input Rectifier and DC-Link Filter	C.3
C.2.2	Bridge Legs.....	C.4
C.2.3	DC-Link Dump	C.6
C.2.4	Auxiliary Power Supplies	C.7

C.3	Measurement Circuits	C.8
C.3.1	Rotor Position Measurement	C.8
C.3.2	Phase Current Measurement	C.9
C.3.3	Over-Current Detection	C.10
C.3.4	DC-Link Voltage Measurement.....	C.10
C.4	Controller Circuits.....	C.11
C.4.1	Front Panel Controls and Analogue Input	C.11
C.4.2	Analogue Output.....	C.11
C.4.3	Phase Current Input.....	C.12
C.4.4	Encoder Input and System Control	C.13
C.4.5	TMS320E14 Processor Board.....	C.15
C.4.6	Commercially Produced Boards	C.17

Appendix D - Controller Software

D.1	Processor Configurations	D.1
D.1.1	Two Processor Configuration	D.1
D.1.2	Three Processor Configuration	D.2
D.2	Lookup Tables.....	D.3
D.3	Parameter Monitoring	D.4
D.4	Reference Calculation	D.5
D.5	Current Controllers	D.6
D.6	Flux-linkage Controllers	D.8
D.7	Speed Measurement	D.9
D.8	Speed Controllers.....	D.10
D.9	Random Number Calculation.....	D.11

Appendix E - An Introduction to Genetic Algorithms

E.1	Basics.....	E.1
E.2	Terms and Search Strategy.....	E.1
E.3	Theorems.....	E.3

Appendix F - List of Symbols

F.1	Abbreviations	F.1
F.2	Matrix and Vector Symbols	F.1
F.3	General Symbols	F.2

References and Publications	xix
--	------------

List of Figures

Figure 1.1	Typical torque characteristics of three and four phase machines.	1.8
Figure 1.2	Cross section of a five phase 10:8 switched reluctance machine showing flux paths.....	1.9
Figure 1.3	Converter circuits with dissipative energy dumps.	1.14
Figure 1.4	Standard two-quadrant converter circuits.	1.15
Figure 1.5	Novel converter circuits.	1.16
Figure 2.1	View of the end-windings in a fully-pitched winding and single-tooth winding switched reluctance machine..	2.2
Figure 2.2	Unrolled machine showing winding arrangements and a single phase excited.	2.3
Figure 2.3	Machine cross section showing tooth overlap for inductance estimation.....	2.4
Figure 2.4	Idealised self and mutual inductances for a three phase 6:4, fully-pitched winding machine.	2.5
Figure 2.5	Derivation of torque for a single-tooth winding machine with unipolar phase currents.	2.8
Figure 2.6	Magnetic flux plot for a fully-pitched winding machine with two phase on unipolar excitation.	2.8
Figure 2.7	Derivation of torque for a fully-pitched winding machine with unipolar two phase on excitation.	2.9
Figure 2.8	Magnetic flux plot for a fully-pitched winding machine with two phase on bipolar excitation.	2.10
Figure 2.9	Derivation of torque for a fully-pitched winding machine with bipolar two phase on excitation.	2.10
Figure 2.10	Magnetic flux plot for a fully-pitched winding machine with three phase on bipolar excitation.	2.11
Figure 2.11	Derivation of torque for a fully-pitched winding machine with bipolar three phase on excitation.	2.11
Figure 2.12	Fully-pitched winding machine cross section showing winding arrangement and flux.	2.14
Figure 2.13	Single-tooth winding machine cross section showing winding arrangement and flux.	2.14
Figure 2.14	Estimation of machine phase winding lengths.....	2.18
Figure 2.15	Comparison of simplified voltage and current waveshapes at base speed.	2.26

Figure 3.1	Experimental machine loading arrangement.....	3.2
Figure 3.2	Arrangement of control and power equipment.	3.4
Figure 3.3	Power converter circuit arrangement.	3.6
Figure 3.4	Control signals and processor bus arrangement.....	3.8
Figure 3.5	Current control software function block diagram.	3.9
Figure 4.1	Simplified hysteresis current controller.	4.2
Figure 4.2	Simplified constant off time current controller.....	4.3
Figure 4.3	Simplified PID current controller.	4.4
Figure 4.4	Closed loop current control system block diagram.....	4.5
Figure 4.5	PID controller response without anti-windup.....	4.7
Figure 4.6	PID controller response with anti-windup.	4.7
Figure 4.7	PID controller response with adaptive gain setting.	4.8
Figure 4.8	Transformed current controller block diagram.	4.9
Figure 4.9	Comparison of real and transformed current controllers.	4.10
Figure 5.1	Derivation of torque from fully-pitched winding parameters.....	5.1
Figure 5.2	Fully-pitched winding currents and equivalent single-tooth winding currents.....	5.2
Figure 5.3	Derivation of torque from single-tooth winding parameters.	5.3
Figure 5.4	Phase current reference waveform showing control angles and current demand level.	5.4
Figure 5.5	Real fully-pitched winding phase current - I_a , and equivalent single-tooth winding phase current - I_1	5.5
Figure 5.6	Real fully-pitched winding flux-linkage - ψ_a , and equivalent single-tooth winding flux-linkage - ψ_1	5.6
Figure 5.7	Total shaft torque -Torque, and torque contribution of equivalent single-tooth winding phase 1 - T_1	5.6
Figure 5.8	Voltage control of the fully-pitched winding switched reluctance machine. Phase voltage - V_a , and equivalent single-tooth winding phase voltage - V_1	5.8
Figure 5.9	Voltage control of the single-tooth winding switched reluctance machine. Phase voltage V_1 , and equivalent fully-pitched winding phase voltage V_a	5.8
Figure 5.10	Real and equivalent single-tooth winding flux-linkage. Voltage control of the single-tooth winding machine - ψ_{stw} , and voltage control of the fully-pitched winding machine - ψ_{fpw}	5.9

Figure 5.11	Real and equivalent single-tooth winding phase current. Voltage control of the single-tooth winding machine - I_{stw} , and voltage control of the fully-pitched winding machine - I_{fpw}	5.9
Figure 5.12	Total shaft torque. Voltage control of the single-tooth winding machine - T_{stw} , and voltage control of the fully-pitched winding machine - T_{fpw}	5.10
Figure 5.13	Equivalent single-tooth winding flux-linkage/current locus for voltage control of the fully-pitched winding machine.	5.10
Figure 5.14	Single-tooth winding flux-linkage/current locus for voltage control of the single-tooth winding machine.	5.11
Figure 5.15	Low speed real and equivalent current waveforms.....	5.14
Figure 5.16	Low speed fully-pitched winding ψ, i locus.....	5.14
Figure 5.17	Low speed equivalent single-tooth winding ψ, i locus.....	5.15
Figure 5.18	Low speed equivalent single-tooth winding current, flux-linkage, and torque.....	5.16
Figure 5.19	The effects of sub-optimum conduction widths at low speed.....	5.17
Figure 5.20	High speed real and equivalent current waveforms.	5.18
Figure 5.21	High speed real and equivalent current, voltage, and flux-linkage waveforms.	5.19
Figure 5.22	High speed fully-pitched winding ψ, i locus.	5.20
Figure 5.23	High speed equivalent single-tooth winding ψ, i locus.	5.20
Figure 5.24	Four quadrant torque/speed envelope for $I_{dem} = 10A$	5.23
Figure 5.25	Four quadrant shaft power envelope for $I_{dem} = 10A$	5.24
Figure 5.26	Torque/speed envelope for $I_{dem} = 5A$ and $10A$	5.25
Figure 5.27	Shaft power envelope for $I_{dem} = 5A$ and $10A$	5.25
Figure 5.28	Phase current at low and medium speed.	5.27
Figure 5.29	High speed operation, 1350rpm, 28Nm.	5.28
Figure 5.30	Torque/speed envelopes for $I_{dem} = 20A$, conduction widths unlimited and limited to 120°	5.28
Figure 5.31	Torque/speed envelopes for the experimental fully-pitched winding and single-tooth winding switched reluctance drives	5.29
Figure 6.1	Derivation of torque from fully-pitched winding parameters with bipolar two phase on excitation.	6.2
Figure 6.2	Derivation of torque from single-tooth winding parameters with bipolar two phase on excitation.	6.2
Figure 6.3	Derivation of torque from fully-pitched winding parameters with bipolar three phase on excitation.	6.3

Figure 6.4	Derivation of torque from single-tooth winding parameters with bipolar three phase on excitation.	6.4
Figure 6.5	Phase current reference waveform showing control angles and current demand level.	6.4
Figure 6.6	Simulated real and equivalent single-tooth winding parameters for trapezoidal phase currents.	6.6
Figure 6.7	Simulated real and equivalent single-tooth winding parameters for squarewave phase currents.	6.7
Figure 6.8	Simulated real and equivalent single-tooth winding parameters for sinusoidal phase currents.	6.8
Figure 6.9	Low speed current and torque for trapezoidal conduction.	6.11
Figure 6.10	Low speed flux-linkage/current locus for trapezoidal conduction.	6.11
Figure 6.11	Low speed current and torque for squarewave conduction.	6.12
Figure 6.12	Low speed flux-linkage/current locus for squarewave conduction.	6.13
Figure 6.13	Low speed current and torque for sinusoidal phase currents.	6.13
Figure 6.14	Low speed flux-linkage/current locus for sinusoidal phase current.	6.14
Figure 6.15	High speed current, voltage and flux for bipolar conduction, 12Nm.	6.15
Figure 6.16	Derivation of bipolar voltage control waveform from unipolar voltage control.	6.16
Figure 6.17	High speed current, voltage and flux for bipolar conduction, 15Nm.	6.16
Figure 6.18	Waveform derived from bipolar 120° single-tooth winding current.	6.17
Figure 6.19	Low speed current, flux and torque, for bipolar 180° conduction equivalent single-tooth winding current.	6.18
Figure 6.20	Torque/speed characteristics for 10A peak phase currents.	6.19
Figure 6.21	Torque/speed characteristics for 10A RMS phase currents.	6.19
Figure 6.22	Shaft power characteristics for 10A RMS phase currents.	6.20
Figure 6.23	Low speed current and torque for star connected trapezoidal conduction.	6.21
Figure 6.24	High speed current and voltage for star connection and sinusoidal phase currents.	6.22
Figure 6.25	Torque/speed characteristics for star connected phases.	6.22
Figure 6.26	Low speed operation with delta connected phases and sinusoidal line currents.	6.23

Figure 7.1	Unity feedback system model.....	7.3
Figure 7.2	Block diagram of the motor flux control system.	7.4
Figure 7.3	Motor flux controller transfer function block diagram.	7.5
Figure 7.4	Simplified discrete time model of motor flux control system.	7.7
Figure 7.5	Flux-linkage controller for a single-tooth winding switched reluctance machine using current reference waveforms.	7.8
Figure 7.6	Flux-linkage controller for a single-tooth winding switched reluctance machine using flux reference waveforms.	7.8
Figure 7.7	Flux-linkage controller for a fully-pitched winding switched reluctance machine using single-tooth winding flux reference waveforms.	7.9
Figure 7.8	Flux-linkage reference position offset. Phase current I_a for nominal 240° unipolar conduction.....	7.10
Figure 7.9	Flux-linkage reference position offset. Phase current I_a for nominal sinusoidal conduction.	7.11
Figure 7.10	Comparison of measured phase current controller responses.	7.13
Figure 7.11	Flux-linkage controller operation at 90rpm 50Nm, showing real and equivalent single-tooth winding parameters.	7.15
Figure 7.12	Flux-linkage controller operation at 90rpm 50Nm, showing real current and flux-linkage.	7.15
Figure 7.13	Flux-linkage controller operation at 90rpm 50Nm, showing phase voltage and flux-linkage.	7.16
Figure 7.14	Flux-linkage controller operation on the limit of control at 560rpm.	7.17
Figure 7.15	Flux-linkage controller operation beyond the limit of control at 680rpm.	7.18
Figure 7.16	Flux-linkage controller response for a constant flux ramp reference.....	7.20
Figure 7.17	Flux-linkage controller with flux ramp reference at 730rpm 32Nm.	7.21
Figure 7.18	Maximum flux-linkage/current loci for a given peak current and peak to peak flux excursion.....	7.22
Figure 7.19	Flux-linkage loci for the experimental single-tooth winding machine with flux-linkage offsets of 0.0Wb(a) and 0.1Wb(b).	7.22
Figure 8.1	Constant current and constant torque machine characteristics.	8.2
Figure 8.2	Simulated constant torque operation, constant torque ramp method.	8.6
Figure 8.3	Simulated constant torque operation, constant current ramp method.	8.7

Figure 8.4	Simulated constant torque operation, minimum RMS current method.	8.8
Figure 8.5	Constant torque operation at 30Nm 57rpm.....	8.10
Figure 8.6	Constant torque operation equivalent single-tooth winding flux-linkage/current loci.....	8.11
Figure 8.7	Current flux-linkage and torque for constant torque operation at 30Nm 57rpm..	8.13
Figure 8.8	Three dimensional representation of voltage vectors.	8.16
Figure 8.9	Machine 30Nm flux-linkage profile and flux-linkage characteristics for various 30Nm constant torque control methods.....	8.19
Figure 8.10	Constant torque flux profiles showing the effects of controller saturation.....	8.19
Figure 8.11	Flux-linkage profile composed of constant flux ramps.	8.20
Figure 8.12	Simulated constant flux ramp operation.	8.22
Figure 8.13	Genetic Algorithm flow chart.	8.24
Figure 8.14	Constant torque flux ramps for low speed.	8.30
Figure 8.15	Constant torque flux ramps for medium and high speed.	8.31
Figure 8.16	Operation at 110rpm and 50Nm with GA optimised flux ramps.....	8.31
Figure 8.17	Motoring at 140rpm and 15Nm with GA optimised flux ramps.....	8.32
Figure 8.18	Generating at 530rpm and 20Nm with GA optimised flux ramps.....	8.33
Figure 8.19	Speed control system.	8.33
Figure 8.20	Desired speed controller gain characteristics.....	8.35
Figure 8.21	Speed control system block diagram.	8.36
Figure 8.22	Open loop system gain and phase response.....	8.36
Figure 8.23	Speed controller response to demand and load step changes.....	8.37
Figure 9.1	Motoring torque speed envelopes for 10A RMS phase currents.....	9.1
Figure 9.2	Simulated data showing the variation in torque with advance angle.....	9.2
Figure 9.3	Measured variation in average torque with advance angle for unipolar and sinusoidal operation.	9.3
Figure 9.4	Measured variation in average torque with pulse width and advance angle.....	9.3
Figure 9.5	Torque per RMS ampere for standard modes.	9.4

Figure 9.6	Peak to peak torque ripple for standard modes.....	9.4
Figure 9.7	Torque per RMS ampere for standard and constant current modes.....	9.5
Figure 9.8	Peak and RMS current for standard and constant torque modes giving 30Nm.....	9.5
Figure 9.9	Experimental arrangement for acoustic noise measurements.	9.6
Figure 9.10	Variation in acoustic noise with speed, unloaded, and loaded with unipolar and sinusoidal currents.	9.7
Figure 9.11	Variation in acoustic noise with dc-link voltage for unipolar operation.	9.8
Figure 9.12	Acoustic noise for different 10A RMS current waveforms.	9.9
Figure 9.13	Acoustic noise for different current waveforms each producing approximately 30Nm.	9.10
Figure A.1	Flux-linkage measurement circuit.....	A.3
Figure A.2	Flux-linkage testing at the aligned position.	A.4
Figure A.3	Flux-linkage testing at the aligned position.	A.4
Figure A.4	Three dimensional plot of fully-pitched winding switched reluctance machine flux-linkage characteristics.	A.5
Figure A.5	Fully-pitched winding switched reluctance machine flux-linkage characteristics for constant rotor position.	A.6
Figure A.6	Fully-pitched winding switched reluctance machine flux-linkage characteristics for constant phase current.	A.6
Figure A.7	Single-tooth winding switched reluctance machine flux-linkage characteristics for constant rotor position.	A.7
Figure A.8	Three dimensional plot of fully-pitched winding switched reluctance machine torque characteristics for two phase on unipolar excitation.	A.8
Figure A.9	Fully-pitched winding switched reluctance machine torque characteristics for constant phase current. Two phase on unipolar excitation.	A.8
Figure A.10	Comparison of electrical and mechanical energy.	A.9
Figure B.1	Mounting arrangement of experimental SR, and DC load machines.	B.1
Figure B.2	DC load machine torque constant at low field currents.	B.3
Figure B.3	DC load machine torque constant at high field currents.	B.3
Figure B.4	Machine test rig rotating losses.....	B.5
Figure B.5	Torque calculation circuit parameters.....	B.6
Figure C.1	Load and measurement circuit.	C.1

Figure C.2	Power converter block diagram.....	C.3
Figure C.3	Input rectifier and dc-link filter.....	C.4
Figure C.4	Converter bridge leg circuit.....	C.5
Figure C.5	DC-Link Dump Circuit.....	C.7
Figure C.6	Auxiliary power supplies.....	C.8
Figure C.7	Rotor Position Encoder.....	C.8
Figure C.8	Phase current measurement circuit.....	C.9
Figure C.9	Phase over-current detection circuit.....	C.10
Figure C.10	DC-Link Voltage Divider.....	C.10
Figure C.11	Front panel controls and DC-link voltage input.....	C.11
Figure C.12	Analogue output circuit.....	C.12
Figure C.13	Phase current input circuit.....	C.13
Figure C.14	Position encoder input and system control sub-unit.....	C.14
Figure C.15	TMS320E14 Memory interface circuit.....	C.15
Figure C.16	TMS320E14 Control and PWM Generation.....	C.16
Figure D.1	Two control processor configuration.....	D.1
Figure D.2	Three control processor configuration.....	D.2

List of Tables

Table 2.1	Machine winding length ratios.....	2.19
Table 2.2	Measured experimental machine torque ratios for equal loss.....	2.23
Table 5.1	Comparison of voltage control for fully-pitched winding and single-tooth winding switched reluctance drives.....	5.12
Table 5.2	Comparison of experimental and simulated results for operation at 1225rpm.....	5.21
Table 5.3	Comparison of converter ratings for fully-pitched winding and single-tooth winding switched reluctance drives.....	5.22
Table 6.1	Comparison of torque performance for simulated unipolar and bipolar phase currents.....	6.9
Table 8.1	Comparison of phase currents for an average shaft torque of 30Nm.....	8.9
Table A.1	Physical parameters of the experimental fully-pitched winding, and single-tooth winding, switched reluctance machines.....	A.2
Table B.1	Name-plate rating of the DC load machine.....	B.2
Table C.1	Load circuit power ratings.....	C.2
Table D.1	Monitoring outputs available.....	D.4

Chapter 1

Introduction

The first recorded use of a switched reluctance machine was in 1838 when Davidson used it to drive a locomotive on the Glasgow to Edinburgh railway. In the 1920's the stepper motor was invented by Walker in Aberdeen, and had many features in common with modern stepping and reluctance motors (Miller [1.1]). At this point although the foundations were laid suitable power switching and control devices were not available.

Patents filed in the US by Bedford [1.2] in the early 1970's describe electronic commutation, power circuit configurations and synchronisation of the phase currents to rotor position, all fundamental features of contemporary switched reluctance drives. During the 1970's and 80's there was a great deal of effort put into developing switched reluctance and stepping motor drives, most notably at Leeds and Nottingham Universities. Messrs Acarney, Bausch, Blenkinsop, Bryne, Corda, Davis, Fulton, Harris, Hughes, Lawrenson, Ray, Rieke and Stephenson are most notable. Between them they have published many notable papers including Bausch *et al* [1.3], Lawrenson *et al* [1.4], Byrne *et al* [1.5] and Hughes *et al* [1.6], which have laid the basis of modern switched reluctance and stepping drives.

Switched Reluctance Drives Ltd., now part of Emerson Electric, has played an important part in the commercial development of switched reluctance drives and their market acceptance in Europe. Following the rapid pace of development in the 1970's and early to mid 80's, work has concentrated on refining the basic principles. Research and development has concentrated on aspects such as improvements in efficiency, machine thermal design, the reduction of torque ripple and acoustic noise, elimination of sensors, and the simplification of power and control circuits.

After many years with no fundamental change in the winding configuration of switched reluctance machines, a new winding configuration was published by Mecrow [1.7]. The fully-pitched winding configuration as it was named results in a machine with closely coupled phases, in contrast to the de-coupled phases of a conventional switched reluctance machine. Covered by Patent number 9 126 206.3 (December 1991) the new winding configuration offers increased torque per unit copper loss, when compared to a conventional switched reluctance machine. At the start of this Ph.D., project work

completed by Mecrow included the experimental measurement of flux-linkage and torque characteristics on a prototype fully-pitched winding switched reluctance machine, and the outline of three potential excitation patterns. It is at this point that the author started development work on fully-pitched winding switched reluctance drives.

1.1 Objectives

At the start of this project the objectives were few in number, and broad in coverage. As often happens in research, as the project progressed the objectives grew in number and became more focused. Stated below are the original broad aims:-

- Gain a deep understanding into the operating principles of fully-pitched winding switched reluctance drives.
- Design and build an electronic drive suitable for controlling the experimental fully-pitched winding switched reluctance machine.
- Implement various machine excitation schemes and gain an understanding of their operation.
- Evaluate and compare the different modes of operation.
- Improve the performance of switched reluctance drives by working on commonly perceived problems such as torque ripple and acoustic noise.

As will ^{be} shown in the following chapters of this thesis, all of the objectives have been met.

1.2 The Contribution to Knowledge

The following points outline what are believed to be new areas of work, previously unpublished by other authors. Most, but not all relate to the operation of fully-pitched winding switched reluctance drives.

- Real time operation of a fully-pitched winding switched reluctance machine.
- Unipolar operation of a fully-pitched winding switched reluctance machine.
- Bipolar operation of a fully-pitched winding switched reluctance machine.
- Flux-linkage control of a switched reluctance machine.
- Experimental comparison of published methods for the constant torque operation of switched reluctance machines.
- Constant torque operation using 'flux ramps'.

- Use of a Genetic Algorithm for the optimisation of switched reluctance machine control parameters.
- A comparison of the acoustic noise produced by various excitation schemes.

1.3 Overview of the Thesis

This thesis has two main sections: the Chapters and the Appendices. Chapters describe the theory, simulation and experimental results obtained for various modes of switched reluctance drive operation. In general, as the chapters progress, the operating modes become more complex and advanced in nature. The appendices detail aspects such as machine characteristics, circuit configurations, and the drive hardware and software. Although not fundamental to the understanding of the chapters, the appendices may provide answers to questions such as; "How was that done?", and "Why is that?" which may occur to the reader whilst reading the chapters.

This chapter concludes with a look at the state of the art in switched reluctance drives, before **Chapter 2** introduces the concept of fully-pitched winding switched reluctance machines. Comparisons are made between fully-pitched winding switched reluctance machines and single-tooth winding switched reluctance machines, and a set of transformations are presented which greatly simplify the understanding of fully-pitched winding switched reluctance machine operation.

The physical implementation of the experimental drive is covered in **Chapter 3**, followed by a review of current control methods, and a description of the method used for the experimental drive in **Chapter 4**.

The scene having been set in previous chapters, the simplest mode of operation for a fully-pitched winding switched reluctance machine - unipolar operation, is described in **Chapter 5**. This is developed further in **Chapter 6** which covers bipolar operation. Bipolar operation gives a much wider scope in the choice of current waveform and circuit configuration with several possibilities explored.

More advanced control methods are introduced in **Chapter 7** which describes the flux-linkage control of a switched reluctance machine. An in depth exploration of constant torque operation can be found in **Chapter 8**. Published methods are reviewed and some evaluated experimentally. Building on the work of Chapter 7, a new method of constant torque operation is introduced which is based on flux-linkage control. A Genetic Algorithm is used to optimise parameters for constant torque operation, and a speed controller implemented which uses the optimised flux-linkage waveforms.

The different modes of operation covered in previous chapters are compared in **Chapter 9** before drawing final conclusions in **Chapter 10**.

Appendices cover machine characteristics, the experimental rig, power and control circuits and software, and give an introduction into Genetic Algorithms.

1.4 Electrical Machine Drives

Electrical machine drives play a vital role in the world today, with applications in almost every aspect of life. Domestic and automotive applications provide a vast market for low power applications, with manufacturing and process industries providing a lower volume higher value market for medium and high power drives. Although a considerable simplification, the market for electrical drives can be split into three areas: single speed drives, variable speed drives and servo drives; each with successively higher levels of performance and complexity.

1.4.1 Single Speed Drives

Single speed drives represent by far the largest use of electrical machines with applications from fans and pumps to conveyer belts and shavers. The type of machine generally depends on the supply available; for battery and automotive applications brushed DC machines are used, single phase AC supplies use split phase or shaded pole induction motors, or series connected commutator motors (universal motors). Whilst for three phase AC supplies induction motors are used, except at very high powers where a synchronous machine may be employed.

Although having a fixed running speed, some applications require a soft start facility to reduce the mechanical or electrical stresses associated with start up. For AC motors this is usually achieved by phase control of the applied voltage.

1.4.2 Variable Speed Drives

Variable speed drives represent a large and varied market from washing machines and windscreen wipers, through conveyers and mine winders, to rolling mills and rail traction. This was once the preserve of DC machines with variable resistances or Ward-Leonard supplies, but the advent of electronics has opened up a wide range of machine types for variable speed operation. The market is dominated by brushed DC machines and AC induction motors, with synchronous machines (including permanent magnet) and switched reluctance machines accounting for most of the remainder. The European market for variable speed drives was estimated at \$1550 million in 1992 by Frost and

Sullivan (Scharf [1.8]), with DC drives accounting for approximately 40% of the market (Drives and Controls magazine, July/August 1995).

In contrast to single speed drives, which usually have no or very little electronics associated with them, modern variable speed drives contain a large amount of electronics often integrated into a small package. The rapid pace of electronics development over the last twenty years has had a major impact on variable speed drives, most notably in the widespread adoption of AC induction motors in preference to brushed DC machines. Power semiconductors capable of switching efficiently at high frequency, and microprocessors capable of executing many million instructions per second, have enabled complex control schemes to be implemented and given AC induction motor drives a level of dynamic performance previously only available from DC drives.

Great interest is currently being shown in variable speed drives for applications that were previously fixed speed because they can offer substantial energy savings in some cases. Fans and pumps run at a fixed speed with the flow rate altered by mechanical throttling can give quite dramatic energy savings by employing a variable speed drive. Energy savings of 42% and 64%, with a pay back period of 18 months have been reported (Drives and Controls magazine, April 1995).

Other fixed speed applications such as conveyers which require soft starting are increasingly being replaced with variable speed drives. This allows soft starting at full torque, and the possibility of variable running speed which gives flexibility.

1.4.3 Servo Drives

The title 'Servo Drive' is usually applied to a high performance variable speed drive which often has the ability to control position as well as speed. Servo drives offer the highest performance of common drive systems, with a typical torque bandwidth of 1000Hz (Drury [1.9]). Three types of machine are commonly employed: brushed DC, and permanent magnet synchronous machines with either trapezoidal (brushless DC), or sinusoidal (brushless AC), voltage characteristic. When using a brushed DC machine a high frequency chopper is required rather than a line commutated converter to achieve 'servo' performance. Typical applications include robotics, machine tools and high resolution positioning systems. It is estimated by Frost and Sullivan that the world wide market for servo drives will be worth \$1813 million in 1999 (Drives and Controls magazine, June 1995).

1.5 The state of the art in Switched Reluctance Drives

Before looking at current research topics and drive technology, the basic parameters that define a switched reluctance machine will be investigated, in relation to their selection for particular applications.

1.5.1 Fundamental Machine Parameters

A switched reluctance machine is defined fundamentally by the number of electrical phases and number of stator and rotor teeth which it has. The selection of these parameters is defined by several factors, the most important of which is the application.

The Number of Stator and Rotor Teeth

For a switched reluctance machine with m phases, in order to minimise the switching frequency the number of stator teeth n_s , and rotor teeth n_r , should be chosen so that:-

$$LCM(n_s, n_r) = m \cdot n_r \quad (1.1)$$

Where $LCM()$ denotes the lowest common multiple. The 'step' or 'stroke' angle between successive aligned positions is given by:-

$$\theta_{step} = \frac{2\pi}{m \cdot n_r} \quad (1.2)$$

And each phase is excited with an electrical frequency given by:-

$$f_{phase} = \frac{n_r \cdot \omega}{2\pi} \quad (1.3)$$

As the rotor moves through one step or stroke the mechanical energy delivered to the load is equal to the change in co-energy, $\Delta W'$ (Joules). Average running torque (Nm) is then given by:-

$$T = \frac{m \cdot n_r}{2\pi} \cdot \Delta W' \quad (1.4)$$

From these equations it can be seen that increasing the tooth number increases the average torque and excitation frequency. Thus in theory torque can be maximised by using a machine with a large number of stator and rotor teeth. In practice two main effects limit the tooth number: firstly, having a large number of teeth leads to a complex structure with a low ratio of aligned to unaligned inductance thus reducing the co-energy $\Delta W'$, and secondly, as the electrical frequency increases less time is available to

excite each phase which increases the converter VA requirement. The higher electrical frequency also increases eddy current losses.

A general requirement for industrial variable speed drive applications is the ability to self start in either direction which requires three or more phases. For the speed range 0-10,000rpm common tooth numbers are 6:4, 8:6 and 12:8, where 12:8 is effectively a four magnetic pole version of 6:4. A four phase 8:6 machine has the advantage of significantly reduced torque ripple in comparison to the three phase 6:4 and 12:8 arrangements.

Applications covering 10,000rpm and above favour designs with a lower electrical frequency in order to minimise converter rating and iron losses. A three phase 6:4 configuration is a popular choice as it has an electrical frequency half that of 8:6 and 12:8 arrangements.

For low speed applications such as direct drive robotics, high torque and a small step angle are desirable. This can be achieved by increasing the tooth number. For a given number of phases the tooth number can be increased by having a four or eight pole magnetic design with short flux paths. An alternative method is to retain a two pole magnetic design and have multiple teeth per stator pole, as developed by Finch *et al* [1.10]. Although having a large number of teeth e.g. 24:32 (Miller [1.1]) gives a large number of strokes per revolution, 96 in this case, the fine tooth structure leads to a low inductance ratio reducing the torque capability.

Although the stator and rotor iron of a switched reluctance machine are very simple in basic shape and construction, the small air gap required to maximise inductance ratio and hence output is often cited as a manufacturing cost penalty. In a comparison made by Miller *et al* [1.11] between various types of machine with the same air gap, switched reluctance machines were shown to have a torque density considerably higher than induction machines. This then contradicts the assertion that switched reluctance machines have a manufacturing cost penalty due to small air gap tolerances.

The number of phases

The selection of phase number is determined to a great extent by cost and performance. In very cost sensitive applications where performance in terms of torque ripple and starting ability are not particularly important, single and two phase drives are most appropriate. Single phase drives cannot self start from all rotor positions and can only produce positive torque for at most half a cycle, thus torque ripple is high and the load needs to have sufficient inertia to coast round to the next cycle. One method developed

to ensure starting is to use a 'parking magnet' (Horst [1.12]) which holds the rotor in the torque-producing region at standstill. Two single phase machine designs are investigated by Chan [1.13] for use in a ceiling fan. One has a pressed steel stator and rotor and radial flux, and the other laminated teeth and axial as well as radial flux paths. He also describes two power circuits, one using a phase controlled Triac directly from the mains. Although very simple and low cost the drives meet the performance level required. However no starting method is covered.

Although a completely symmetrical two phase machine has no self starting ability from the aligned position, various designs exist which can self start in one direction only. These include the 'stepped air gap' and 'snail cam' motors (see Miller [1.1]) both of which have tapered air-gaps, and a design by Byrne [1.14] which uses controlled saturation. The large winding slots and unaligned air gap enable machines to be designed with low copper losses and high inductance ratio, thus maximising the torque per ampere and efficiency. With their simplicity and low parts count, two phase switched reluctance drives can be expected to become more common in applications requiring uni-directional operation and low starting torque, such as fans and central heating pumps.

The ability to start from any position and in either direction requires at least three phases, thus this is the minimum number used for industrial variable speed drives and applications requiring a high starting torque, such as compressors. Four phase machines have a greater overlap between phases and substantially reduced inherent torque ripple compared to three phase machines, as shown in Figure 1.1. Another advantage of four phase machines is that there are always two phases in their positive torque producing region, and by exciting both phases the effective torque per unit copper loss is increased.

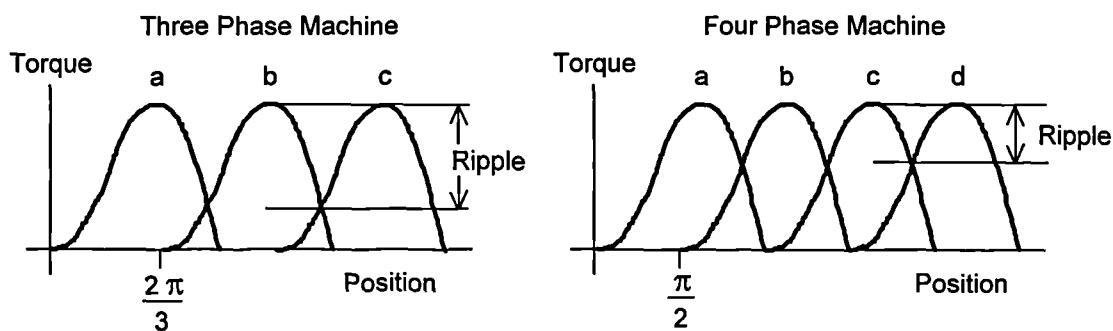


Figure 1.1 Typical torque characteristics of three and four phase machines.
with a single phase excited.

Torque ripple inherent in a machine can be further reduced by having more than four phases. The increased phase overlap also means that there are at least two, and at times three phases, which can all be excited to produce positive torque. Although having more

phases can increase torque density and reduce torque ripple, there are disadvantages to having a high phase number. The number of power devices and electrical connections to the machine both increase in line with phase number, thus increasing the cost and complexity of the drive. Having a high number of phases, and hence stator teeth, leads to a finer tooth structure and reduced inductance ratio which reduces the torque density.

Machines with an odd number of phases can be wound such that unipolar excitation of phases results in short flux paths as described by Michaelides *et al* [1.15]. In this configuration two adjacent phases are excited at a time, such that the preferred flux path is through two adjacent pairs of teeth and a short section of back iron, as shown in Figure 1.2. This configuration has several advantages. Copper losses are lower for a given MMF as two phases are contributing at any one time. Each phase is excited for a period corresponding to a rotation of two steps and thus the electrical commutation frequency is lower. As the flux path length is shorter the air-gap flux will be greater for a given MMF, and the volume of back iron excited at any one time is reduced, thus the iron losses will be lower. In tests on a seven phase 14:12 machine Michaelides *et al* were able to get double the torque at 500rpm for the same copper losses by configuring the windings for short flux paths.

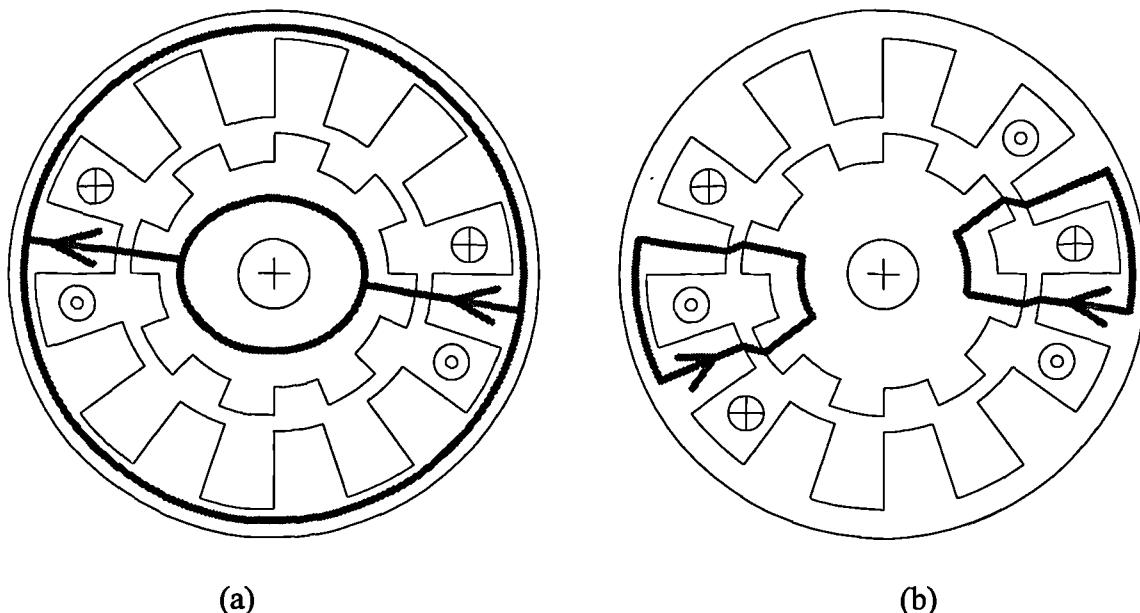


Figure 1.2 Cross section of a five phase 10:8 switched reluctance machine showing flux paths. (a) - Conventional excitation, (b) - Excitation for short flux paths.

1.5.2 Applications

Industrial variable speed drives represents a large, high value market for switched reluctance drives. However they face stiff competition from established technologies

such as DC and induction motor drives. Switched reluctance drives offer several advantages for industrial use including low maintenance requirements, low rotor losses, high torque density, four quadrant operation and a wide speed range. There are however some features which make switched reluctance drives less attractive: the need for a rotor position sensor, the large number of electrical connections to the machine, and the inability to interchange drives and machines. High torque ripple and acoustic noise are also two other commonly quoted disadvantages. Most industrial switched reluctance drives with a base speed of 1500-3000rpm use three or four phase machines, with 6:4, 8:6, or 12:8 tooth configurations. These phase and tooth numbers represent a good engineering trade off between several parameters including the number of power devices and electrical connections, inherent torque ripple, inductance ratio and electrical frequency.

Domestic appliances represent a very large potential market for low power relatively low performance switched reluctance drives. At present most domestic appliances are powered by series wound DC motors, split phase induction motors, and shaded pole induction motors, characterised by low torque density and low efficiency. A switched reluctance motor (designed by Switched Reluctance Drives Ltd.) for a food processor achieved a torque density three times that of the series DC motor it replaced (Blake *et al* [1.16]). Other potential domestic applications for switched reluctance drives include vacuum cleaners, washing machines, and power tools. Although the motor itself can be much smaller, and cheaper to produce in volume, this must be offset against the requirement for power and control circuitry. The number of power switches required can be minimised by using a single or two phase motor design, thus keeping the power circuit to a minimum. Highly integrated ASIC control circuits and integrated power modules enable the converter to be implemented in a very compact and low cost manner. The main hurdle remaining in miniaturisation is the dc-link capacitor. With improved manufacturing technology this will continue to shrink, and optimisation of drive design could further reduce the value of capacitor required. With ever higher levels of integration and performance, the size and cost of drive electronics can be expected to fall, making switched reluctance drives the higher performance, lower cost, preferred option.

Legislation on vehicle emissions is causing manufacturers to look at electric vehicles as a means of meeting ever tightening regulations. In addition to the main drive, electrical machines are being evaluated for use as main and auxiliary generators in hybrid vehicles. Although many early electric vehicles used DC motors, several motor types are now under evaluation in the interests of improved reliability, higher power density and reduced cost. Four types of machine: DC, induction, switched reluctance and

permanent magnet are considered by West [1.17] and Chang [1.18] for the main propulsion drive. In this application switched reluctance drives have the advantage of high torque density, high efficiency, simple construction and low maintenance. However the effect of torque ripple on drive train noise is a major concern, so too is the lack of high volume manufacturing experience, and EMC concerns due to the rapidly pulsating phase currents. Despite these problems the Post Office is planning a fleet of 50 vans powered by 30kW switched reluctance drives (West [1.17]). At present induction motor drives are favoured due to their low cost, low torque ripple and noise, and absence of rotor position sensor. The falling costs of permanent magnets are expected to make brushless DC drives the favoured choice in 10 years' time as they offer the highest torque density and efficiency of the four motor types.

Switched reluctance drives are currently being evaluated for aerospace applications including use as a fuel pump drive (Radun [1.19]), and as a gas turbine starter generator (Ferreira *et al* [1.20]). In both applications the maximum operating speed is very high: 25,000rpm for the fuel pump, and 52,000rpm for the generator. Switched reluctance machines have been chosen for these applications because the simple rotor construction enables high speed operation, there is no generation into shorted windings, high torque density is possible, and problems of demagnetisation at the high operating temperature are avoided. As the rotor speed is high a three phase 6:4 design is used in each application to constrain the electrical frequency required. Although switched reluctance machines have no rotor copper loss, the eddy current loss due to the pulsating nature of the magnetic flux can be significant in high speed applications. Reported rotor iron losses for the gas turbine starter generator are 1.03kW at 29,684rpm and a shaft power of 20.4kW.

Although the majority of work being undertaken on switched reluctance drives is for rotary applications work is also being done on linear applications, for example a linear force actuator developed by Corda *et al* [1.21].

1.5.3 Machine Winding Configurations

Doubly salient switched reluctance machines have been developed on the basis of having the phase windings short pitched around a single stator tooth. In this configuration mutual coupling between windings is minimised and torque production is due to the change in self inductance with rotor position.

Still based on the concept of short pitched de-coupled windings, the configuration proposed by Michaelides *et al* [1.15] enables machines with an odd number of phases to operate with short flux paths. Flux now circulates around sets of teeth rather than around

the whole machine as shown in Figure 1.2. The reduction in magnetic path length, and hence reluctance, gives a higher air-gap flux for a given MMF and reduces core losses.

A major departure from the conventional winding configuration came with the publication by Mecrow [1.7] of the fully-pitched winding arrangement. For a three phase machine, phase self inductances are almost independent of rotor position, and torque is derived from changes in mutual inductance with position. This winding configuration is described more fully in Chapter 2, and operation of fully-pitched winding switched reluctance machines is the subject of much of this thesis.

A hybrid winding arrangement is described by Liang *et al* [1.22], in which a switched reluctance machine with conventional short-pitch windings is fitted with an additional commutating winding. The commutating winding is mutually coupled to each of the main phase windings and is used to store the magnetic field energy during commutation rather than returning it to the supply. A further development of this is to replace one of the short-pitch windings with two fully-pitched windings. The resulting 'three phase' machine has two short-pitch windings and two fully-pitched windings (Li *et al* [1.23]). This winding configuration overcomes the disadvantage of having to find space for the commutating winding. In the new hybrid machine the two short-pitch phases are excited in turn, followed by the two fully-pitched windings together. One of the fully-pitched windings is also used as a commutating winding when current transfers between short-pitch phases. Torque is produced from the change in self inductance for two thirds of a cycle and from changes in mutual inductance during the remaining third of a cycle. For both these winding configurations, better high speed performance is demonstrated in comparison to a conventional switched reluctance drive, due to the reduction in energy circulating to the supply during commutation.

1.5.4 Power Converters

Modern power devices have revolutionised the implementation of switched reluctance drive power converters, and the search for simple low cost converters has lead to the development of several novel circuit topologies.

Power Switching Devices

Considerable advances in power switching devices have been made over the last twenty years with new device types being the standard for modern designs. In the past bipolar transistors covered low and medium power applications (up to several tens of kW) with force commutated thyristors covering high power applications. In both cases the circuits were complex and bulky due to the high current base drives required for large bipolar

transistors, and commutating circuits in the case of thyristors. Modern power devices including MOSFETs (Metal Oxide Insulated Field Effect Transistors), IGBTs (Insulated Gate Bipolar Transistors) and GTOs (Gate Turns Off thyristors), are characterised by relative ease of control and rapid switching.

Power MOSFETs offer the highest switching speeds and largest SOA (Safe Operating Area) and require very little gate power. However the conduction losses are quite high, especially for devices rated above 500V. The low current density (once again especially for devices rated at over 500V) means that a large silicon area is required which increases the cost per ampere. Because of the high conduction losses and cost per ampere MOSFETs are generally only used in very high performance low power drives where the high switching speed is required and the losses can be tolerated. An exception to this is low voltage applications (less than approximately 60V to 100V) where the very low 'on' resistance possible gives lower conduction losses than with other devices.

Probably the most important device in the drives market today is the IGBT; with ratings of 500V to 1700V, and 10A to 1000A, devices are available to cover a wide range of applications. The IGBT combines several features of bipolar transistors and MOSFETs to give a low loss easy to drive switch. Conduction losses are similar to a bipolar transistor and can be traded off for switching speed. Although the switching speed is faster than a bipolar transistor, total switching losses are quite high due to a long current 'tail' at switch off, and it is this that limits practical switching frequencies to about 20kHz at low power (<10kVA), and 2kHz at high power (>250kVA). Current density is even higher than for bipolar transistors which means that a smaller silicon area is required for a given rating. This has two main benefits: lower device cost, and lower gate charge compared to a similarly rated MOSFET.

For applications with a dc-link voltage above approximately 1500V, suitable IGBT devices are not available and GTOs are the preferred choice. These have a very low conduction loss and high current density. Before the development of high current IGBTs, GTOs covered a wide power range, even down to 22kW (the Tasc 180 switched reluctance drive described by Ray *et al* [1.24] was rated at 22kW and used a GTO power circuit). Now however GTOs are reserved for the highest power applications, with single devices rated at up to 4.5kV, 3kA. Although little gate current is required to turn on and keep on a GTO, the turn off gain is in the region of 5 which for an anode current of 2.5kA would require a gate current of 500A, although only for 10 μ s or so. As well as requiring a high current capability, GTO gate drive circuits are considerably more complex than those required for an IGBT, especially if operation over a wide

temperature range is required. Switching losses are also much higher than for IGBTs, with low inductance snubber circuits being required.

With the promise of conduction losses comparable to a thyristor, low switching losses, and a charge controlled gate, the MCT (MOS Controlled Thyristor) promises to be the nearest thing to an ideal switch yet developed. As yet though they have not gained market acceptance.

Circuit Configurations

For very low cost applications it is desirable to minimise the number of switching devices. The circuits in Figure 1.3 use a single switch per phase and dissipate stored magnetic energy. This limits the maximum drive power to less than 100W for practical purposes. The zener diode dump circuit has the advantage of a constant phase voltage during switch off which decreases the current fall time. It also defines the switch peak off state voltage.

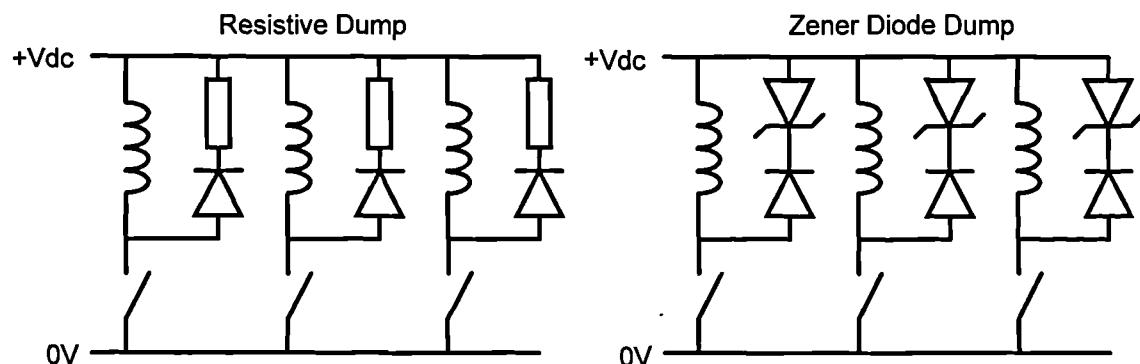


Figure 1.3 Converter circuits with dissipative energy dumps.

Figure 1.4 shows three standard power circuit configurations used for switched reluctance drives over the whole power range. In all of the circuits the energy stored in the phase windings at the end of the conduction interval is returned to the supply. These circuits are also suitable for use with switched reluctance generators. Each circuit configuration has its advantages and disadvantages.

The bifilar winding converter has the advantage of having a single ground referenced power switch and a single diode per phase. There are however several disadvantages. As the machine requires two sets of windings per phase there is less winding area available for each winding, which reduces the machine output for a given size. There are also three or four electrical connections required per phase in comparison to two per phase for other converters. For the converter, a switch voltage rating at least twice the dc-link voltage is required, which can limit the choice of devices for applications with a three

phase 440V or 660V supply. Imperfect coupling between windings means that a snubber is required to dissipate the energy stored in the stray inductance, which reduces efficiency and makes this circuit unsuitable for high power drives.

The split supply circuit, as with the bifilar circuit, has a single switch and diode per phase. As each phase is clamped to the supply the device voltage ratings are determined by the dc-link voltage. This configuration has the potential for the lowest number of electrical connections required, by having a single common connection to give $(n+1)$ total connections. The main disadvantage of this configuration is that there must be an even number of phases, the control circuit must balance the phase currents to keep the centre point voltage at the correct level, and DC operation of a single phase (as required for holding against an applied torque) is not possible.

Probably the most common power circuit, the asymmetric half-bridge, can be used with any number of phases and has no special control limitations or requirements. The main disadvantage of this circuit is that two switches and two diodes per phase are required. There are also $2n$ electrical connections required to the machine.

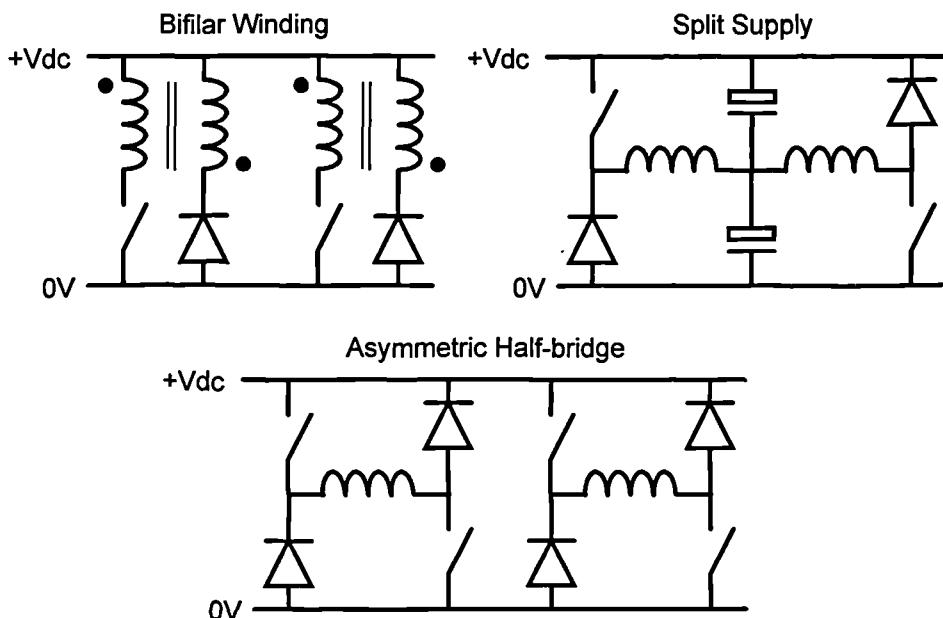


Figure 1.4 Standard two-quadrant converter circuits.

In an effort to reduce device count and therefore cost, much research effort has been put into developing novel power converter configurations. In Europe Pollock and Williams have done considerable work in this area, especially for machines with five or more phases. Three novel power circuit configurations are shown in Figure 1.5 and are discussed here.

The 'C-Dump' circuit has $(n+1)$ switches with the '+1' being part of an auxiliary converter. The circuit is very similar to the simple circuits shown in Figure 1.3 but rather than dissipating the stored energy it is dumped into a capacitor (C_d). An auxiliary Buck converter formed by S_d , D_d and L_d , then returns the energy to the supply from the dump capacitor. The dump capacitor is usually maintained at about twice the dc-link voltage so that the switch off phase voltage is approximately equal to the dc-link voltage. For motoring only applications the dump energy is less than the output power and so the dump circuit switch has a lower power rating than the phase switches. The main disadvantages of this circuit are the need for an auxiliary controller, and that failure of the dump circuit could lead to device over-voltage.

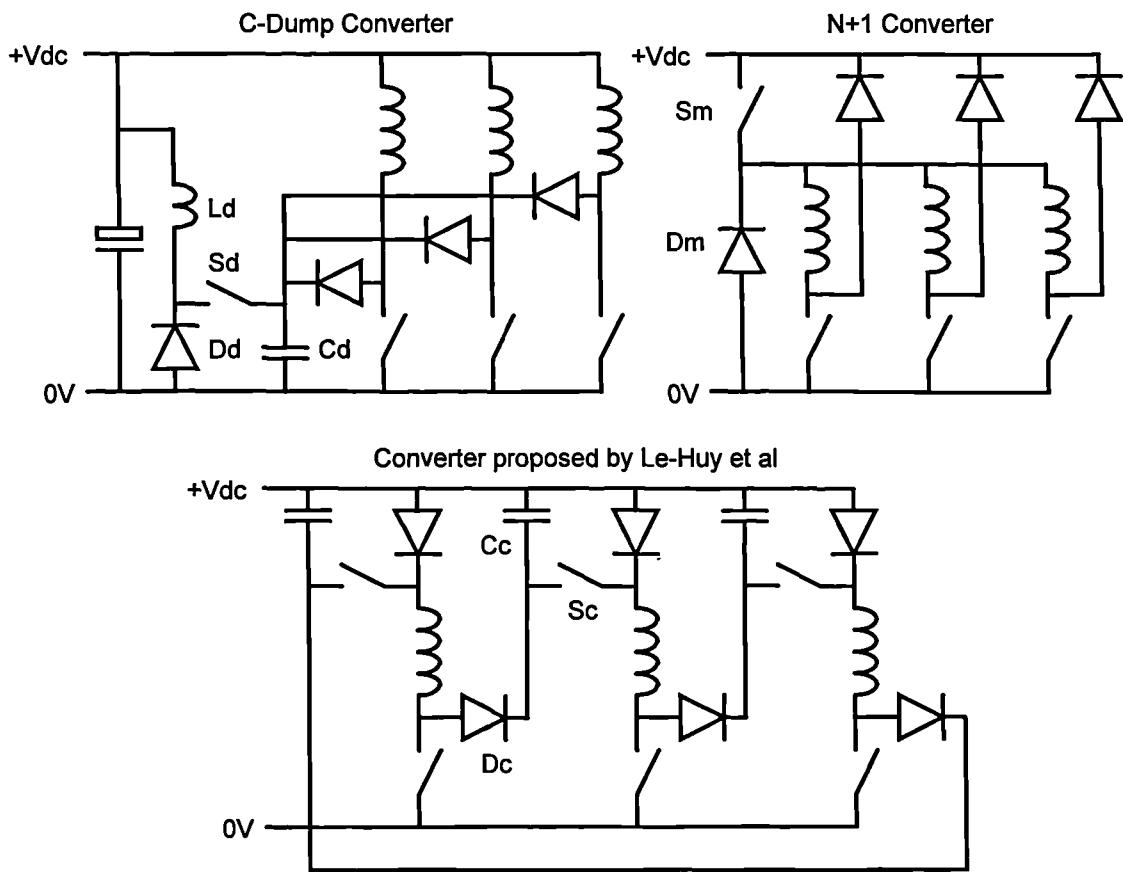


Figure 1.5 Novel converter circuits.

The $(n+1)$ converter circuit shown is one of several configurations described by Pollock *et al* [1.25] and Krishnan *et al* [1.26]. This circuit is based on the asymmetric half-bridge but with one high-side switch shared by two or more phases. The main switch (S_m) is used to control the phase current by PWM or some other modulation scheme, whilst the phase switches select the active phase. This circuit has the advantage that device voltages are clamped to the dc-link, and the number of electrical connections to the machine is only $(n+1)$. The main disadvantage of this circuit is that the phase

voltages are not fully independent, and if more than one phase is enabled at a time the individual phase currents cannot be controlled. This means that the power circuit is best suited to applications and machines which do not require overlapping current pulses: but with careful selection of converter topology and control strategy this is not a major limitation. Pollock also describes a $2(n+1)$ converter which is essentially a two phase version of that shown in Figure 1.5 repeated as many times as are required for the number of phases, as well as other novel converters for 6, 7 and 9 phases.

Although still having $2n$ power switches the converter proposed by Le-Huy *et al* [1.27] only requires n fully controllable switches. It is intended for low cost applications using thyristors for the commutation switches (S_c). In a similar manner to the 'C-Dump' circuit the stored energy is returned to a capacitor (C_c). However rather than returning the energy to the supply, it is delivered into the next phase. At the end of a conduction period the main switch is turned off, stored energy then flows from the phase winding into the commutation capacitor C_c . When the next phase is switched on both the main switch and the commutation switch (S_c) are switched on. Energy then flows from the capacitor (C_c) into the phase with the increased voltage (due to the voltage on C_c) speeding up turn on of the phase. When the capacitor is discharged current flows from the supply through the diode to the phase. This circuit is only suitable for motoring in one direction, as energy cannot return to the supply and the phase sequence is fixed.

1.5.5 Control of Switched Reluctance Drives

In many respects the control of switched reluctance drives has changed little over the years apart from the method of implementation. New control methods are now receiving attention, spurred on by the availability of the processing power required to implement them.

Control Algorithms

Basic control of switched reluctance drives is still based on two modes of operation: current chopping at low speed, and single pulse voltage control at high speed. This method is relatively simple to implement and requires little real time computation. The required switching times and current levels are stored in some form of memory and are generally derived from extensive testing of a prototype drive.

In an effort to improve the dynamic performance and torque ripple obtained with the basic angle and current control method, several modern control techniques have been applied to the switched reluctance machine. Taylor [1.28] uses a composite control structure with feedback linearisation to improve the performance of a position control

system. Feedback linearistaion gives a linear torque response and enables linear control principles to be applied. A non-linear adaptive model is used by Amor *et al* [1.29], to reduce torque ripple and improve position control. Parameters are adapted on line, with torque calculated from flux and current.

The use of Neural networks, Reay *et al* [1.30] and Fuzzy logic, Reay *et al* [1.31] have been investigated for torque ripple reduction. The Neural network and Fuzzy inference rules both give a non-linear current/position relationship, which is trained so as to minimise the torque ripple. A potential problem with this method is the large number of weights and rules required to give constant torque operation over a range of torque values and speeds.

Variable structure, or sliding mode control is used by Buja *et al* [1.32] to reduce torque ripple and improve speed control. Sliding mode control is based on rapidly switching the control transfer function so as to keep the system state vector 'sliding' on a particular plane. The plane used by Buja *et al* is that defined by zero speed error and the maximum phase current. Simulated results presented for fixed control angles show a significant reduction in torque ripple.

A significant advance in switched reluctance drive control would be the development of a scheme for calculating the instantaneous motor torque and optimum excitation waveforms to give constant torque, without the need for extensive motor parameterisation. This would enable drives and machines to be interchanged as with induction motor and DC drives.

Silicon Implementation

There are four basic technologies used to implement drive controllers: analogue, microcontroller/microprocessor, DSP (Digital Signal Processor) and ASIC (Application Specific Integrated Circuit).

Totally analogue circuits are limited to very low performance switched reluctance drives because it is difficult to program in the variation of phase advance and current level required for higher performance. However with the addition of some simple digital circuitry to give reference angles and currents, good drive performance can be obtained (Sugden *et al* [1.33]).

Single chip microcontrollers enable almost all control circuitry to be integrated into a single chip by combining processor, memory, analogue voltage measurement, timers and pulse generators. Popular controllers for motor drive applications include members of the 8051 family, the 8096 and 80196, 80166 and 68300 series. In some applications it

is cheaper to use a microcontroller than analogue circuitry due to the much reduced component count. Flexibility afforded by the use of a processor can also reduce the cost of implementation. Modern microcontrollers offer enough performance to implement conventional control schemes, and with built in communications ports are very attractive for some applications, most notably automotive applications.

With more complex and computationally intensive control schemes such as those for sensorless operation (Elmas *et al* [1.34]), control using neural nets or fuzzy logic (Reay *et al* [1.31]), and constant torque operation (Rochford *et al* [1.35]), insufficient processing power is available from a simple microcontroller or microprocessor. This has led to the increasing adoption of DSPs for implementation of control algorithms, especially in research work. There has been spectacular growth in the processing power available over the last ten or fifteen years and a corresponding drop in the price of that performance. Taking the Texas Instruments range of DSPs as an example, the first device, the TMS320C10 was introduced in 1982; a 16 bit fixed point device it could perform 5MIPS (Million Instructions Per Second). A more recent device the TMS320C40, a 32 bit floating point processor, can perform up to 275MFLOPS. Another recent device, the TMS320C32, offers 40MFLOPS performance for a volume price of \$10 a piece. With the decreasing price/performance ratio of DSPs they can be expected to be increasingly adopted for commercial drives.

By implementing only those functions actually required, ASICs can offer the lowest silicon cost of all technologies; however this only really applies to volume applications. ASIC technology has been widely used in designs by Switched Reluctance Drives Ltd (Sugden *et al* [1.33]) usually accompanied by a simple microcontroller for memory and house-keeping functions. Like microprocessors and DSPs, the price/performance ratio of ASICs continues to increase with ever larger devices available. Mixed signal ASICs such as the LMB1008 (Miller *et al* [1.36]), and the availability of microprocessors, such as the 8051 and Z80, as ASIC macro-cells, enable all control functions to be integrated into a single chip. ASIC technology is suitable for implementing all but the most numerically intensive control algorithms.

1.5.6 Improving Drive Performance and Overcoming Prejudices

Three often cited (and often repudiated) criticisms of switched reluctance drives are that they have high torque ripple, produce more acoustic noise than other drives, and in contrast to the simple and robust machine, need a precision rotor position sensor. Although progress has been made on these areas, switched reluctance drives have not acquired this reputation without there being some truth in the matter.

Reduction of Torque Ripple

The reduction of torque ripple has received considerable interest, especially when applied to three phase machines which have a higher inherent torque ripple than machines with four or more phases. Torque ripple reduction is one area where the controller by itself can make a big impact on performance. Several authors have published work on a variety of control methods, several of which are reviewed and investigated in this thesis. To avoid repetition work will not be repeated here, and the reader is directed to Chapter 8 section 8.2.

Reduction of Acoustic Noise

As with torque ripple reduction, noise reduction is a popular research area. The acoustic noise produced by an electric drive has two fundamental frequency components; that due to the electrical cycle frequency and proportional to speed, and that due to the frequency at which the applied voltage is modulated which may be a fixed PWM frequency.

Acoustic noise due to the voltage modulation frequency has received most attention in relation to induction motor drives because it is the predominant noise source (Wallace *et al* [1.37]). Wallace *et al* found that it was the continuous fixed tone that caused loss of hearing sensitivity. Garcia-Otero *et al* [1.38] propose using harmonic elimination to avoid exciting motor resonances at frequencies where the human ear is most sensitive. In low power drives the noise can be eliminated completely by using an ultra-sonic modulation frequency. With the development of soft-switching resonant converters (Kim *et al* [1.39]) and as power devices become faster it can be expected that it will be possible to use ultra-sonic modulation frequencies in higher power drives.

Cameron *et al* [1.40] conducted an extensive investigation into the origins of acoustic noise in switched reluctance machines. Their most important conclusions were that the major noise source was stator vibrations caused by ovalising forces, torque ripple was not a major noise source, nor was the voltage modulation frequency. To avoid excitation of motor resonances Cameron *et al* propose the use of control angle dither so as to spread the exciting frequency spectrum. An alternative method proposed by Wu *et al* [1.41] (see also Wu *et al* [1.42]) is based on the observation that stator vibrations are excited when a phase is de-energised. They propose having a two stage switch off process, phased so that the stator vibrations caused by each step cancel each other out. Experimental results show that this method substantially reduces the amplitude of stator vibrations.

Elimination of Rotor Position Sensors

The torque output of a switched reluctance drive is critically dependent on the accurate synchronisation of excitation to rotor position, and hence some means of determining rotor position is required. An optical encoder coupled to the machine shaft is the most commonly employed method. This may be as simple as a single punched disk and slotted opto-coupler, or as complex as a high resolution multi-track encoder. However despite the potential simplicity of a rotor position sensor there are many reasons why it is desirable to eliminate the position sensor including the following:-

- Reduction of system costs
- Increased system reliability
- Reduced machine volume
- Elimination of sensor wiring
- To enable operation in high temperature or dusty environments

There has been considerable research work done in this area in the last ten years and a variety of methods developed. Most of the methods described here and others have been reviewed by Ray *et al* [1.43].

Ehsani *et al* [1.44] use a frequency modulation scheme to detect the variation of inductance with rotor position. In this scheme an unexcited phase forms part of a low power oscillator circuit, with the frequency measured by a microprocessor or analogue comparator via a frequency to voltage converter. This method suffers from the problem of how to connect a phase to the low power oscillator, and mutual coupling effects from excited phases.

The effects of mutual coupling are used to advantage by Husain *et al* [1.45] who use the position dependence of mutual inductance to determine rotor position. The voltage coupled into an inactive phase from the PWM voltage applied to another phase is sampled, and from its amplitude the rotor position determined. The effects on mutual inductance of saturation are not taken into account, and there are problems in applying the method to three phase machines (the method was developed for a four phase machine) due to asymmetric leakage inductances and excitation. A further problem is that at high speed the phase voltages are no longer pulse width modulated thus removing the measuring signal.

Originally developed by Acarnley for stepper motors the method used by Panda *et al* [1.46] uses the variation in current rise and fall times to determine position. Effectively the method is based on the variation of phase inductance with position and this is

reflected in the current rise and fall times. Although the effects of average current level and speed on times are shown graphically, Panda does not make it clear how these effects would be accounted for in practice.

The methods used by Mvungi *et al* [1.47] and Ray *et al* [1.43] use machine flux-linkage characteristics to determine position from measurement of flux-linkage and current. Mvungi *et al* use a phase which is not being used for torque production and determine the flux-linkage required to reach a low current threshold by voltage integration. A single dimensional lookup table then gives the position. Compensation of mutual inductance effects is achieved by addition of a non-linear correction signal to the flux-linkage. Ray *et al* use similar methods, using an inactive phase at low speed, but applied to an active phase at high speed. In each case the flux-linkage and current at a particular estimated position are measured. These values are then compared with the flux-linkage characteristics for a known position, and from this a position error determined, which is then used to correct the estimated position.

Most of the foregoing methods use an inactive phase for injection of a diagnostic signal. This may not be possible at high speed if a voltage squarewave is applied to the phase by the power converter. The other common feature of these methods is that a single phase is used at any one time to obtain position information. By using a real time motor model to 'observe' position, information is gathered from all phases and under all conditions. State observers for switched reluctance machines have been proposed by Lumsdaine *et al* [1.48] and Elmas *et al* [1.34]. Measured voltages are applied to the system model in order to estimate the phase currents. These are then compared with the measured phase currents and the error used to correct the estimated position and speed. Both use a linear motor model, and a first order linear load model. Lumsdaine *et al* make a further simplification and assume infinite load inertia and hence constant speed operation in order to simplify the observer. This is justified by the larger bandwidth of the model compared to the system. Results presented by Elmas *et al* are for operation with low phase currents in order to avoid non-linear saturation effects although Lumsdaine *et al* state that results were satisfactory even though non-linear effects were not modelled.

Another observer method originally developed for brushless DC drives by Ertugral *et al* [1.49] can also be applied to switched reluctance drives. This method uses machine flux-linkage characteristics to estimate phase current from measured flux-linkage and estimated position. The estimated currents are then compared with the actual currents and the error used to correct the estimated position and the flux-linkage. As actual

machine flux-linkage characteristics are used, the method is equally applicable to brushless DC and switched reluctance machines.

Although several of the methods described are computationally intensive, especially those using observers, the falling cost of processing power will make sensorless operation more attractive with time.

1.5.7 Switched Reluctance Drives in the Market Place

In the industrial sector switched reluctance drives face very stiff competition from well established and proven technologies such as DC and induction motor drives. At present there are few manufacturers in Europe offering switched reluctance industrial variable speed drives. Drives available from Allenwest, Sicmemotori and British Jeffrey Diamond are all designs licensed from Switched Reluctance Drives Ltd.

The domestic market offers the largest potential sales volume for switched reluctance drives. At present electronic drives have made little impact and there is no established new technology. Switched reluctance drives are thus well positioned to take over from series wound ^{commutator} motors and single phase induction motors as the domestic work horse, using very simple and low cost one and two phase drives.

In the following paragraphs the main rivals to switched reluctance drives for variable speed drive applications are considered briefly, with reference to current research topics and their market position.

Brushed DC Drives

Brushed DC drives were once the dominant technology in industrial variable speed drives but the robustness, low maintenance, and improving performance of induction motor drives has seen the market share for DC drives drop dramatically. Industrial DC Drives using phase controlled rectifiers offer excellent performance and are still specified for applications such as paper making where smooth torque control is a priority. Predictive control maximises dynamic performance which is limited to some extent by the line frequency. Low power servo drives fed from high frequency choppers offer the highest levels of performance for position control systems

Induction Motor Drives

Induction motors fed from voltage source inverters are the de-facto standard for industrial variable speed drives. With simple open loop V/F control dynamic

performance is poor but with the development of field-oriented control by Leonhard [1.50], performance can equal or surpass that of a separately excited DC machine.

Vector controlled drives are now offered by all the major drives manufacturers although some only use space vector modulation rather than full blown field oriented control. In the search for still higher performance ABB has introduced a range of drives which combine the principles of Direct Torque Control with a real time motor model (ABB [1.51]) implemented in a high speed custom processor. Some manufacturers are also using Fuzzy logic to improve drive performance, Mitsubishi for example (Mitsubishi [1.52]) use it to minimise power consumption in its' A140 drives by reducing the steady state V/F ratio if possible.

Research work on induction motor drives continues to develop different methods of vector control, many with the aim of simplifying implementation, for example Casadei [1.53]. Parameter estimation is also receiving considerable attention, one method is that of an observer (Du *et al* [1.54]) which it is claimed gives better performance than a Kalman Filter implementation. Condition monitoring is of increasing interest to industry and methods to detect various faults are being developed. Burnett *et al* [1.55] use harmonic analysis to detect broken rotor bars, whilst Penman *et al* [1.56] use a neural network approach to detect broken bars, supply faults and mechanical wear.

Brushless DC and AC Drives

There are two basic varieties of permanent magnet synchronous machine drives, those with a trapezoidal flux characteristic: brushless DC, and those with a sinusoidal flux characteristic: brushless AC. Both types are used for servo type applications because they can offer a high torque density, high torque to inertia ratio, and the large effective air gap (due to the magnets) gives a low phase inductance allowing fast current slew rates.

Brushless DC drives are usually three phase, and supplied with bipolar currents lasting for 120° each half cycle. Brushless AC drives are fed with sinusoidal currents and give lower torque ripple than brushless DC drives. However new control methods (French *et al* [1.57]) allow smooth torque to be produced by brushless DC machines.

Brushless DC drives can also be used for low cost applications using a simple position commutated converter without current control. This gives terminal characteristics similar to a permanent magnet brushed DC motor.

As with switched reluctance drives there is considerable interest in the reduction of sensors. By using information on the present converter switching state and a real time

motor model individual phase currents can be estimated from a single dc-link current sensor as described by French *et al* [1.58]. Extending the motor model to include a flux-linkage based observer, rotor position can also be estimated (French [1.59]). This opens up the possibility of a drive having no position sensor and only one current sensor.

Brushless drives offer high torque density, and high efficiency due to the absence of rotor losses. However temperature and mechanical shock limitations preclude the use of permanent magnets in some environments. At present magnet materials are also expensive which is important for low cost applications, but with magnet prices expected to fall this factor will diminish in importance.

Synchronous Machine Drives

In this context the term synchronous machine drives is used to cover a variety of synchronous reluctance motors and permanent magnet reluctance motors. Synchronous machines have several attractive features; absence of rotor losses, robust rotor construction suitable for high speed operation, low torque ripple, vector control principles may be applied to give excellent dynamic performance and standard induction motor converter hardware can be used.

Many types of synchronous reluctance motor have been developed with axially laminated types receiving most attention due to the high d-q axis inductance ratio they offer. Hybrid reluctance and permanent magnet designs have also been investigated (Miller *et al* [1.11]). Comparisons of synchronous and induction motor variable speed drives shows that each have their strengths and weaknesses (Lessmeier *et al* [1.60]) with no clear winner.

1.6 Conclusions

Although switched reluctance machines have been around for some time it is only in the last twenty years that real progress has been made and high performance drives become available. Semiconductor technology has had a major role in this and further developments can be expected to make low power drives cheaper and more compact, and give high performance variable speed drives better torque and noise performance. Research work continues to strive for better performance with fewer devices and sensors, trends which will make switched reluctance drives more competitive in a very tough market place.

Chapter 2

Fully-Pitched Winding Switched Reluctance Machines

2.1 What is a Fully-pitched Winding Switched Reluctance Machine

Switched reluctance machines generally have a doubly-salient structure, with different numbers of teeth on the stator and rotor. The coils of each phase winding are usually short pitch wound around a single stator tooth, so that the phases have little mutual magnetic coupling. A fully-pitched winding switched reluctance machine retains the iron structure of a conventional switched reluctance machine, but the coils of each phase winding are fully-pitch wound around several stator teeth. This results in a set of overlapping windings, similar to those in an induction motor, which have a strong mutual magnetic coupling.

A fully-pitched winding switched reluctance machine may be considered to be the dual of a conventional single-tooth winding switched reluctance machine, because whilst the single-tooth winding machine derives most of its shaft torque from changes in self inductance with rotor position, a fully-pitched winding machine derives most of its torque from changes in mutual inductance with position.

A single-tooth winding switched reluctance machine produces positive torque from periods of rising self inductance, thus each phase can only make a positive torque contribution for at most half a cycle. In contrast, by appropriate selection of phase currents, a fully-pitched winding switched reluctance machine can produce positive torque from periods of rising *and* falling mutual inductance. It is thus possible to utilise all the phases, all of the time, to produce positive torque.

Each of these points, and others, are considered in more detail in the following sections.

2.1.1 The Mechanical and Winding Arrangement

The physical structure of a fully-pitched winding switched reluctance machine is very similar to that of a conventional single-tooth winding switched reluctance machine. If a

cross section were to be taken in the middle of the lamination stack, little, if any, difference would be evident between the two machines. However looking at the end-windings would reveal the fundamental difference between the two types of machine.

As with a conventional switched reluctance machine, a fully-pitched winding switched reluctance machine has a doubly-salient iron structure, a wound stator, and no rotor windings or bars. A conventional switched reluctance machine usually has the coils of each phase winding wound around a single stator tooth, the resulting short pitched coils have short end-windings that do not overlap each other. A fully-pitched winding switched reluctance machine has the coils of each phase winding wound around as many stator teeth as there are electrical phases, e.g. the windings for a three phase machine span three stator teeth. The resulting fully-pitched windings have overlapping end-windings around the circumference of the machine. As the end-windings must traverse the width of several stator teeth and slots, rather than a single tooth width, they are necessarily longer in a fully-pitched winding machine. This will be covered in more detail later in this chapter. These arrangements are best illustrated by example, in this case a three phase 6:4 machine.

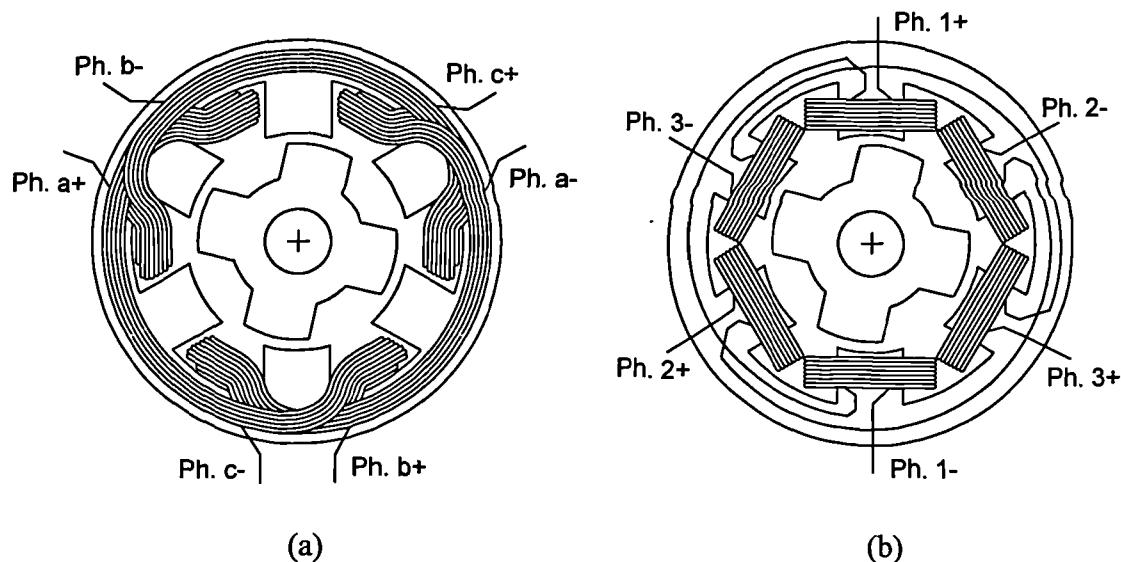


Figure 2.1 View of the end-windings in a fully-pitched winding switched reluctance machine (a), and single-tooth winding switched reluctance machine (b).

The general winding arrangement of fully-pitched winding, and single-tooth winding, switched reluctance machines is shown in Figure 2.1. Each of the three single-tooth windings, phases 1, 2, and 3, comprise of two series connected coils, each short pitched around a single stator tooth. There are a total of N_{123} turns per phase, half in each coil. The windings of a single phase occupy half the space in each of four slots. The three phase fully-pitched winding machine has phases a, b, and c, each of which comprises of

a single coil of N_{abc} turns, fully-pitched around three stator teeth. Figure 2.2 shows the two types of machine with the stator unrolled, this helps emphasise the coupled nature of the fully-pitched winding switched reluctance machine. Looking first at the single-tooth winding machine diagram (upper), it can be seen that magnetic flux passing through any of the stator teeth will only link a single coil. By contrast magnetic flux in any of the stator teeth of the fully-pitched winding machine (lower) will link all three phases.

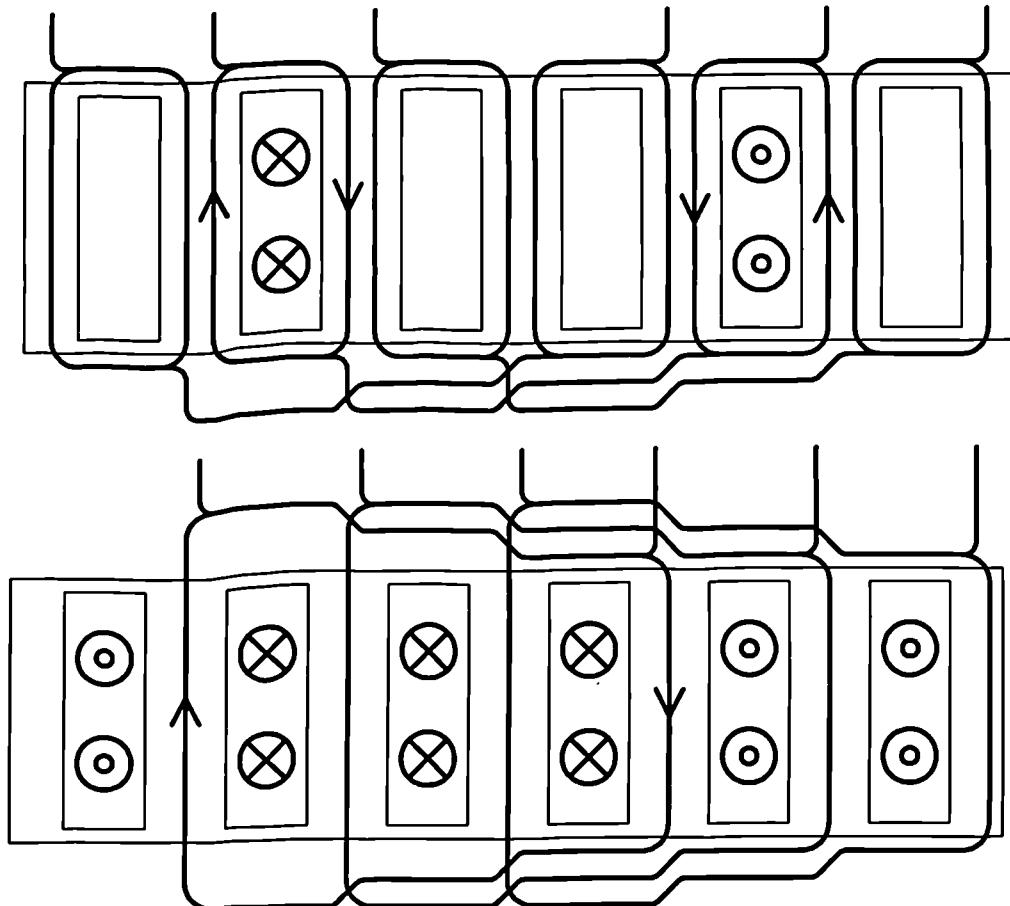


Figure 2.2 Unrolled machine showing winding arrangements and a single phase excited. Upper - single-tooth winding switched reluctance machine. Lower - fully-pitched winding switched reluctance machine.

2.1.2 Self and Mutual Inductances

All reluctance machines work on the principle of changing magnetic circuit reluctance with rotor position, in order to generate torque. This change of reluctance can be expressed as a change in the self and mutual inductances of a set of windings with rotor position. It is therefore useful to derive an approximation to the phase self and mutual inductances in a fully-pitched winding switched reluctance machine, as published by Mecrow [2.1]. The machine is considered to be ideal with a non-saturating iron circuit

of infinite permeability. Stator and rotor tooth widths are equal, as are the stator teeth and gaps. Furthermore, it is assumed that magnetic flux crosses the air gap only where stator and rotor teeth overlap, and there is no fringing of the field. Once again a three phase 6:4, fully-pitched winding machine will be considered. A cross section of the machine is shown in Figure 2.3, with tooth overlaps of x and y . The relationship between fully-pitched phases a and b is considered, other phase pairs follow a similar pattern.

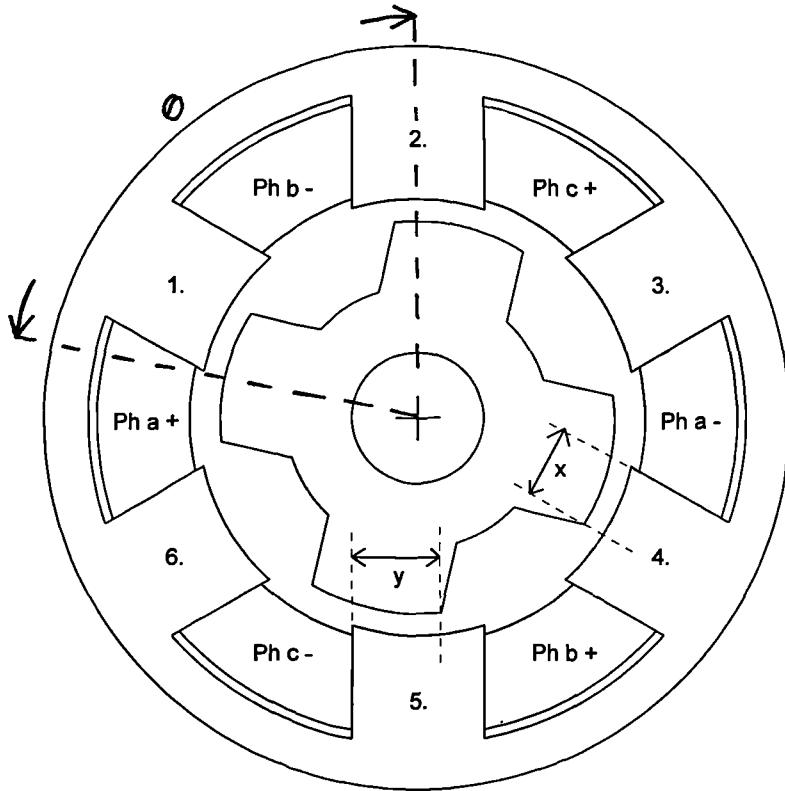


Figure 2.3 Machine cross section showing tooth overlap for inductance estimation.

The self inductance of phase a can be estimated to be a leakage component, L_{lk} , plus a component determined by the tooth overlap follows:-

$$L_{ph} = L_{lk} + \frac{\mu_0 N^2 \alpha \beta_L}{2G} \quad (2.1)$$

Where N is the number of turns per phase, α is the stack length, and G is the air gap length. From Figure 2.3 it can be seen that for flux linking the axis of phase 'a' the air gap area is proportional to the sum of the tooth overlaps, hence:-

$$\beta_L = x + y \quad (2.2)$$

As the rotor rotates the overlap of one pair of teeth increases, as the other pair decreases, with the sum remaining constant at one tooth width. The phase self inductance is

therefore independent of rotor position in this idealised case, and is equivalent to the fully aligned inductance in a conventional single-tooth winding machine.

The mutual inductance between two phases is dependant on the proportion of flux that links both phase windings. If we consider phase 'a' carrying positive current, such that flux circulates around the machine, through stator teeth 1, 2, and 3, through the rotor, and back around through stator teeth 4, 5, and 6. It can be seen from Figure 2.3 that flux passing through stator tooth 4 links phase 'b' in the same sense as phase 'a' whilst flux in stator teeth 5 or 6 links phase 'b' in the opposite sense to phase 'a'. The mutual inductance between phase 'a' and 'b' can thus be approximated by:-

$$M_{ab} = \frac{\mu_0 N^2 \alpha \beta_M}{2G} \quad (2.3)$$

where

$$\beta_M = x - y \quad (2.4)$$

As the rotor rotates the mutual inductance changes in sign and magnitude as the dimensions change from $x = 0$, $y = \text{one tooth width}$, to $x = \text{one tooth width}$, $y = 0$. The mutual inductance between other pairs of phases varies in a similar manner to 'a' and 'b' but are displaced from by a rotation of one stator tooth pitch. Based on equation 2.1 and equation 2.3 the machine self and mutual inductances can be plotted with respect to rotor position, as shown in Figure 2.4.

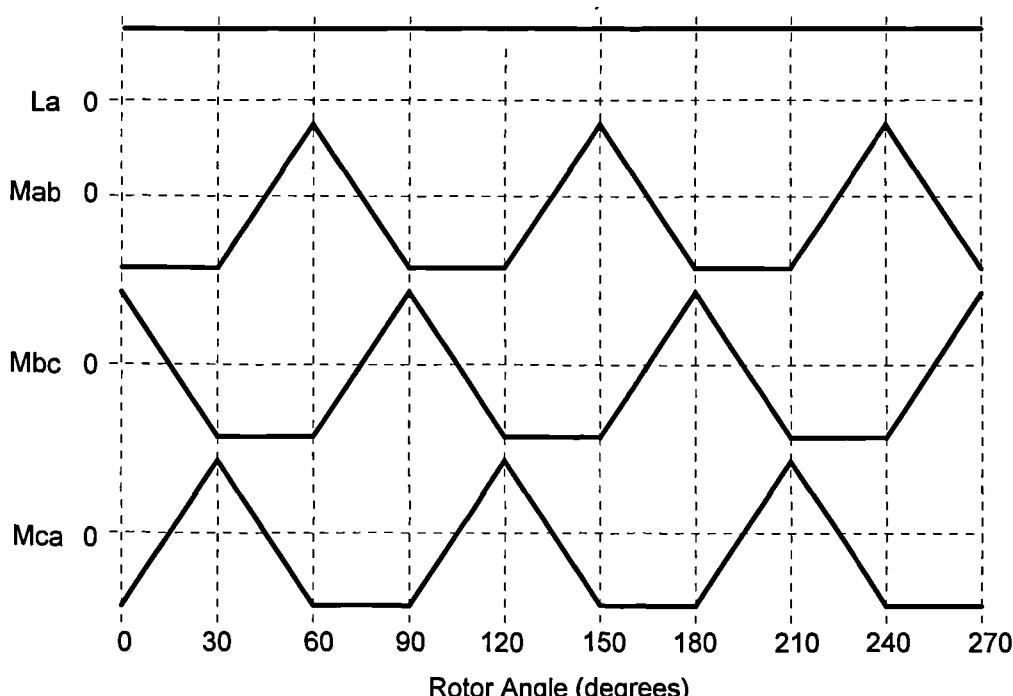


Figure 2.4 Idealised self and mutual inductances for a three phase 6:4, fully-pitched winding machine.

2.1.3 The Method of Torque Production

The fundamental operating principle of a fully-pitched winding switched reluctance machine is the same as that of a conventional single-tooth winding switched reluctance machine: Current is passed through windings placed in the stator slots to produce a driving MMF, the resulting magnetic flux circulates around the machine linking the windings and thus storing energy in the magnetic field. The stator and rotor are shaped such that as the rotor rotates, the reluctance of the magnetic circuit varies. This change in reluctance alters the energy stored in the magnetic field for a given phase current, and it is this change in energy that gives rise to shaft torque. For a reluctance machine the generalised equation for the shaft torque produced by each phase is of the form:-

$$T(\theta, i) = \frac{\partial}{\partial \theta} W'(\theta, i) \Big|_{i=const.} \quad (2.5)$$

where the co-energy, W' , is calculated from:-

$$W'(\theta, i) = \int_0^i \psi(\theta, i) di \quad (2.6)$$

For a real saturating machine, the flux linkage, and therefore also the co-energy and torque, are non-linear functions of rotor position and phase current. These equations are therefore best solved by numerical methods. However an insight into the principles of machine operation may be obtained by considering an idealised, linear, non-saturating machine. For this machine the flux linkage may be expressed in terms of the machine inductance matrix and phase current vector:-

$$\begin{pmatrix} \psi_1 \\ \psi_2 \\ .. \\ \psi_n \end{pmatrix} = \begin{pmatrix} L_1 & M_{21} & .. & M_{n1} \\ M_{12} & L_2 & .. & M_{n2} \\ .. & .. & .. & .. \\ M_{1n} & M_{2n} & .. & L_n \end{pmatrix} \begin{pmatrix} I_1 \\ I_2 \\ .. \\ I_n \end{pmatrix} \quad (2.7)$$

If this expression for flux-linkage is used in the equations for co-energy and torque given above, then the total electromagnetic torque expressed in matrix format is:-

$$T = \frac{1}{2} \begin{pmatrix} I_1 \\ I_2 \\ .. \\ I_n \end{pmatrix}^t \cdot \frac{\partial}{\partial \theta} \left[\begin{pmatrix} L_1 & M_{21} & .. & M_{n1} \\ M_{12} & L_2 & .. & M_{n2} \\ .. & .. & .. & .. \\ M_{1n} & M_{2n} & .. & L_n \end{pmatrix} \begin{pmatrix} I_1 \\ I_2 \\ .. \\ I_n \end{pmatrix} \right] \quad (2.8)$$

An idealised single-tooth winding switched reluctance machine is usually considered to have negligible mutual coupling between phases, and so for a three phase machine the torque can be simply expressed as:-

$$T = \frac{1}{2} i_1^2 \frac{dL_1}{d\theta} + \frac{1}{2} i_2^2 \frac{dL_2}{d\theta} + \frac{1}{2} i_3^2 \frac{dL_3}{d\theta} \quad (2.9)$$

This is the familiar torque equation for a conventional single-tooth winding switched reluctance machine, with the torque produced from changes in self inductance with rotor position, and proportional to phase current squared.

None of the inductance terms in an idealised fully-pitched winding switched reluctance machine can be considered negligible. For a three phase machine with the windings arranged as shown in Figure 2.3 the phase self inductance's are independent of rotor position (equation 2.1 and Figure 2.4) and therefore do not appear in the simplified torque equation given below:-

$$T = i_a i_b \frac{dM_{ab}}{d\theta} + i_b i_c \frac{dM_{bc}}{d\theta} + i_c i_a \frac{dM_{ca}}{d\theta} \quad (2.10)$$

From this equation it can be seen that torque is produced from changes in mutual inductance with position, and is proportional to the product of phase current pairs. The fully-pitched winding switched reluctance machine could thus be considered to be the dual of a single-tooth winding switched reluctance machine.

2.1.4 Phase Conduction Sequences

In the interests of simplicity the method of torque production with different phase conduction sequences will be examined for idealised non-saturating, three phase switched reluctance machines.

The operation of a conventional single-tooth winding switched reluctance machine will be examined first. Derivation of torque is illustrated in Figure 2.5 for typical unipolar phase currents, with 120° conduction intervals. Each of the three phases is considered to have rising self inductance for $1/3$ of a cycle, and by synchronising the phase current with the period of rising inductance positive torque is produced. Summing the torque contributions of all three phases gives constant shaft torque.

When operating as described above, the MMF in a single-tooth winding switched reluctance machine is concentrated in the slots either side of the stator teeth that are carrying most of the magnetic flux. This pattern of MMF is produced by what can be

considered to be the simplest operating mode for a fully-pitched winding switched reluctance machine - unipolar currents carried in two phases at a time.

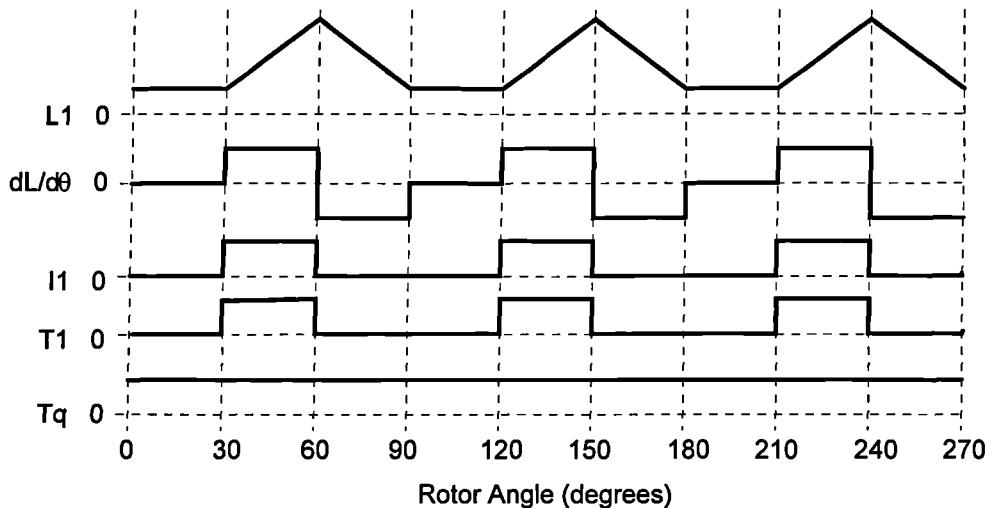


Figure 2.5 Derivation of torque for a single-tooth winding machine with unipolar phase currents.

Unipolar Excitation

The magnetic flux plot for unipolar two phase on excitation is shown in Figure 2.6.

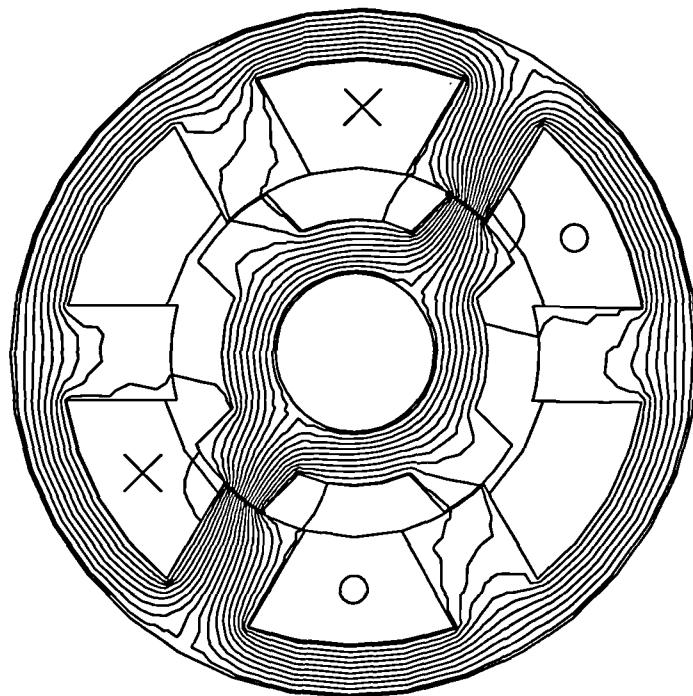


Figure 2.6 Magnetic flux plot for a fully-pitched winding machine with two phase on unipolar excitation.

As with the single-tooth winding switched reluctance machine, the driving MMF is concentrated in the slots either side of the stator teeth carrying most of the flux. For this mode of excitation one electrical cycle corresponds to a rotation of one rotor tooth pitch, and each phase conducts for 240° , with 120° between each phase. Phase currents are synchronised to the rotor position such that whilst the mutual inductance between a pair of phases is rising they are both conducting, as illustrated in Figure 2.7. Positive torque is produced during this interval, with the sum from each pairing of phases giving constant shaft torque.

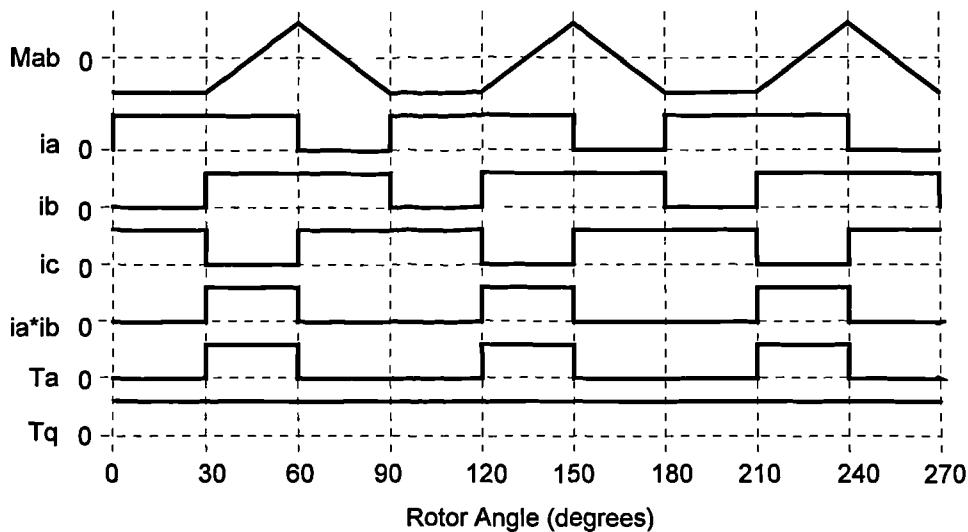


Figure 2.7 Derivation of torque for a fully-pitched winding machine with unipolar two phase on excitation.

Bipolar Excitation

The simplest mode of operation with bipolar phase currents is that with two phases excited at a time. The magnetic flux plot for 'bipolar two phase on' excitation is illustrated in Figure 2.8, this shows that the driving MMF is simultaneously exciting adjacent stator teeth. For bipolar operation one electrical cycle corresponds to a rotation of two rotor tooth pitches. Each phase conducts for 120° during each half cycle, with current reversing every half cycle, as illustrated in Figure 2.9. As with unipolar operation each phase is displaced from the others by 120° . Phase currents are synchronised to the rotor position such that whilst the mutual inductance between a pair of phases is falling they are both conducting. The product of different phase current polarities, and the falling mutual inductance, results in positive torque. Contributions from each pairing of phases adds to give constant shaft torque.

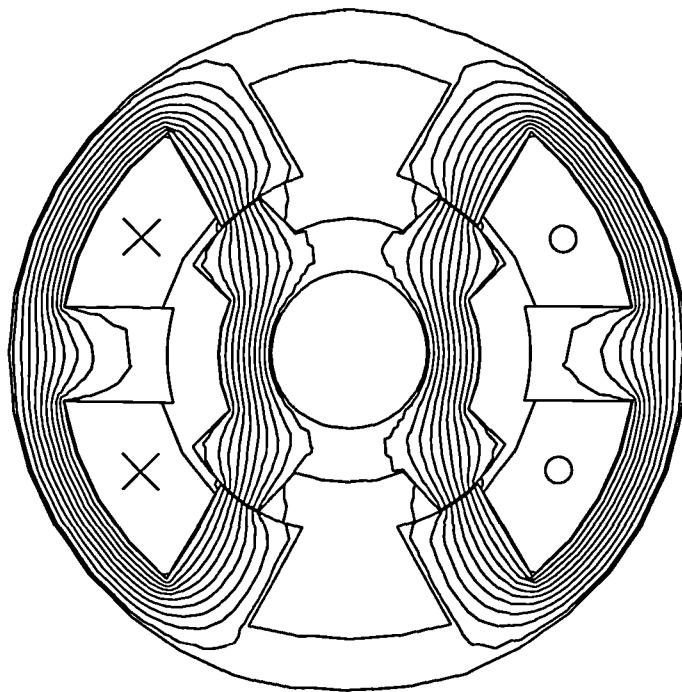


Figure 2.8 Magnetic flux plot for a fully-pitched winding machine with two phase on bipolar excitation.

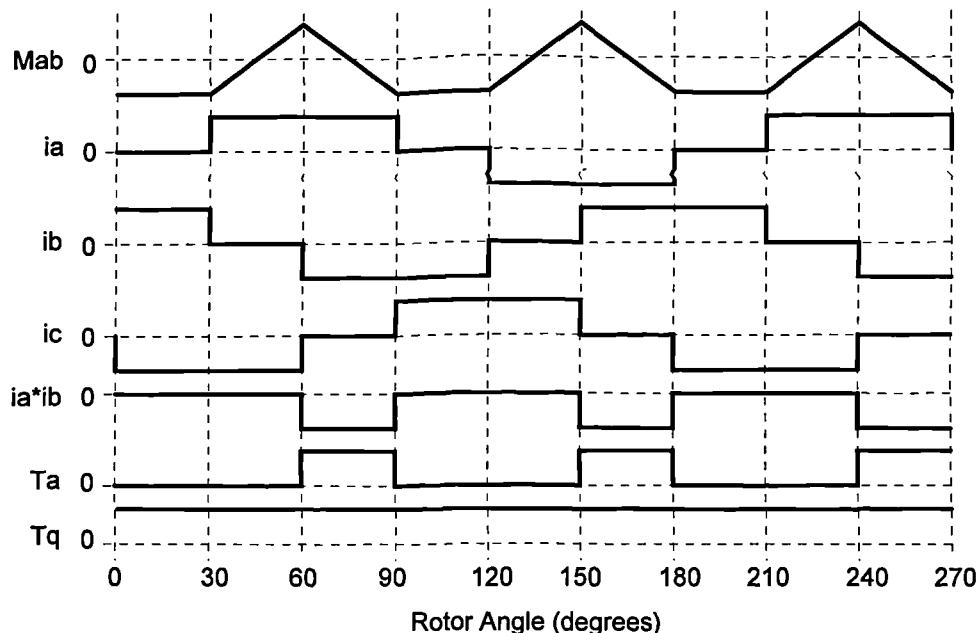


Figure 2.9 Derivation of torque for a fully-pitched winding machine with bipolar two phase on excitation.

For a three phase fully-pitched winding switched reluctance machine it is possible to use all phases simultaneously to produce positive torque. The magnetic flux plot for bipolar three phase on excitation is shown in Figure 2.10, note that the windings in every slot are being used to drive flux around the machine.

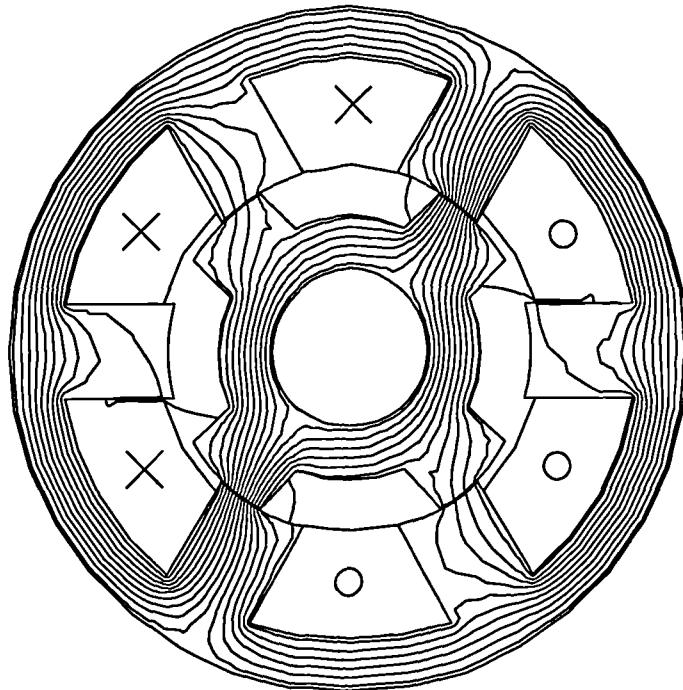


Figure 2.10 Magnetic flux plot for a fully-pitched winding machine with three phase on bipolar excitation.

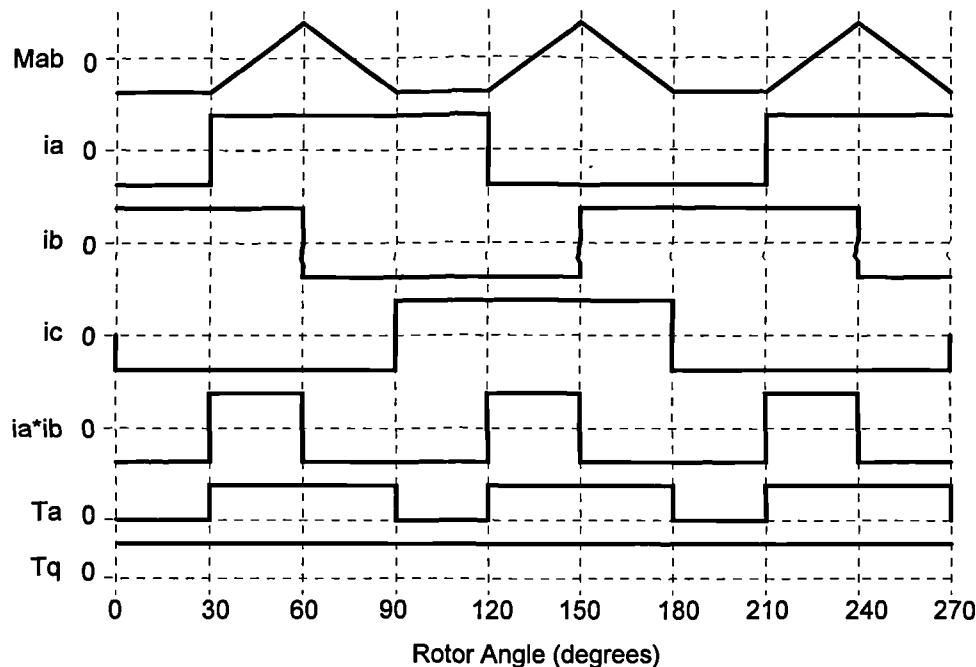


Figure 2.11 Derivation of torque for a fully-pitched winding machine with bipolar three phase on excitation.

In this mode of operation each phase is excited with a bipolar squarewave as shown in Figure 2.11. As with two phase on bipolar operation, one electrical cycle corresponds to a rotation of two rotor teeth. Each phase conducts for 180° during each half cycle, with

the current reversing every half cycle. The conduction sequence in each phase is displaced by 120° . Phase currents are synchronised to the rotor position such that whilst the mutual inductance between a pair of phases is rising they conduct current in the same direction, and whilst the mutual inductance is falling they conduct current in opposite directions. Positive torque is thus produced from periods of rising and falling mutual inductance. Each pair of phases now produces overlapping torque contributions which sum to give constant shaft torque.

To summarise: In the case of unipolar operation, two phases are excited at any one time, each phase conducts for 240° , and one electrical cycle corresponds to a rotation of one rotor tooth pitch. Positive torque is produced from periods of rising mutual inductance only. For bipolar two phase on excitation, two phases are excited at any one time, each phase conducts for 120° during each half cycle, with the current reversing on each half cycle. Each electrical cycle corresponds to a rotation of two rotor teeth, and positive torque is produced from periods of falling mutual inductance only. And finally for bipolar three phase on excitation, all phases are excited all of the time, the current reverses on alternate half cycles, and one electrical cycle corresponds to a rotation of two rotor teeth. Positive torque is produced from periods of rising and falling mutual inductance.

2.2 The Equivalent Machine Transformations

In a real machine, the electromagnetic properties of the phase windings are non-linear functions of rotor position and phase current. These non-linearities make it inappropriate to use inductance and current matrices for the calculation of co-energy and torque. Furthermore, little insight into the method of torque production in a fully-pitched winding switched reluctance machine is gained from the use of inductance parameters, as it is not easy to predict from these which pairs of stator and rotor teeth are being excited to contribute useful torque. A fully-pitched winding switched reluctance machine is, by its nature, electromagnetically a closely coupled machine, so to solve the circuit equations of an m phase machine, m coupled non-linear differential equations must be solved simultaneously. This makes simulation and analysis of fully-pitched winding machines using the actual phase parameters problematical. In contrast, a single-tooth winding switched reluctance machine is usually assumed to have negligible coupling between phase windings (Preston *et al* [2.2]) and so the circuit equation of each phase can be solved independently, from a single non-linear differential equation.

The general concept of using an alternative network in order to simplify analysis was introduced by Kron [2.3]. He showed that a set of transformation, or "connection", matrices could be used to relate circuit parameters in the real network to those in a simplified, but equivalent network. Using the concept of an 'equivalent single-tooth winding switched reluctance machine', a set of transformations can be used to relate flux-linkage and current, in the real fully-pitched winding machine, to those in the equivalent single-tooth winding machine. This enables the phases to be de-coupled, and machine operation can be modelled using techniques used for conventional single-tooth winding machines.

The 'equivalent single-tooth winding switched reluctance machine' must have a magnetic circuit with identical physical dimensions and electromagnetic properties to its fully-pitched winding counterpart. This implies that both machines have the same number of stator and rotor teeth, and each must have the same number of electrical phases. The following assumptions are necessary for the implementation of the transforms:-

- (i) The magnetic field is only a function of the net MMF in each slot, and the actual distribution of conductors within any one slot is of no significance.
- (ii) Saturation in the back iron of both the rotor, and stator, is insignificant. This assumption is commonly used in the simulation of switched reluctance machines, as it enables de-coupled modelling of each phase.

Both current and flux linkage transformations make no assumptions other than those above, and therefore remain valid even when the rotor and stator teeth become saturated.

The transformations relating flux-linkage and current, in a fully-pitched winding switched reluctance machine, to those of an equivalent single-tooth winding switched reluctance machine were first introduced by Mecrow [2.4], with further explanation given by Barrass *et al* [2.5]. The transformations are derived from examination of the MMF per slot, and flux per stator tooth, in the fully-pitched, and equivalent single-tooth winding machines. This is most easily seen by reference to diagrams showing the machine cross-section and winding arrangement. For simplicity a three phase 6:4 machine is considered, although the principle can be applied to any arrangement.

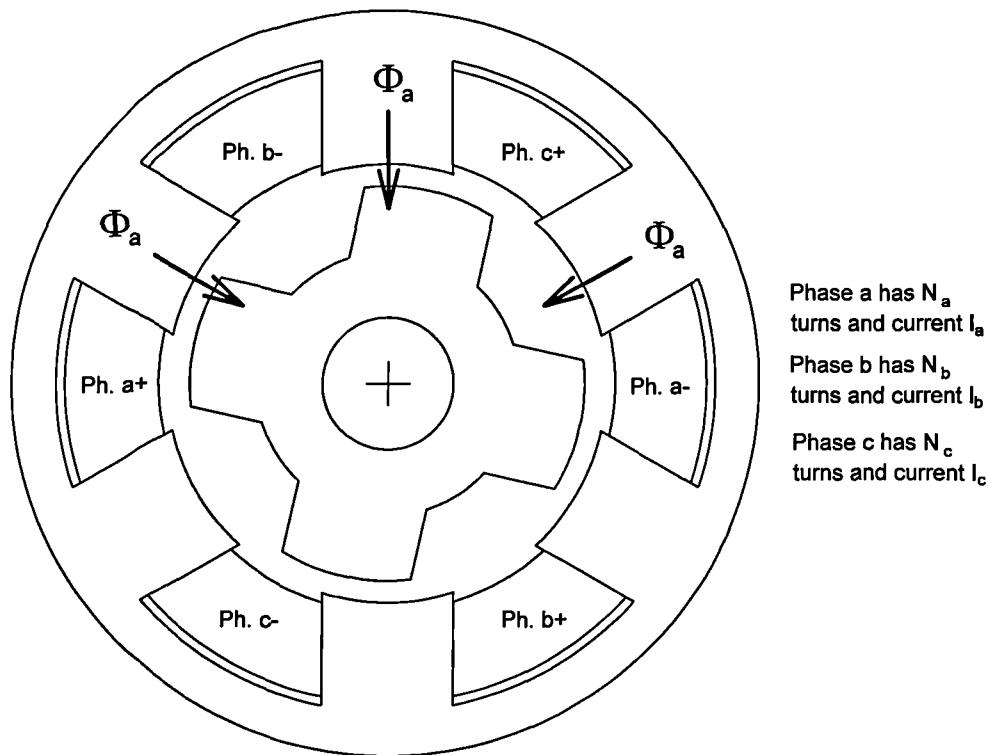


Figure 2.12 Fully-pitched winding machine cross section showing winding arrangement and flux.

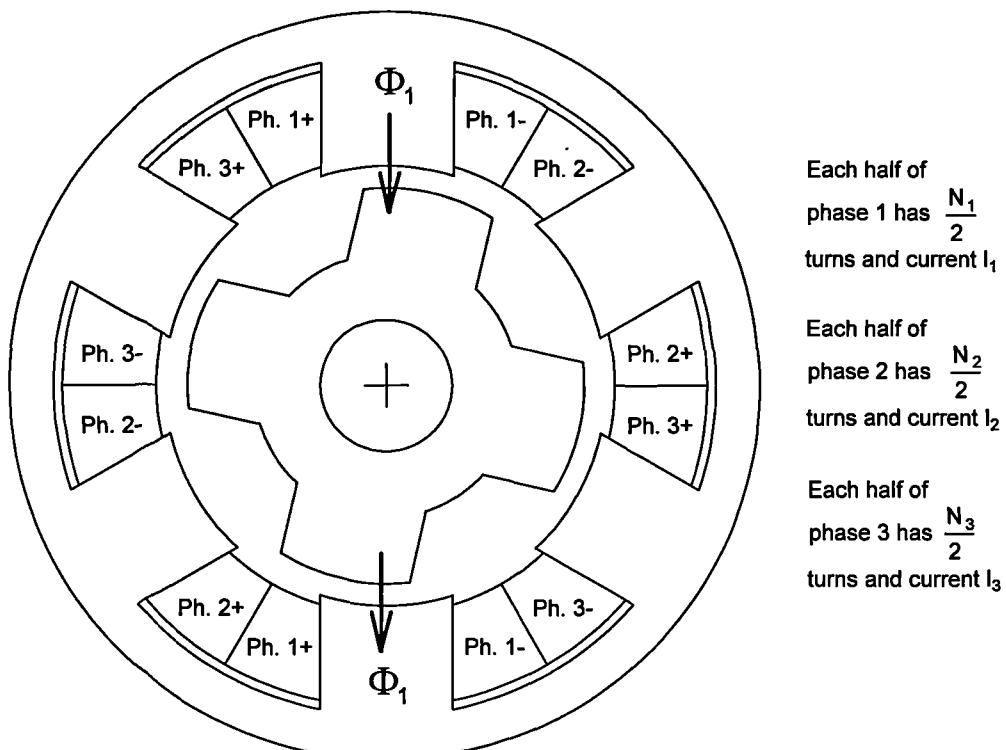


Figure 2.13 Single-tooth winding machine cross section showing winding arrangement and flux.

Considering first the derivation of an equation relating the phase currents in the two machines, based on analysis of slot MMFs. By inspection of Figure 2.12 and Figure 2.13 it can be seen that equating the total MMF in the top right-hand slot of each machine:-

$$-2I_c N_c = I_1 N_1 + I_2 N_2 \quad (2.11)$$

similarly for the next two slots going clockwise around the stator

$$-2I_a N_a = I_2 N_2 + I_3 N_3 \quad (2.12)$$

$$-2I_b N_b = I_1 N_1 + I_3 N_3 \quad (2.13)$$

The factor of two arises because for the fully-pitched winding machine shown, each slot carries all of the turns for one phase, whereas for the single-tooth winding machine, each slot carries only half of the total phase turns.

Now if $N_1 = N_2 = N_3 = N_{123}$, and $N_a = N_b = N_c = N_{abc}$, then expressed in matrix format

$$\begin{pmatrix} I_a \\ I_b \\ I_c \end{pmatrix} = \frac{-N_{123}}{2N_{abc}} \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} I_1 \\ I_2 \\ I_3 \end{pmatrix} \quad (2.14)$$

Now considering the equation relating the flux-linkages, based on the analysis of flux-linkages and the flux per stator tooth. Looking at the axis of fully-pitched phase 'a', if we consider flux passing through this axis and circulating around the machine via the rotor core and stator back iron, it can be seen that the same flux circulating in the equivalent single-tooth winding machine must link the single-tooth windings of phases 1,2 and 3, as the flux passes through the stator teeth. With the rotor in the position shown, it would be expected that most flux would link the phase 1 winding, although this is not necessary for the derivation. Equating stator tooth flux, and the winding(s) which it links:-

$$\Phi_a = \Phi_1 - \Phi_2 - \Phi_3 \quad (2.15)$$

similarly for the other fully-pitched phases

$$\Phi_b = -\Phi_1 + \Phi_2 - \Phi_3 \quad (2.16)$$

$$\Phi_c = -\Phi_1 - \Phi_2 + \Phi_3 \quad (2.17)$$

Now flux-linkage is the product of flux and the number of turns that it links, which using the same equalities for the number of turns results in:-

$$\begin{pmatrix} \Psi_a \\ \Psi_b \\ \Psi_c \end{pmatrix} = \frac{N_{abc}}{N_{123}} \begin{pmatrix} +1 & -1 & -1 \\ -1 & +1 & -1 \\ -1 & -1 & +1 \end{pmatrix} \begin{pmatrix} \Psi_1 \\ \Psi_2 \\ \Psi_3 \end{pmatrix} \quad (2.18)$$

Neglecting the effects of resistive voltage drop, voltage is simply the time derivative of flux linkage, therefore it follows that:-

$$\begin{pmatrix} V_a \\ V_b \\ V_c \end{pmatrix} = \frac{N_{abc}}{N_{123}} \begin{pmatrix} +1 & -1 & -1 \\ -1 & +1 & -1 \\ -1 & -1 & +1 \end{pmatrix} \begin{pmatrix} V_1 \\ V_2 \\ V_3 \end{pmatrix} \quad (2.19)$$

For the experimental 12:8 three phase fully-pitched winding, and single-tooth winding machines, the total number of turns per phase are equal and so $N_{abc} = N_{123}$, substituting this in the above equations gives:-

$$\mathbf{i}_{123} = \mathbf{C} \cdot \mathbf{i}_{abc} \quad (2.20)$$

and

$$\Psi_{123} = \mathbf{C}^{-1} \cdot \Psi_{abc} \quad (2.21)$$

where

$$\mathbf{C} = \begin{pmatrix} +1 & -1 & -1 \\ -1 & +1 & -1 \\ -1 & -1 & +1 \end{pmatrix} \text{ and it's inverse } \mathbf{C}^{-1} = \begin{pmatrix} 0 & -\frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & 0 & -\frac{1}{2} \\ -\frac{1}{2} & -\frac{1}{2} & 0 \end{pmatrix} \quad (2.22)$$

It should be noted that the transformation matrix \mathbf{C} and it's inverse do not exist for all fully-pitched winding configurations. For a winding arrangement similar to that shown, i.e. with the phase windings concentrated in single slots, then the matrix \mathbf{C} and it's inverse always exist. However, if the windings of a fully-pitched phase are distributed in several slots, then the transformation matrices may be singular.

Extensive use is made of these transformation matrices when analysing the experimental fully-pitched winding machine. By transformation to the equivalent single-tooth winding machine, the fully-pitched winding machine can be simulated, static characteristics predicted, and a much deeper understanding of the operating waveforms and mechanisms gained.

2.3 Comparison of Fully-Pitched Winding and Single-Tooth Winding Switched Reluctance Machines

One of the most important questions to be answered when a new machine type is introduced is: "Does the new machine type offer any advantages over existing types?". It is rare that there is a simple answer to this question; as with so many things in engineering there are many trade-offs to be balanced, and the answer to the general question depends on what aspects of performance are compared, and under what conditions.

No attempt will be made to make a detailed comparison of the fully-pitched winding switched reluctance machine to what could be considered the main rivals for switched reluctance machines in the industrial market place; DC machines, induction motors, permanent magnet synchronous machines, and synchronous reluctance motors. Several of these machine types have been compared to conventional single-tooth winding switched reluctance machines by Miller [2.6] and Harris *et al* [2.7].

When comparing a fully-pitched winding switched reluctance drive to a single-tooth winding switched reluctance drive there are several performance indexes to be considered; machine performance, cost factors such as copper volume, and the rating of the power converter required to drive the machine. These aspects will be explored in the following sections.

2.3.1 Phase Winding Lengths

One of the advantages of conventional single-tooth winding switched reluctance machines is that they have a simple winding arrangement, with short end-windings (Fulton *et al* [2.8]) and a coil shape that lends itself to winding off the machine. The short end windings result from each coil being wound around a single stator tooth only, the end-windings therefore only span the width of a single stator tooth. This is in contrast to a fully-pitched winding switched reluctance machine, which has coils whose end-windings span as many stator teeth as there are electrical phases. The end-windings are therefore much longer than those in a single-tooth winding machine, and are similar in appearance to those in an induction motor. It is however possible to reduce the end-winding length for a given machine diameter, and number of phases. If there is only one winding per phase then the end-winding must span 180° mechanically, as in a three phase 6:4 SRM. But by doubling the number of stator and rotor teeth, e.g. from 6:4 to 12:8, the end-winding only needs to span 90° mechanically (c.f. changing from a 2 pole, to 4 pole winding in an induction motor), thereby reducing the end-winding length.

Figure 2.14 shows how the winding lengths of fully-pitched winding and single-tooth winding machines may be compared.

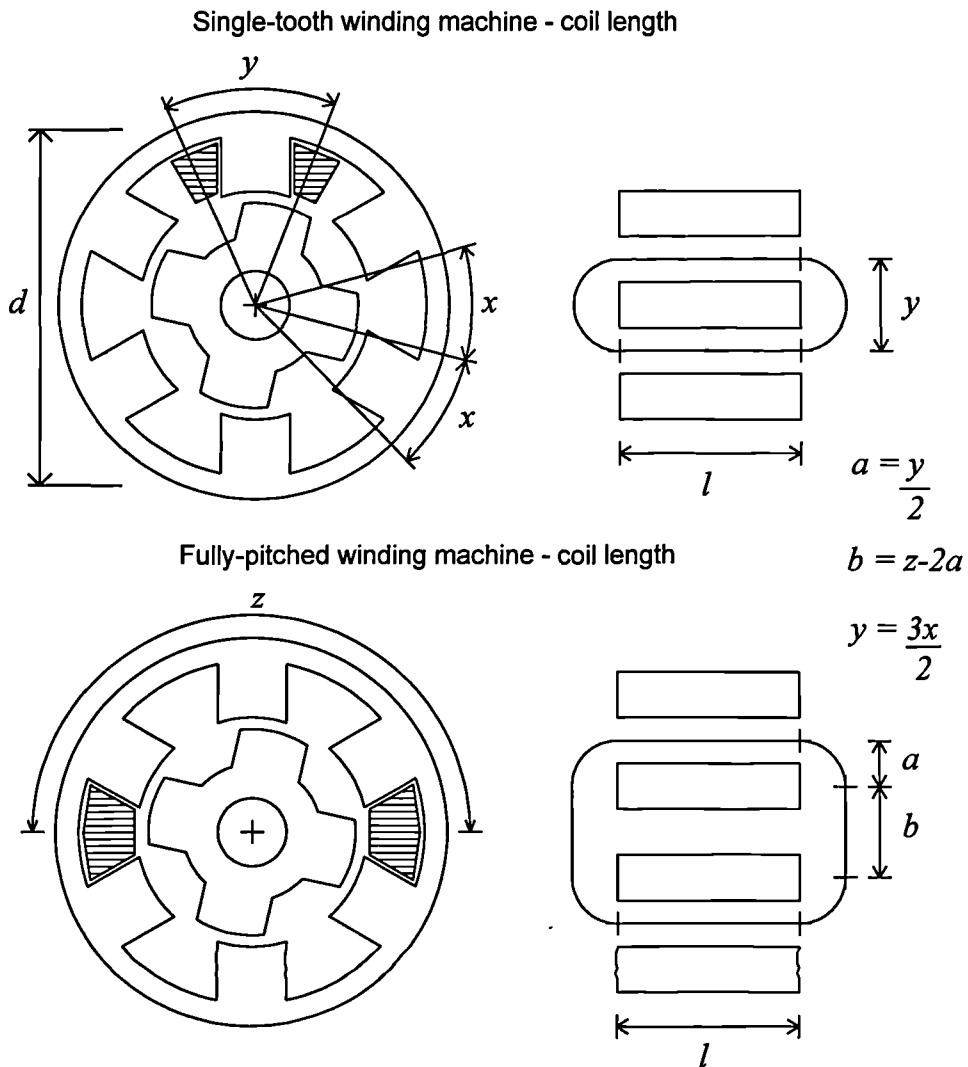


Figure 2.14 Estimation of machine phase winding lengths.

If it is assumed that the stator tooth widths and tooth gaps are equal, then the arc length x is given by:-

$$x = \frac{\pi \cdot d}{2n_s} \quad (2.23)$$

With reference to Figure 2.14 the length of one turn of a single-tooth coil can be estimated from:-

$$L_{st} = 2l + \frac{3\pi^2 d}{4n_s} \quad (2.24)$$

and for a fully-pitched coil from:-

$$L_{fp} = 2l + \pi \cdot y + 2b \quad (2.25)$$

which substituting values for a and b gives:-

$$L_{fp} = 2l + \frac{\pi d}{n_s} \left(2m + \frac{3\pi}{4} - \frac{3}{2} \right) \quad (2.26)$$

where l is the stack length, and d is the outside diameter. From these equations comparisons can be made for different machine dimensions and tooth numbers, Table 2.1 gives winding length estimates and ratios for three phase machines, with different stack lengths and tooth numbers.

Machine type	Tooth number	Winding length to stack length ratio			Active length to total winding length ratio		
		$l=0.5d$	$l=d$	$l=2d$	$l=0.5d$	$l=d$	$l=2d$
Single-tooth winding SRM	6:4	4.46	3.23	2.62	0.448	0.618	0.764
	12:8	3.23	2.61	2.31	0.618	0.764	0.866
Fully-pitched winding SRM	6:4	9.17	5.58	3.79	0.218	0.358	0.527
	12:8	5.58	3.79	2.89	0.357	0.527	0.690
Ratio of fully-pitched to single-tooth	6:4	2.05	1.73	1.45	0.486	0.578	0.689
	12:8	1.73	1.45	1.25	0.578	0.687	0.797

Table 2.1 Machine winding length ratios.

Note that for the values given above, the active winding length is taken to be that which rests in the slot itself, in other words the whole of the end-windings are taken to be inactive in terms of torque production. For the experimental three phase 12:8 machines, which have a stack length to outside diameter ratio of approximately 0.93, the ratio of winding lengths is 1.44. This figure corresponds very well with the estimated figure of 1.45 for $l = d$ given in the table.

It is clear from the figures in Table 2.1 that the additional end-winding length associated with the fully-pitched windings can extract a large penalty in terms of winding length. For example, a short ($l=0.5d$) 6:4 machine is predicted to have a winding length double that of its short-pitch wound equivalent. However, by using a long machine ($l=2d$), with a four pole magnetic configuration the longer end-windings only add 25% to the overall length. Clearly if any advantages of the fully-pitched winding switched reluctance machine are to be exploited, it would be advantageous to have a stack length greater than the diameter and four pole magnetic circuit configuration.

2.3.2 Theoretical Copper and Torque Density

In a conventional single-tooth winding switched reluctance machine with one phase excited, the total MMF is given by:-

$$\mathcal{J}_{st} = N_1 I_1 \quad (2.27)$$

For a fully-pitched winding switched reluctance machine with two phases excited, the total MMF is the sum of that produced by the two phases:-

$$\mathcal{J}_{fp} = N_b I_b + N_c I_c \quad (2.28)$$

If the two machines have an equal number of turns per phase, and the two phase currents in the fully-pitched winding machine are equal, then to obtain the same MMF in each case the phase currents must have the following relationship:-

$$I_b = I_c = \frac{I_1}{2} \quad (2.29)$$

Using the same diameter wire in each case will result in the same overall slot fill-factor, assuming the machines have the same dimensions. For equal MMFs, the current density in the windings of the fully-pitched machine must therefore be half of that in the single-tooth winding machine. As the current density in the windings is halved, the ohmic losses per unit length are reduced by a factor of four. Under these conditions the copper could be considered to be under-utilised. However there are two phases excited at a time in the fully-pitched winding machine, and each winding is longer. Assuming that for equal MMFs the fully-pitched and single-tooth winding machines give the same torque, then the relative copper losses for the same shaft torque are given by:-

$$\frac{P_{abc}}{P_{123}} = \frac{R_{abc}}{2R_{123}} \quad (2.30)$$

The output available from a given machine size is determined essentially by its thermal characteristics. Comparisons made on the basis of equal copper losses are therefore appropriate for the fully-pitched and equivalent single-tooth winding machines, as they are mechanically equivalent, differing only in winding configuration. For equal copper losses the ratio of phase current, between a single-tooth winding machine with one phase excited, and a fully-pitched winding machine with two phases excited is given by:-

$$\frac{I_{abc}}{I_{123}} = \sqrt{\frac{R_{123}}{2R_{abc}}} \quad (2.31)$$

In general, if there are m phases of the fully-pitched winding machine excited the phase current ratio is:-

$$\frac{I_{abc}}{I_{123}} = \sqrt{\frac{R_{123}}{m \cdot R_{abc}}} \quad (2.32)$$

Thus as the number of conducting phases increases, the current per phase goes down by the square root of this number. However, the product of conducting phases and the current per phase, goes up with the square root of the number of conducting phases. Thus providing the additional MMF available can be utilised, the shaft torque will increase with the number of phases conducting for a given total copper loss.

A better understanding of the effects of winding length and copper mass can be obtained by comparing the total MMF available in fully-pitched and single-tooth winding machines which have the same dimensions, and are constrained to have equal copper losses. For machines with windings of N turns, each having length l , and wound with wire of area A , the ratio of copper volumes between the two machines is given by:-

$$\frac{Vol_{abc}}{Vol_{123}} = \frac{N_{abc}l_{abc}A_{abc}}{N_{123}l_{123}A_{123}} = k_{cu} \quad (2.33)$$

The turn length ratio is defined by:-

$$\frac{l_{abc}}{l_{123}} = k_{lpt} \quad (2.34)$$

Where the subscripts 'abc' and '123' relate to fully-pitched, and single-tooth winding machines respectively. Assuming that the resistivity of the wire used in each machine is the same, then the ratio of phase resistances is:-

$$\frac{R_{abc}}{R_{123}} = \frac{N_{abc}l_{abc} \cdot A_{123}}{N_{123}l_{123} \cdot A_{abc}} = \frac{k_{cu}}{k_{lpt}^2} \quad (2.35)$$

Substituting the resistance ratio above into equation 2.32, and multiplying by the number of turns per phase to obtain the ampere turns, the MMF ratio is given by:-

$$\frac{\mathcal{I}_{abc}}{\mathcal{I}_{123}} = \sqrt{\frac{m \cdot k_{cu}}{k_{lpt}^2}} \quad (2.36)$$

By allowing the copper volume to increase in proportion to the turn length ($k_{cu} = k_{lpt}$), for example by using the same diameter wire and number of turns per phase, then the MMF ratio will go down with the square root of turn length ratio. If however the machine is constrained to have the same copper volume ($k_{cu} = 1$), then the MMF ratio will go down in direct proportion to the turn length ratio. This more pessimistic relationship would not occur in reality for the following reason. Due to the longer end-winding length in the fully-pitched winding machine, for equal copper volumes the proportion of copper in the slots must be reduced (other factors remaining constant). This in turn will reduce the slot fill-factor resulting in a sub-optimum machine design. To correct for the reduced slot fill-factor, the rotor and stator diameter split ratio may be increased by reducing the stator slot depth. The resulting increase in rotor diameter will lead to more torque for a given MMF.

2.3.3 The Experimental Machine Copper and Torque Density

Using equations given in section 2.3.2 the experimental fully-pitched winding and equivalent single-tooth winding switched reluctance machines can be compared. The experimental fully-pitched winding switched reluctance machine uses the same diameter of wire and has the same number of turns per phase as the experimental single-tooth winding switched reluctance machine, but the increased winding length gives phase resistance of 1.147Ω (at 20°C) compared to 0.797Ω . Comparing the copper losses in the two machines when producing the same MMF and torque, using equation 2.30, the loss ratio is:-

$$\frac{P_{abc}}{P_{123}} = \frac{1.147\Omega}{2 \times 0.797\Omega} = 0.72 \quad (2.37)$$

So although the winding resistance is 44% higher, for an equal MMF per-slot, the copper losses in the fully-pitched winding machine are only 72% of those in the single-tooth winding machine.

Substituting phase resistance values into equation 2.32 to calculate the phase current ratios, the torque production for equal copper loss can be compared. Phase current ratios are 0.589 and 0.481, for two and three phase excitation respectively. These figures were used by Mecrow [2.4] to compare torque production for equal copper losses at several current levels, with results summarised in Table 2.2.

Single-tooth winding phase current (A)	Single-tooth winding machine peak torque (Nm)	Fully-pitched winding machine increase in peak torque relative to single-tooth winding machine (%)		
		Unipolar 2 phase on	Bipolar 2 phase on	Bipolar 3 phase on
10	24.7	44	58	91
15	48.2	38	41	59
20	72.5	34	34	39

Table 2.2 Measured experimental machine torque ratios for equal loss.

These figures show that compared on the basis of equal loss, the experimental fully-pitched winding machine shows gains in torque of between 34% and 91% over the single-tooth winding machine. This increase in torque is however at the expense of increased copper mass, up from 8.34kg to 12.0kg (an increase of 44%) which in turn has increased the total active mass (windings plus laminations) by approximately 9%. Thus in terms of peak torque per unit copper mass, the fully-pitched winding machine produces between 93% and 132% of the single-tooth winding machine torque, and compared on the basis of torque per unit total active mass the figures range from 123% to 175%.

Comparing the experimental machine MMF ratios with equation 2.36, using copper weight and turn length ratios of 1.44, the MMF ratio is 1.18 for two phase on excitation, and 1.44 for three phase on excitation. If the number of turns per phase in the fully-pitched winding machine was reduced to give a copper volume ratio of unity, the MMF ratio would fall to 0.98 and 1.20, for two and three phase on excitation respectively. It should be noted however that this is for the sub-optimum condition of only partially filled slots.

2.3.4 Power Converter Rating

It is difficult to derive an analytical equation for the power converter rating of a switched reluctance drive, because for any real machine the flux linkage and torque are non-linear functions of phase current and rotor position. In order to give an accurate figure for converter rating it is necessary to take measurements from a prototype drive, or simulate drive operation using measured or calculated motor characteristics. Converter rating, expressed in terms of the number of switching devices and their voltage and current ratings, is also critically dependant on the number of motor phases, and the number of switching devices per phase. Furthermore, the actual device rating relative to the machine phase current is dependant on the device type, as each has different conduction and switching loss characteristics.

Drives using Single-Tooth Winding Switched Reluctance Machines.

Work by Miller [2.6] assumes that at base speed the phase current will just maintain a 'flat topped' profile, and that under these conditions the rate of change of flux linkage during the constant current region is equal to the dc-link voltage. A piecewise linear model is used for the machine flux linkage characteristics in order to derive an expression for the converter VA required. When substituting various constants for a prototype 5hp drive Miller obtained an estimated figure of 11kVA/kW, based on peak current and voltage. From measurements on the actual drive Miller obtained a figure of 10.5kVA/kW, an improvement of 5%. In other work by Harris *et al* [2.7] and Ray *et al* [2.9], which both consider a four phase 7.5kW SR drive, figures reported for peak converter rating are 11.2kVA/kW and 9.7kVA/kW respectively. These figures were derived from simulations or actual measurements, and the variation of 15% between them for drives with the same 'nameplate' rating illustrates the variation small changes to machine design and control strategy can make.

Each of the three papers above also give ratings for inverter fed induction motors. Figures for peak converter rating vary between 9.25kVA/kW and 11.4kVA/kW, depending on values assumed for the power factor, and ratio of fundamental to peak current. The variation of 23% is also partly explained by an efficiency factor used by Ray *et al* (η in equation 6) to take into account operation from a PWM converter that is not used by Miller. Taking this into account the figures range from 10.4kVA/kW to 11.4kVA/kW, a variation of 10%. Now comparing the average value for induction motor drives of 10.9kVA/kW, with the average value for switched reluctance drives of 10.45kVA/kW, it appears that the switched reluctance drive may have a small advantage. The switched reluctance drive power circuits considered in these three papers have one controlled switch per phase, in this case a total of four. By comparison

a standard induction motor drive requires a total of six controlled switches. Thus in addition to a small advantage in overall device rating, the switched reluctance drives considered require fewer controlled switches and their attendant drive and protection circuitry.

Drives using Fully-Pitched Winding Switched Reluctance Machines.

If we now consider a fully-pitched winding switched reluctance drive, the converter rating relative to a conventional single-tooth winding switched reluctance drive depends very much on the mode of operation. For unipolar operation the same power circuit configuration can be used in each case, there is thus no penalty in device number and the converter rating is determined by details of the phase current and voltage waveforms, which will be investigated later. For bipolar operation the number of devices depends on circuit configuration and current waveform required. For fully independent control of phase currents a full 'H' bridge of four devices per phase, or a single bridge leg of two devices and a split supply will be required. Note that with a split supply the effective phase voltage is halved and so the phase current, and hence device rating, needs to be doubled for the same shaft power. Both these circuit configurations require a total device peak VA rating twice that of unipolar drives, for the same peak voltage and current. By using a star or delta connection, only two devices per phase are required to obtain bipolar phase currents. The connection topology does however place constraints on the phase current waveshapes which may not be beneficial to torque production. It should be noted that all bipolar circuit configurations have at least two controlled switches per phase, but for unipolar operation it is possible to have a single controlled switch per phase, thus reducing the amount of drive and protection circuitry required.

Unipolar Fully-Pitched Winding Switched Reluctance Drives

The effects of the actual current and voltage waveforms on the total device VA rating can only be determined accurately by simulation or measurement. However it is possible to make some approximations which allow figures for fully-pitched winding and single-tooth winding switched reluctance drives to be compared. It should be emphasised that these approximations can only be used as a rough guide. For simplicity we shall consider three phase machines and use waveform approximations that are representative of the experimental fully-pitched winding switched reluctance drive. The ratio of turns per phase between fully-pitched and single-tooth machines does not effect the VA requirements, as the current and flux-linkage scale inversely for a given output, the turns ratio will therefore be taken as unity.

For two phase on unipolar operation at low speed only half the current per phase is required in the fully-pitched winding machine to produce the same MMF, but each phase is excited for twice as long. The average phase current supplied by the converter is thus equal. As the peak to peak flux linkage in the fully-pitched winding machine is twice that in the single-tooth winding machine, the product of peak current per phase and peak to peak flux linkage are equal. From this it may appear that the VA ratings are also likely to be equal. However, this does not take into account phase conduction overlap which becomes significant approaching base speed.

For a flat-topped current profile at base speed, as described by Miller for a conventional single-tooth winding switched reluctance machine, the corresponding waveforms for a fully-pitched winding switched reluctance machine are shown in Figure 2.15. The current and voltage waveforms illustrated for the fully-pitched winding machine are calculated from the equivalent single-tooth winding machine waveforms using the transformation equations given in section 2.2.

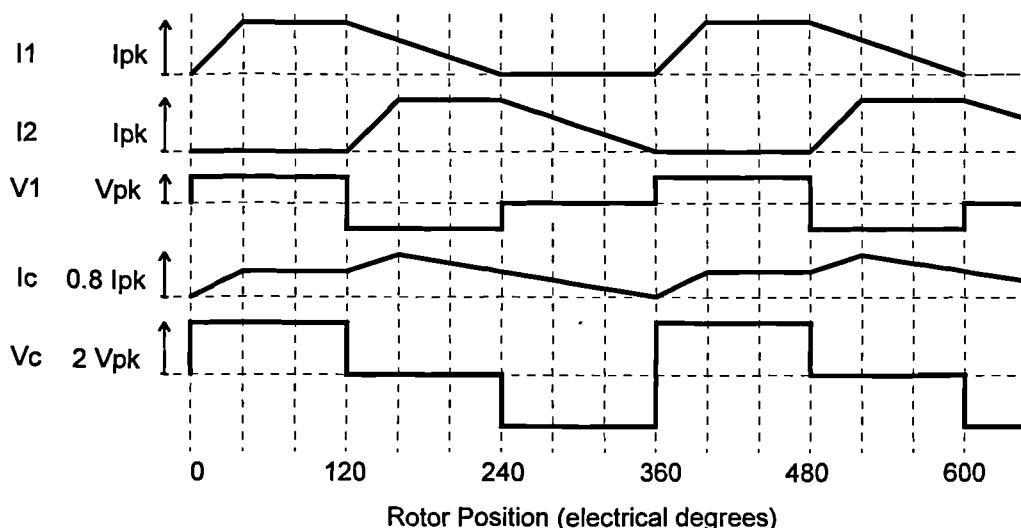


Figure 2.15 Comparison of simplified voltage and current waveshapes at base speed.

The effects of phase current overlap in the single-tooth winding machine show up as a peak in the middle of the fully-pitched winding machine phase current waveform. This increases the peak fully-pitched winding current to approximately 0.8 times that in the single-tooth winding, compared to a factor of 0.5 with no overlap. Peak VA requirement per device is thus 1.6 times greater in the fully-pitched winding switched reluctance drive. This very pessimistic figure is not realised in practice because the fully-pitched winding machine controller would be programmed to reduce the peak phase current by application of negative phase voltage. Also, the maximum flux linkage excursion is not being utilised in either machine as the phase voltage is zero for one third of a cycle. When examining factors due to incomplete phase voltage modulation, simple direct

comparisons are not possible if the applied voltage pulse width is increased beyond 120° .

This is because the modes of operation for fully-pitched winding and single-tooth winding switched reluctance drives are different, accurate comparisons of converter switch VA are thus only possible by simulation or measurement, due to the complex shapes of the phase current and flux linkage-current loci.

Bipolar Fully-Pitched Winding Switched Reluctance Drives

The situation for bipolar operation is even more complex than for unipolar operation because of the wide variety of possible excitation patterns. A feature of many bipolar excitation patterns is that the corresponding equivalent single-tooth winding phase current has a conduction period greater than 180° , thus there must be times when pairs of teeth are excited so as to oppose the desired torque. Another feature of overlapping equivalent single-tooth winding current periods is that they produce overlapping periods of flux linkage, which have the effect of increasing the peak fully-pitched winding flux-linkage. The overlapping periods of current and flux, with the additional complication of negative torque contributions, means that without detailed machine characteristics it is not possible to compare the device VA requirements of a bipolar fully-pitched winding switched reluctance drive, with those of a single-tooth winding switched reluctance drive. As different switching devices (or pairs of devices) supply positive and negative phase current, over a complete cycle the average current per device is halved, mitigating to some extent the requirement for twice as many controlled switches, although the peak current rating remains unchanged.

For sinusoidal operation comparisons may be drawn with the operation of synchronous reluctance motors. The power factor, and hence VA requirement for a given output, is dependant on the ratio of direct and quadrature axis reactance. Switched reluctance machines are designed to be operated with parts of the magnetic circuit heavily saturated at rated output. This reduces the effective direct axis inductance, saliency ratio and hence power factor.

The above discussion illustrates that without simulation or measurement, it is difficult to predict with any accuracy the converter switch VA rating required for a fully-pitched winding switched reluctance drive, due to machine non-linearity and complex phase current and flux linkage waveforms. It is however possible to draw some broad conclusions. For unipolar converters the total switch peak VA rating for fully-pitched winding switched reluctance drives, and single-tooth winding switched reluctance drives, is likely to be similar. Comparing fully-pitched winding switched reluctance drives with unipolar or bipolar converters, the penalty of having twice as many

controlled switches is likely to be greater than gains from increased phase utilisation, and so overall will require a greater total switch peak VA rating.

2.4 Summary and Conclusions

The concept of a fully-pitched winding switched reluctance machine has been introduced and its winding arrangement illustrated. Simple modelling of phase inductance values shows that for a three phase machine the self inductance is constant, whilst the mutual inductance varies with rotor position. The fully-pitched winding switched reluctance machine can thus be considered the dual of a single-tooth winding switched reluctance machine as torque is derived from changes in mutual, rather than self inductance. Basic current waveshapes and slot excitation patterns required for torque production have been shown. Transformation equations allowing electrical parameters in the fully-pitched winding switched reluctance machine to be related to those in an equivalent single-tooth winding switched reluctance machine have been stated. By de-coupling the windings of a fully-pitched winding machine, these transformations allow the machine to be modelled as a conventional single-tooth winding machine. Comparisons between fully-pitched winding and single-tooth winding switched reluctance machines show that the longer end windings in a fully-pitched winding machine increase the phase resistance and hence loss per ampere. However based on equal copper loss the experimental fully-pitched winding machine is able to produce more MMF, and hence torque. In terms of converter rating detailed analysis is required to differentiate single-tooth winding and fully-pitched winding drives, with unipolar currents, but bipolar currents are likely to require a higher kVA/kW.

Chapter 3

Drive Implementation

3.1 Implementation of the Experimental Drive

An electric drive system comprises of four main elements; the control electronics, the electrical power converter, the electrical machine, and a mechanical load. For experimental purposes the mechanical load is provided by another electrical machine coupled to the shaft of the experimental machine. Each of these systems are described in the following sections. More detailed circuit diagrams and mechanical details can be found in Appendix A for the experimental switched reluctance machines, Appendix B for the mechanical loading rig, and Appendix C for the control and converter circuits.

3.1.1 Experimental machines and loading arrangement

Two switched reluctance machines are used for all the experimental and simulated results in this thesis; a fully-pitched winding switched reluctance machine, and its equivalent single-tooth winding switched reluctance machine. These are referred to as the 'fully-pitched winding machine' and the 'single-tooth winding machine'. The single-tooth winding machine is produced by Allen-West in the UK, to a design licensed from SR Drives Ltd. It is an Allen-West model SR75, which is a doubly-salient machine with 12 stator poles and 8 rotor poles. It has a three phase winding, and when mated with the Allen-West drive electronics is rated at 7.5kW, providing 47Nm at 1500rpm. The experimental fully-pitched winding switched reluctance machine was produced by taking another Allen-West model SR75, and replacing the conventional three phase single-tooth windings with a set of three phase fully-pitched windings. Wire of the same diameter was used, and the total number of turns per phase was also kept the same. The fully-pitched arrangement results in longer end-windings but these were accommodated within the original end castings, without any modifications to the stator stack or housing. The experimental fully-pitched winding and equivalent single-tooth winding machines are thus identical, apart from the winding configuration, and small differences due to manufacturing tolerances, and lamination magnetic properties.

To enable the experimental machine to be mechanically loaded another electrical machine is used. This has the advantage over a friction or powder brake that the torque can be easily controlled over a wide speed range, the output power need not be dissipated as heat, and the switched reluctance machine can also be driven as a generator. A separately excited, brushed DC machine was chosen as the load machine as it is easily controlled, and readily available. The load machine is electrically connected to another DC machine which forms part of a Ward-Leonard set. This arrangement provides a low impedance supply to the DC load machine, and allows power to flow in either direction. The DC machine on the Ward-Leonard set is mechanically coupled to an AC induction motor, supplied from a fixed line voltage and frequency. Figure 3.1 illustrates this general arrangement.

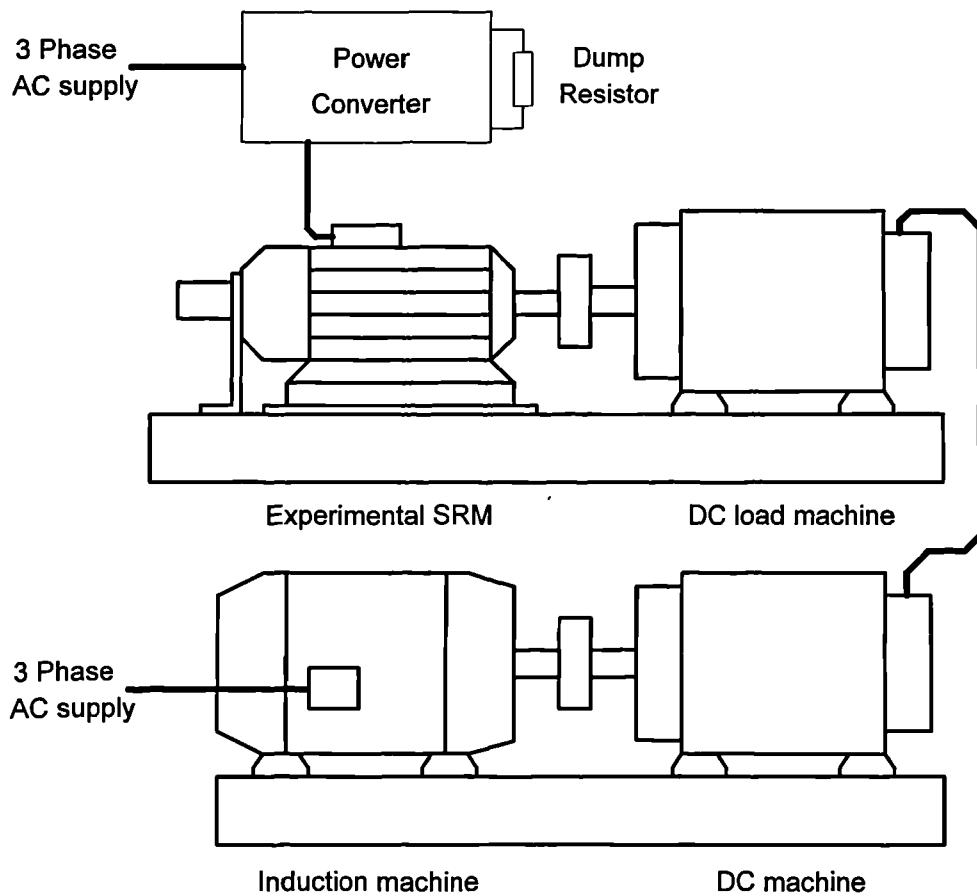


Figure 3.1 Experimental machine loading arrangement.

Each DC machine has a separately excited field circuit (see Appendix B) with the armature circuits inter-connected. As shown, each pair of machines are mounted on a common bed-plate. When the experimental machine is motoring power circulates from the supply, through the SR machine, the two DC machines, and the induction motor, back to the supply. Power is then only required to meet system losses. As it was anticipated that the majority of the time would be spent with the SR machine motoring,

the SR drive power converter was built with a single quadrant front end. Thus when acting as a generator, power from the SR machine is dissipated in the dc-link dump resistors, and the induction motor must supply the SRM shaft power (dissipated in the dc-link dump resistors) and system losses.

Metering of load circuit parameters, in conjunction with knowledge of load machine and mechanical parameters, allows the switched reluctance machine shaft power to be calculated. Details of the loading rig parameters and torque calculation methods can be found in Appendix B.

3.1.2 Control Systems and Power Electronics

The basic design philosophy behind the control and power circuits was to produce a system that was flexible enough to implement all the anticipated machine excitation schemes, without requiring hardware modifications. To this end control functions are implemented in software and the power circuit has a full 'H' bridge per phase to give full independent control of phase currents.

Figure 3.2 shows the general arrangement of the control and measurement equipment used for tests on the experimental machine. Most of the control circuits are mounted in a 19" rack cabinet, with connections to the host PC, VDU terminal, and power converter.

Host System and Drive Control Processors

Although not part of the real time control system, the host personal computer (an IBM compatible DOS machine) acts as a file store and allows programs to be written, compiled, and downloaded to the control electronics. A serial connection to the oscilloscope allows waveforms to be downloaded to the PC for storage and analysis.

There are three microprocessors used to control the switched reluctance machine, two of which have a direct connection to the host PC. The main control board and VME bus master is a Motorola 68030 board. This has an RS232 serial connection to the host PC, and another serial link to a VDU monitor for command entry and de-bugging. The second processor with a direct link to the host PC is the Texas TMS320C31, this has a proprietary high speed serial link which provides full downloading and symbolic de-bugging facilities. The third processor, a Texas TMS320E14, does not have a direct connection to the host PC. Code is downloaded to it via the 68030 which has access to the TMS320E14 memory area. All three processors are connected together on a VME bus which also allows access to some the input/output (I/O) circuits.

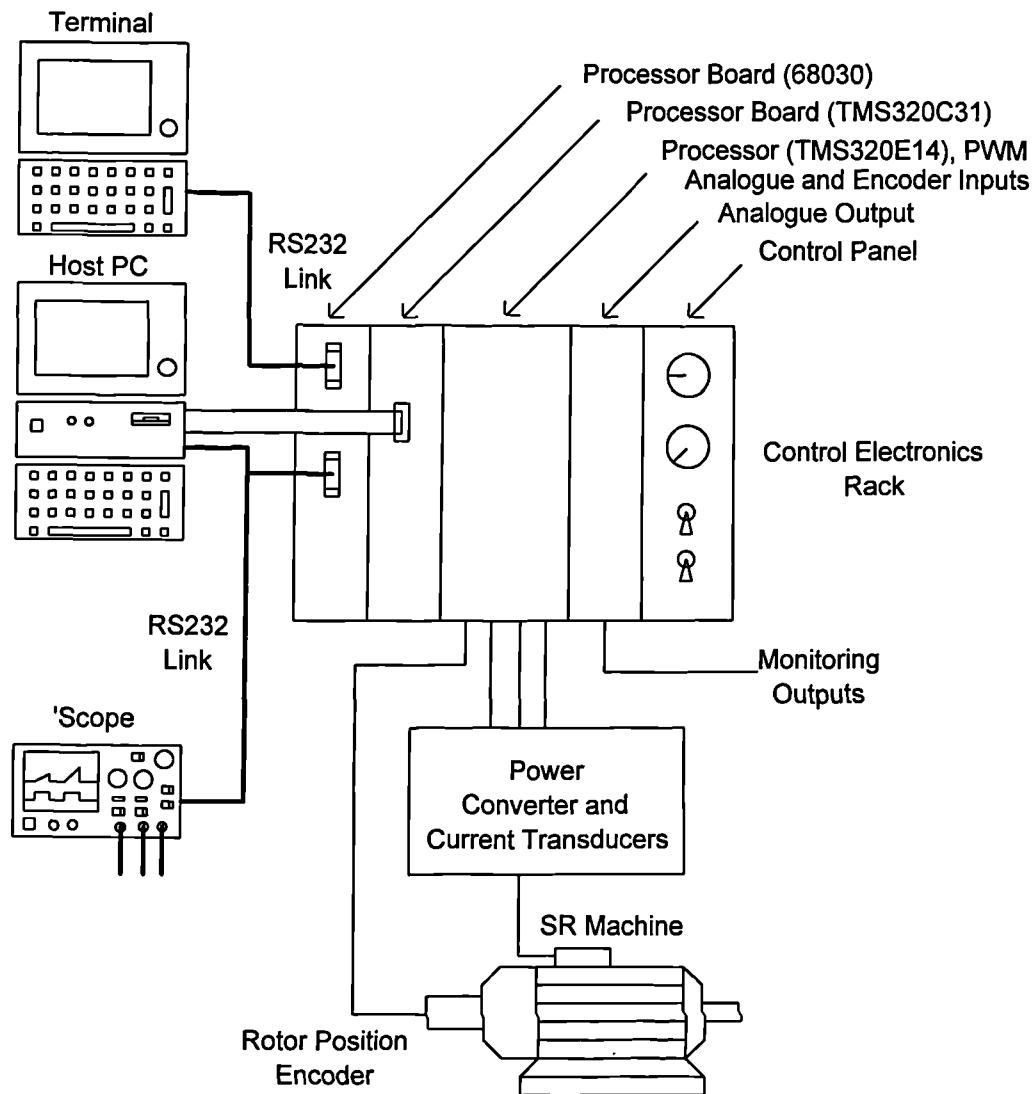


Figure 3.2 Arrangement of control and power equipment.

Control System Inputs and Outputs

There are three main inputs to the control system:-

- Phase current.
- Rotor position.
- Manual control settings.

and two main outputs:-

- Power converter pulse width modulated (PWM) gate drive signals.
- Analogue monitoring outputs.

Apart from the phase current inputs and PWM outputs, which come directly from the TMS320E14, all other signals are accessed via the VME bus. A digital input board is used for the rotor position encoder which also provides digital protection inputs, and a speed monitoring output. Front panel switches and rotary controls allow parameters such as advance angle and current demand to be adjusted manually. These values are digitised by an eight channel analogue input board which is also used to measure the dc-link voltage. A six channel analogue output board allows simultaneous monitoring of 'invisible' parameters (such as flux-linkage), as well as those which can be measured directly.

Phase current is digitised using 12bit analogue to digital converters with a conversion time of approximately $8\mu\text{s}$. These are accessed directly by the TMS320E14 as a digital input port. This processor is equipped with several counter timers, which in this application are used for PWM generation, and speed measurement by timing of position encoder signals. The PWM frequency was chosen to be 10kHz (see Appendix C for reasoning) and the TMS320E14 can generate three phase PWM signals at this frequency without intervention from the CPU.

There are several hardware protection features; the TMS320E14 has a hardware watchdog which is configured to reset the processor and disable all power devices if it times out. After power on reset, or a reset initiated by the 68030, all power devices are disabled, and the TMS320E14 is held in reset to prevent spurious operation. Phase currents are also monitored and if any exceed a pre-set limit, once again the power devices are disabled.

The Power Converter

The power converter has four main elements; the input rectifier, dc-link filter, dc-link dump, and inverter bridges as shown in Figure 3.3. A simple diode bridge is used to rectify the three phase supply and provide the dc-link voltage. Nominal dc-link voltage is 500V which corresponds to a line to line voltage of 355V RMS. At rated power line currents are 15A RMS. The series inductor limits inrush current and improves line current form factor. In conjunction with the dc-link capacitor the inductor forms the dc-link filter which reduces voltage ripple on the dc-link. The main functions of the dc-link capacitor are to act as an energy storage reservoir, and present a low supply impedance to the inverter bridges. When acting as a generator or braking, energy is transferred to the dc-link from the machine. This energy cannot be returned to the supply and must be dissipated to prevent the dc-link voltage from rising beyond safe levels. This function is performed by the dc-link dump circuit, which connects a fixed resistance across the dc-

link when the voltage reaches a pre-set level. Electrical energy is thus dissipated as heat in the dump resistor.

There are three inverter bridges connected to the dc-link, one for each phase of the switched reluctance machine. Each 'H' bridge inverter is made up of two bridge legs, each comprising of two IGBT switches and anti parallel diodes. By appropriate load connection and control of devices, the machine may be configured for operation with a full 'H' bridge per phase, a two transistor chopper per phase - as commonly used in conventional switched reluctance drives, or as a six transistor three phase bridge, with star or delta load connection.

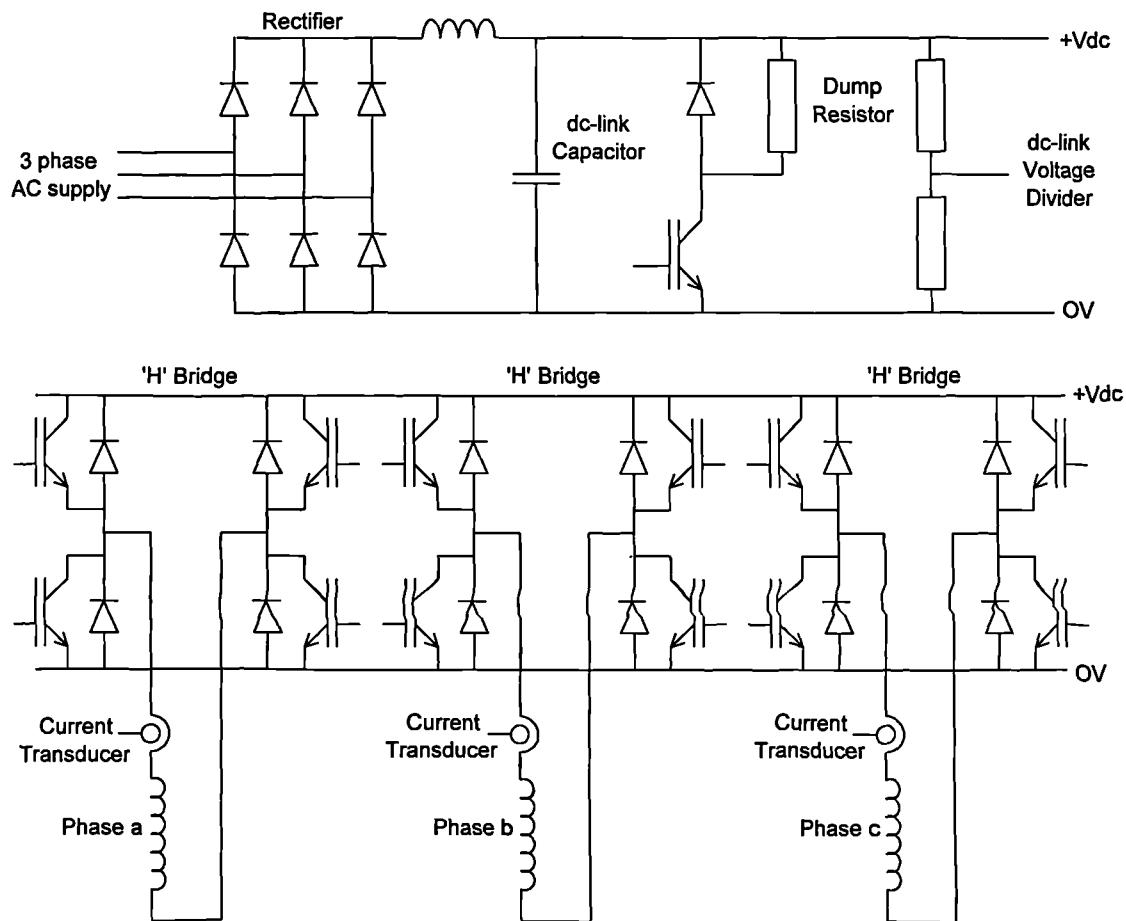


Figure 3.3 Power converter circuit arrangement.

Other circuits associated with the power circuit are a dc-link voltage divider for measurement of the dc-link voltage, Hall effect current transducers for phase current measurement, power switch gate drive circuits and auxiliary power supplies. The phase current transducers used are 50A Hall effect devices with a bandwidth of DC to 100kHz. Two primary turns are used to give full scale output for 25A. Gate drive circuits for the IGBT power switches incorporate over-current protection and have power supply under-voltage lockout. Control signals to the drivers are optically isolated,

with power for the high-side drivers transformer isolated. Auxiliary power supplies provide power for the current measurement circuit, low side IGBT drivers, and high side IGBT drivers through a multi-winding high frequency isolating transformer.

Apart from the commercially produced 68030 and TMS320C31 processor boards all other circuits, including the power converter and TMS320E14 processor board, have been designed and built by the author. Further details of the converter power and control circuits can be found in Appendix C.

3.1.3 Drive Control Software

Although all control boards are connected by a VME bus, as shown in Figure 3.4, not all I/O signals are available on the bus and there is only one VME bus master. This dictates to a great extent how the software must be configured.

Programming Languages

Code for the 68030 and TMS320E14 is written in assembly language, with a software simulator available for the TMS320E14. The TMS320C31 is programmed in 'C' which enables more complex algorithms and data structures to be implemented with ease. Although implemented in 'C' some use of assembly language instructions is required for parts of the initialisation and interrupt routines.

Processor Communication and Computing Power

The Motorola 68030 processor board is the VME bus master and as such it is the only processor on the bus that can initiate data transfers. If other processors wish to transfer data they must signal the bus master, which then controls the data movement. The 68030 is a 32bit CISC processor with good addressing capabilities, but it performs mathematical functions slowly in comparison to a dedicated DSP and is therefore not suitable for implementing three phase digital current control at 10kHz. In addition to the other two processors on the bus, the 68030 has direct access to the rotor position encoder, analogue control inputs and analogue outputs.

With direct access to the phase currents, PWM outputs, and capable of measuring rotor speed, the TMS320E14, a 16bit fixed point DSP, is well suited to the task of high speed digital current control. Clocked at 24MHz it has a cycle time of 166ns, and can perform a multiplication in a single cycle. It does however have a limited address space and addressing capabilities, it is therefore not well matched to programs with large data structures and lookup tables. The third processor is a TMS320C31, 32bit floating point DSP. This has a large address space with extensive addressing capabilities and clocked

at 33MHz performs a floating point multiplication in 60ns. Able to perform complex mathematical functions quickly and capable of handling large lookup tables, the TMS320C31 can implement advanced control algorithms with short cycle times. However, the time available to exploit this processing power is restricted because this processor has no direct access to system variables and must pass all data via the VME bus master.

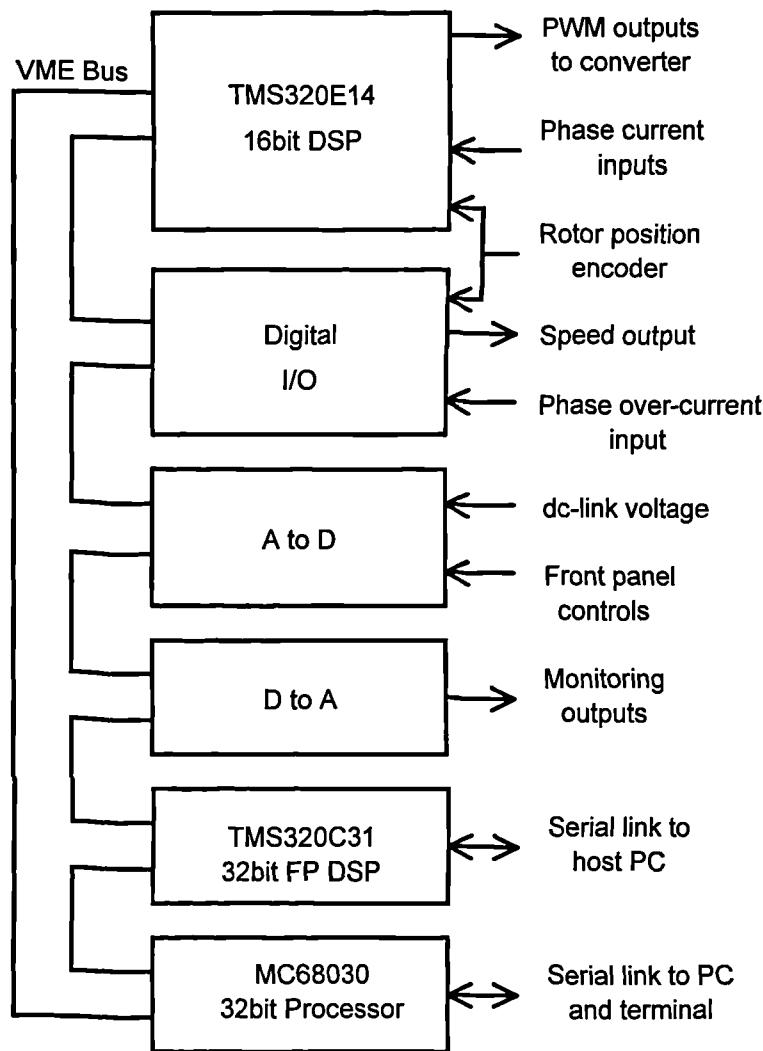


Figure 3.4 Control signals and processor bus arrangement.

Control Tasks

Basic current control of a switched reluctance machine requires two main functions; calculation of the desired phase currents from rotor position and control settings, and control of the real phase currents to the desired values. Figure 3.5 shows in more detail the function blocks required to implement a simple current controlled drive.

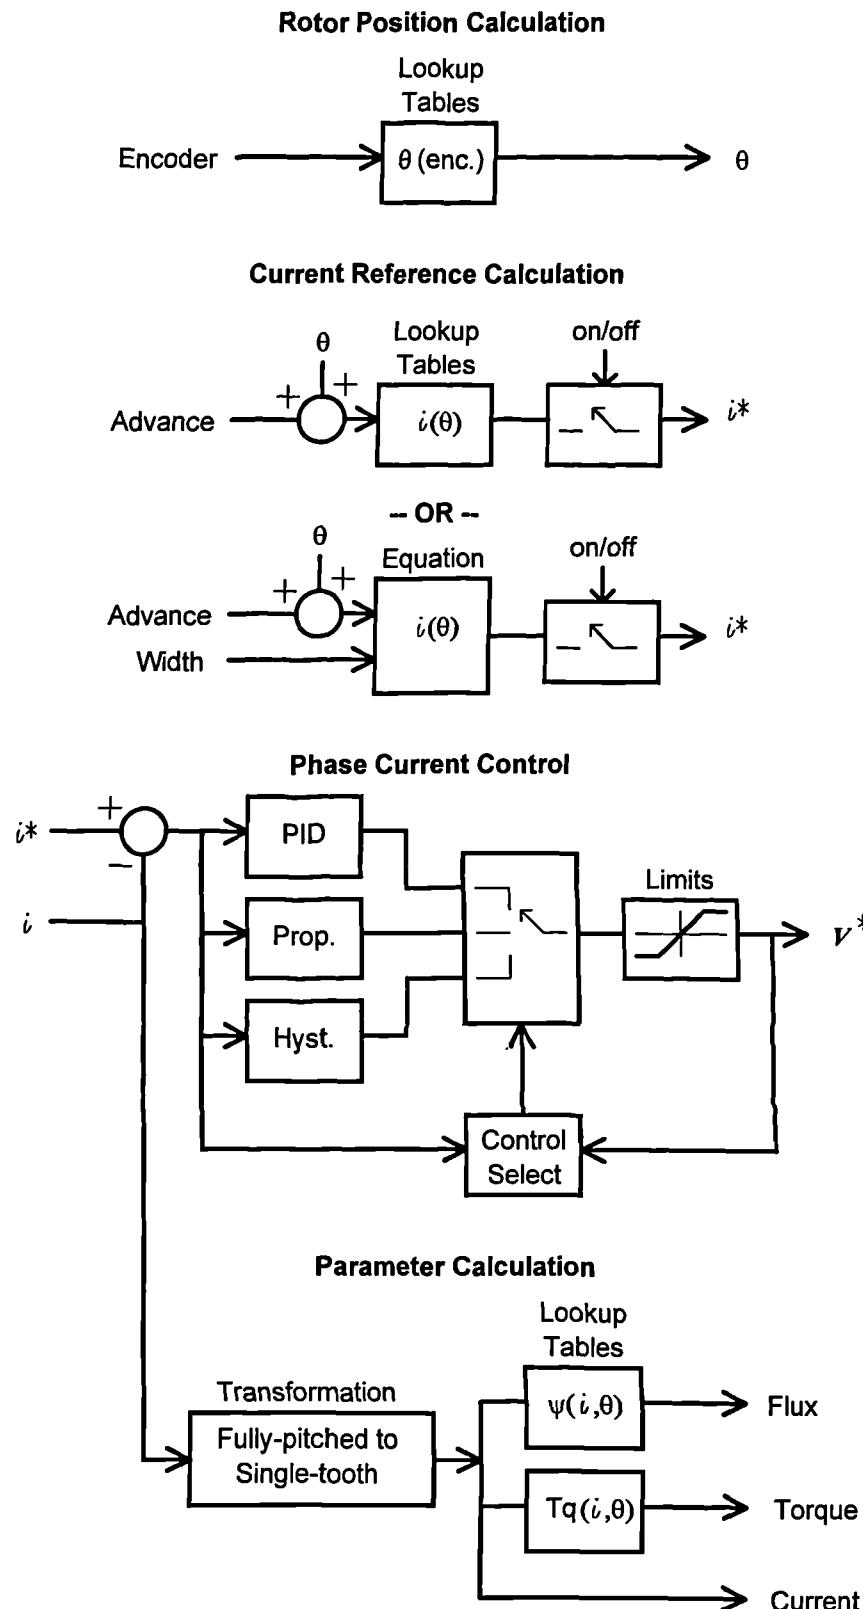


Figure 3.5 Current control software function block diagram.

Rotor position is measured using a gray code output optical encoder. To get a linear relationship between rotor position and measured angle (θ), a lookup table converts from gray code to binary values. Knowing rotor position the phase current references

can then be calculated from equations, or lookup tables. For simple rectangular waveforms such as unipolar two phase on operation, equations are best used as the phase advance angle and conduction width are easily allowed for. With more complex waveforms such as sinusoidal currents, lookup tables are best used as much less calculation is required. Only the phase advance needs to be accounted for as conduction width has no meaning in this case. Having calculated the phase current references, the real current must then be controlled to that value. Figure 3.5 shows a typical implementation using a PID current controller, the gains are adapted to suit the output state and size of the current error. Chapter 4 will examine phase current control in more detail.

Protection Functions

There are three protection functions implemented in software, each designed to protect the power converter by disabling all the inverter power switches. Phase over-current protection is performed by the TMS320E14, which disables the power devices until current falls back below the trip level. The other protection function performed by the TMS320E14 is the software watchdog which must be triggered at regular intervals. Software errors or a system crash cause the watchdog to time out which resets the processor and disables the inverter. The third protection feature is implemented on the 68030, and is designed to protect the power converter if the dc-link dump circuit fails whilst braking or generating. Failure of the dc-link dump will cause the dc-link voltage to rise beyond its nominal value, this is detected and the inverter disabled before the voltage can rise to levels which may damage the power devices.

Processor Configurations

There are two process configurations used for testing the switched reluctance drive; the two control processor configuration, and the three control processor configuration.

The two control processor configuration is used to implement basic current control only. In this configuration the 68030 calculates the phase current references from rotor position and front panel control settings. House keeping functions and monitoring outputs are also handled by the 68030. System timing and phase current control are implemented by the TMS320E14. A third processor, the TMS320C31, is not used as any part of the control process, but can be used to perform monitoring functions such as real-time calculation of flux-linkage and shaft torque.

The three processor configuration is used to implement more complex control schemes, such as flux-linkage control, constant torque operation or speed control. With this

configuration most of the processing tasks are performed by the TMS320C31 including: calculation of rotor position, phase reference value calculation, phase current or flux-linkage control, and real time calculation of flux-linkage and torque. The TMS320E14 controls phase current measurement, generates PWM outputs, performs over-current protection, and determines system timing. Control parameter setting, analogue monitoring outputs, house keeping functions, and data transfer between processors, are all handled by the 68030.

3.2 Summary and Conclusions

The experimental switched reluctance drive hardware has been introduced. Each element of the drive and loading arrangement is designed for maximum flexibility and ease of configuration. This is achieved by using a power circuit, which by control of device signals, can emulate several common configurations, and by implementation of almost all control functions in software.

Chapter 4

Current Control in Switched Reluctance Drives

4.1 Current Control of Switched Reluctance Machines

In most machine drives it is desirable to control phase current, as shaft torque is proportional to phase current or phase current squared.

A switched reluctance machine presents quite a different load to a phase current controller than, for instance, an induction motor or brushless DC machine. Looking first at the load time constant, the doubly-salient and saturating nature of a switched reluctance machine results in a phase time constant which varies considerably with rotor position and phase current, as the magnetic circuit reluctance varies. This is in contrast to induction and brushless DC machines, which tend to have a fairly constant phase inductance, and is quite low in the case of a brushless DC machine (due to the large effective air gap). Secondly, in a brushless DC machine the back emf is a function of speed and rotor position, but in a switched reluctance machine it is also a function of phase current. These factors combine to produce a complex and non-linear load, which makes the optimisation of closed loop current controllers especially difficult.

Current control of a fully-pitched winding switched reluctance machine is further complicated by the strong mutual coupling, which is a non-linear function of rotor position and phase current.

4.1.1 Current Control Methods

Phase current is usually controlled using some form of closed-loop feedback controller, with the overall aim of forcing the actual phase current to track the reference. Switched reluctance drives in general employ a voltage source converter, and use some form of duty cycle modulation to control the phase voltage in response to the phase current error. Several different types of controller have been used; hysteresis, constant off time, and PI or PID being the most popular. Each of these types of controller will now be examined in more detail.

Hysteresis Control

A hysteresis controller could be considered to be the simplest form of current controller, especially when implemented in analogue electronics. This type of controller has been used by Torrey *et al* [4.1], and a typical implementation is shown in Figure 4.1. The controller attempts to maintain the current within a hysteresis band around the desired current, by application of full positive or full negative volts to the phase. For an ideal system with no delays, the controller is unconditionally stable, and there is an inverse relationship between switching frequency and hysteresis band, for a given load time constant and supply voltage. There is no delay in the response to step demand changes. In any real system several factors compromise this ideal, and in many ways desirable, behaviour. The most important limitation is the power switching frequency which must be limited to a value commensurate with the power device used. If using relatively slow devices, such as GTOs, this can lead to high proportions of ripple current. Measurements presented by Ray *et al* [4.2] for example, have a peak to peak ripple of approximately 16A with an average current of 24A. System delays, most relevant in digital implementations, also limit the maximum switching frequency and minimum current ripple which can be achieved. System delays have the effect of delaying the application of a new voltage after the current has passed through a current threshold. This increases the current ripple because during the delay the current continues to ramp further outside the hysteresis band. As the hysteresis band is reduced, system noise starts to become relevant as noise spikes can cause spurious switching, leading to excessive switching losses, increased current ripple and more radiated noise. Despite these practical limitations, the simplicity and unconditional stability of hysteresis controllers make them attractive for switched reluctance drives.

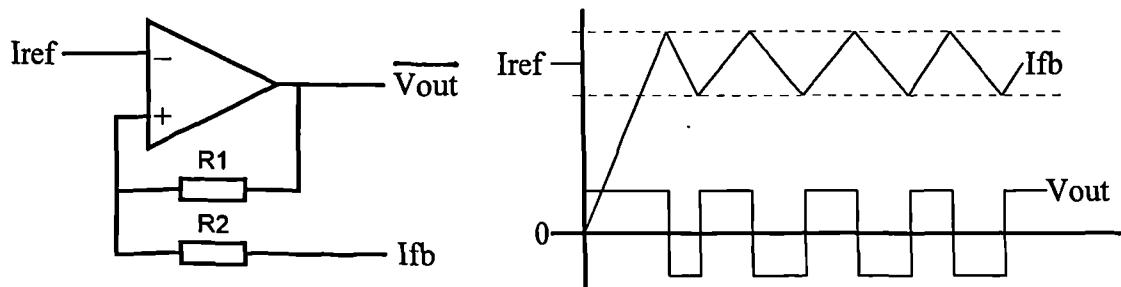


Figure 4.1 Simplified hysteresis current controller.

Constant Off-Time Control

Constant off time control, as used by Sugden *et al* [4.3], is in some respects similar to hysteresis control, but with only an upper current limit. A constant off time controller, as illustrated in Figure 4.2, applies positive phase volts until the current reaches the reference level. At this point the power devices are switched off for a fixed duration and the current falls, after the fixed delay positive phase volts are re-applied. The resulting phase current has a peak value determined by the reference value, with the ripple set by the off time and system time constant. Maximum switching frequency is limited by the off time, which can also delay the response to step demand changes. As with a hysteresis controller a constant off time controller is unconditionally stable, but system delays and noise can compromise the ideal performance.

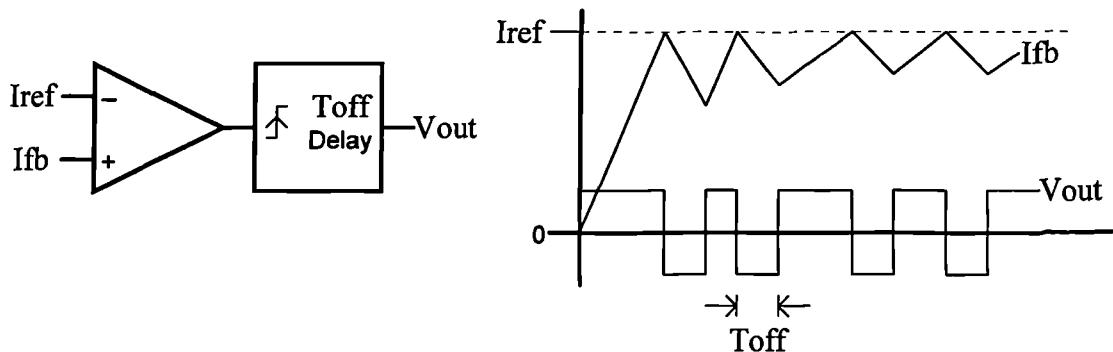


Figure 4.2 Simplified constant off time current controller.

A common feature of hysteresis and constant off time controllers is that the switching frequency varies with load and control parameters. For analogue implementations the main effect of this is to complicate switching loss calculations, but for digital implementations the implications are more severe. The variable switching frequency complicates processor timing, and the asynchronous switching times of different phases can require the processor to action two tasks at the same time. A further complication is that the phase current can only be quantised at a finite sample rate. To be suitable for microprocessor implementation the control algorithm should be implemented with a fixed sample interval, usually synchronous to the sampling of phase current, and each phase switched synchronously. These timing restrictions compromise the operation of digital hysteresis and constant off time controllers, in comparison to their analogue counterparts. Digital control does however make it much easier to implement non-linear and adaptive control systems, and in a highly integrated solution can reduce the component count and overall cost.

PID Control

Fixed frequency pulse width modulation is usually used in conjunction with PI (Proportional plus Integral), or PID (Proportional plus Integral plus Derivative) controllers. In analogue implementations this is commonly achieved by comparing the error amplifier output with a triangle wave reference, and in digital implementations the error amplifier output is used to set the PWM timer pulse width. Figure 4.3 shows the general form of a PID current controller; the phase current error is amplified by the PID gain block (in practical systems this is then limited to a value corresponding to 100% modulation) before being fed to the modulator. Analogue modulators have infinite output voltage resolution, whilst digital modulators have a finite number of possible output voltages set by the timer word length, and usually in the range 8 to 12 bits. The large number of possible output levels (c.f. two levels in hysteresis control) enables the output voltage to be closely matched to the level required to keep the current at the desired level, thus keeping current ripple to a minimum. System stability is dependant on correct selection of the three gain terms in a PID controller and usually involves a trade off of bandwidth for stability. Response to step changes in demand can be slowed by the integral term, although this can be compensated for to some extent by the derivative term. However system stability and noise limit the maximum derivative gain. The design of PID current controllers is covered in more detail later in this chapter.

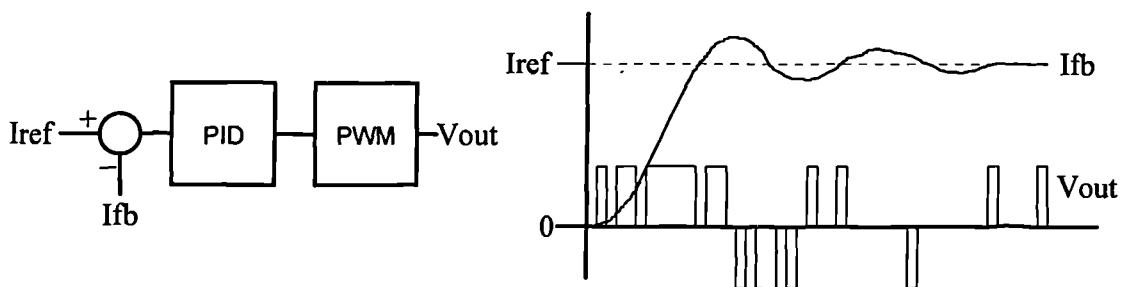
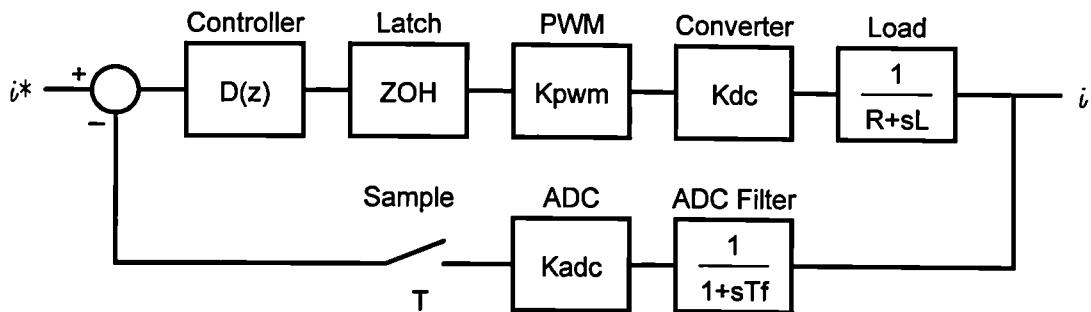


Figure 4.3 Simplified PID current controller.

4.1.2 Current Controller System Model

Figure 4.4 shows the simplified current control system model for one phase. Phase current error is first amplified by the controller block, followed by the pulse width modulator and power converter, before being applied to the load. This is represented as a simple inductive circuit with constant values; in reality there is a back emf and the inductance varies considerably with operating point. The current measurement feedback loop incorporates filtering and sampling, in addition to gain.



K_{pwm} = Pulse width modulator gain

K_{dc} = Power converter gain

R, L = Load machine phase resistance and inductance

T_f = ADC input filter time constant

K_{adc} = Current measurement ADC gain

Figure 4.4 Closed loop current control system block diagram.

The linear model shown in Figure 4.4 is only valid for small signal characteristics because it does not model the non-linear effects of magnetic saturation or controller output limits. Controller signal processing delay and dead time effects are also not shown in this simplified model. There are two main sources of output disturbance not shown in this model; the motional emf, and for fully-pitched winding machines the mutual coupling between phases. Both these effects vary non-linearly with rotor position and phase current.

The multitude of non-linearities in the converter, load, and output disturbance effects, make it impossible to produce a linear model from which to assess bandwidth and stability over a meaningful range of conditions. Simulation incorporating load non-linearity, mutual coupling and motional effects, is required to predict the current step response and detect instability. This type of modelling is outside the scope of this thesis, the simplified model does however provide a starting point from which to optimise controller parameters.

4.1.3 PID Current Controllers

When implemented in analogue electronics, PID controllers are both more complex and more difficult to optimise, than hysteresis controllers. But with a microprocessor implementation they have several advantages. Fixed frequency operation simplifies processor timing, implementation of the controller in software allows control parameters to be dynamically changed, and integrated pulse width modulators are available which produce low jitter outputs without processor intervention. Probably the biggest problem with PID controllers is obtaining a fast, stable response over a wide range of load conditions. For a switched reluctance machine the large variation in load time constant,

motional emf, and in the case of a fully-pitched winding machine the strong mutual coupling between phases, makes controller optimisation especially difficult.

The continuous time transfer function of a PID controller can be expressed as follows:-

$$D(s) = K_p \cdot \left[1 + \frac{1}{sT_i} + sT_d \right] \quad (4.1)$$

The subscripts 'p', 'i' and 'd' refer to the proportional, integral and derivative parts respectively. When transformed to the Z domain using the backward difference method, and with a sample interval of T , the resulting transfer function is:-

$$D(z) = K_p \cdot \frac{a_2 z^2 + a_1 z + a_0}{z(z-1)} \quad (4.2)$$

The three quadratic gain terms are:-

$$a_2 = 1 + \frac{T}{T_i} + \frac{T_d}{T} \quad (4.3)$$

$$a_1 = - \left[1 + \frac{2T_d}{T} \right] \quad (4.4)$$

$$a_0 = \frac{T_d}{T} \quad (4.5)$$

The resulting time domain difference equation used for implementation is:-

$$y_n = y_{n-1} + K_p \cdot (a_2 e_n + a_1 e_{n-1} + a_0 e_{n-2}) \quad (4.6)$$

The coding required to implement this equation can be found in Appendix D.

Although relatively simple to implement, a controller using only the difference equation above, will suffer from 'integral windup' under large signal conditions. The main cause of integral windup is output saturation and non-zero current error. Under these conditions the integral term in the controller continues to increase but the actual output is unable to do so. Once the output has reached the set point, the integral term, and hence the output are both large. The current error must then reverse in order to reduce the integral term, this leads to a large overshoot as the set point is reached, and a longer settling time. This effect is illustrated in Figure 4.5 which shows the phase current step response of a simple PID controller without anti-windup features.

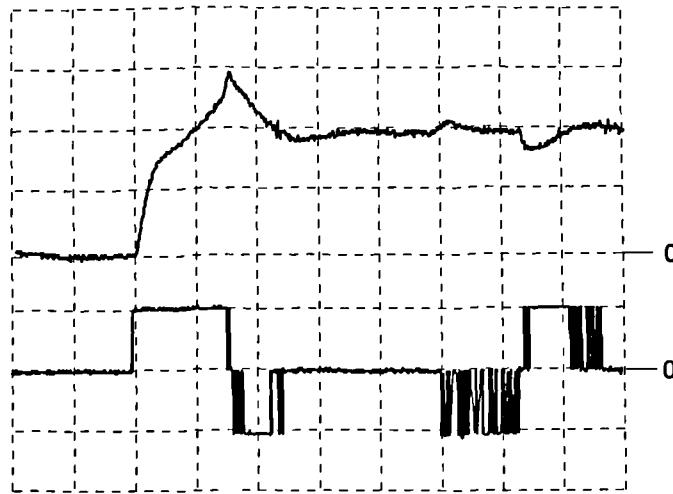


Figure 4.5 PID controller response without anti-windup. Trace 1 - Phase current 5A/div. Trace 2 - Phase voltage 400V/div. Horizontal 5ms/div.

To prevent integral windup the integral gain term is removed when the output saturates, thus preventing the integral term ramping up to a large value whilst the output is saturated. Incorporating this feature has a big impact on the large signal step response as shown in Figure 4.6. The conditions used to produce this figure are as for Figure 4.5, except that the integral gain term is removed whilst the output voltage is saturated.

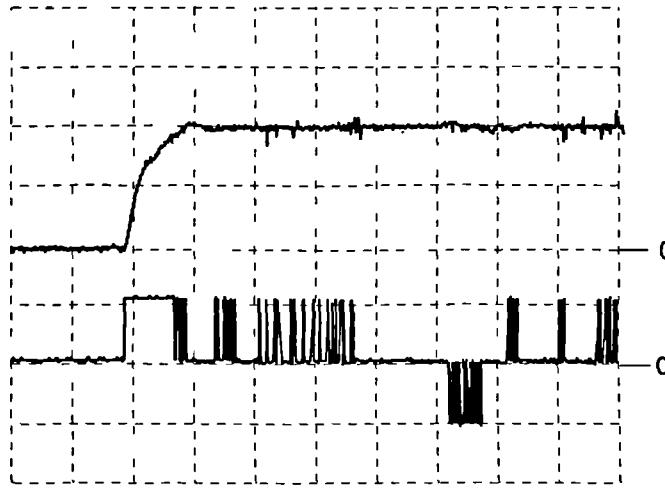


Figure 4.6 PID controller response with anti-windup. Trace 1 - Phase current 5A/div. Trace 2 - Phase voltage 400V/div. Horizontal 5ms/div.

Response to output disturbances caused by commutation is important for the strongly coupled phases of a fully-pitched winding switched reluctance machine. The maximum proportional and integral gain is limited by system stability, and this can lead to a slow response to output disturbances. This can be alleviated to some extent, by increasing the proportional gain if the current error exceeds a pre-set value. Figure 4.7 shows the measured current response for a current demand of 10A peak at 750rpm. The overshoot

evident in the middle of the main conduction interval is due to commutation of another phase. By incorporating the adaptive gain change current overshoot is reduced from 2.5A in trace 1, to 1A in trace 2.

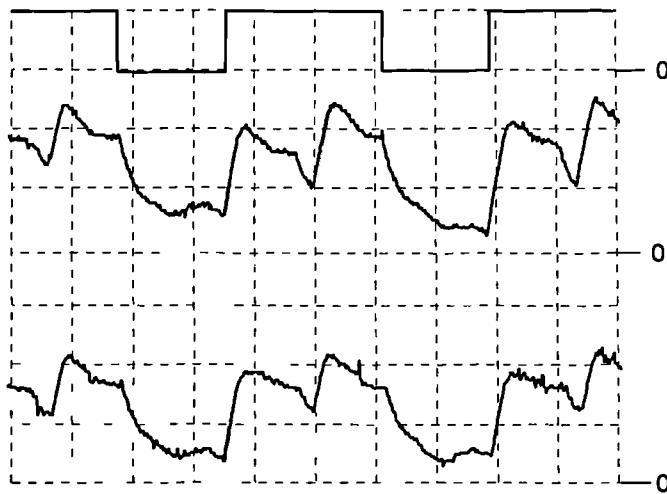


Figure 4.7 PID controller response with adaptive gain setting. Trace 1 - Current reference, 10A/div. Trace 2 - Non-adaptive controller phase current, 5A/div. Trace 3 - Adaptive controller phase current, 5A/div. Horizontal 2.5ms/div.

The PID phase current controller used for experimental drive operation incorporates the anti-windup and adaptive gain changing features *described in this section*. A further enhancement is to run the current control software at twice the PWM output frequency. Although it is not possible to change the output pulse width at a higher sample rate, controller performance is improved because the higher current sample rate and reduced delay both improve controller bandwidth and stability. Further details of current controller implementation can be found in Appendix D.

4.1.4 Transformed Current Controllers

Although designed and implemented separately, the phase current controllers for a fully-pitched winding switched reluctance machine are in fact coupled. The highly non-linear and position dependant coupling between phases would make it very complex to design an optimised, multi-input multi-output state space controller. For this reason the controllers are implemented separately, but are designed to be tolerant of the large output disturbances coupled in from the other phases. If the equivalent single-tooth winding current is controlled, rather than the actual fully-pitched winding currents, then de-coupled control can be achieved. Transformation matrices given in Chapter 2 section 2.2 can be used to calculate the equivalent single-tooth winding currents and voltages, from the actual fully-pitched winding values. The basis of a transformed current controller is to use a PID controller to calculate the equivalent single-tooth winding

voltage from the equivalent single-tooth winding current error, with transformation to and from the actual fully-pitched winding values. This configuration is illustrated in Figure 4.8.

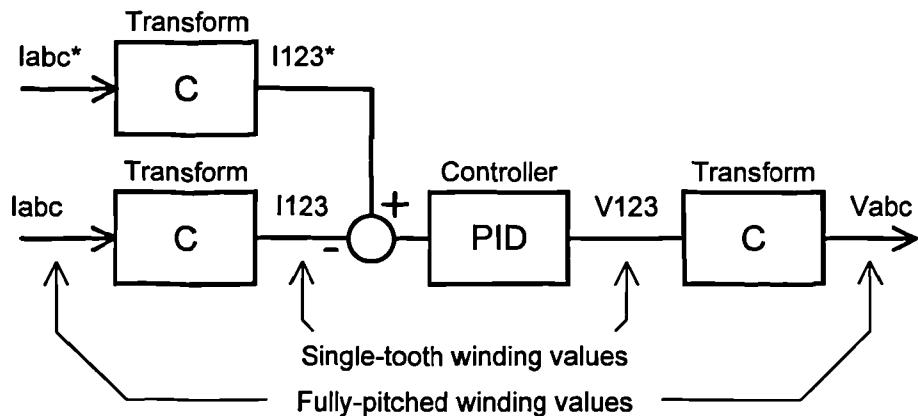


Figure 4.8 Transformed current controller block diagram.

Although simple in theory, in practice there are limits to the degree of de-coupling that can be achieved since it is actually the fully-pitched winding voltages that are applied, and not the equivalent single-tooth winding voltages. If none of the actual output voltages are saturated, then there is full de-coupling between phases. However, if any of the actual phase voltages are saturated then two or more single-tooth winding phase voltages will not equal the desired value. This makes the controller non-linear, and has implications for the integral anti-windup features. The problem is that although the fully-pitched winding voltages saturate at a fixed level (100% duty cycle), the equivalent single-tooth winding voltages are limited to a variable level which is dependant on the actual output voltages. This makes it difficult to detect when a single-tooth winding voltage is being limited, the condition that determines when anti-windup measures should be applied.

Despite these difficulties good controller response can be achieved, as demonstrated in Figure 4.9, which compares the response of a transformed current controller and a standard current controller. In each case the current controller gains and operating conditions are the same. It is evident that the transformed current controller gives a better transient response, in both the actual fully-pitched winding and equivalent single-tooth winding currents. Although performance of the transformed current controller is good under the conditions illustrated, under other operating conditions the transformed current controller is very much worse than the real current controller. This is because of problems with the application of integral anti-windup, as described previously.

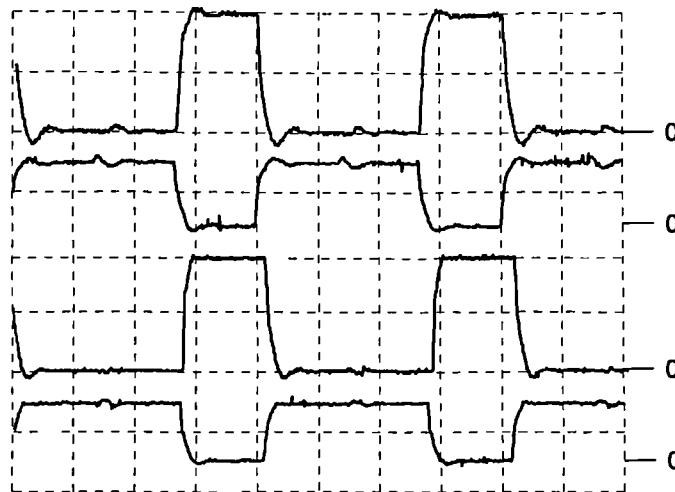


Figure 4.9 Comparison of real and transformed current controllers. Traces 1 and 2 - Real phase current controller, equivalent single-tooth winding current and real phase current. Traces 3 and 4 - Transformed phase current controller, equivalent single-tooth winding current and real phase current. All traces 10A/div vertical, 20ms/div horizontal.

4.2 Summary and Conclusions

Several popular current control methods used for switched reluctance drives have been reviewed. Each type of controller has its advantages and disadvantages. For a simple analogue implementation, hysteresis or constant off time control gives a simple unconditionally stable implementation. A PI or PID controller and PWM output has the advantage of lower current ripple and is better suited to digital implementation, where the fixed frequency simplifies processor timing, and the ability to dynamically alter gains can compensate for load and plant saturation. The controller used for the experimental drive has been introduced, and a simplified plant model used to make initial estimates for controller gain terms. The drive current controller has several features to improve transient performance under large signal conditions, and the effectiveness of these features has been demonstrated.

Chapter 5

Unipolar Operation

5.1 Theory of Operation

5.1.1 The Method of Torque Production

Unipolar operation is the simplest operating mode for a fully-pitched winding switched reluctance drive, and at low speed is very similar to the operation of a conventional single-tooth winding switched reluctance drive. For a three phase fully-pitched winding switched reluctance machine, assuming phase self inductances are constant and the magnetic circuit is linear, then (as previously stated in Chapter 2 section 2.1) the torque equation can be expressed as:-

$$T = i_a \cdot i_b \cdot \frac{dM_{ab}}{d\theta} + i_b \cdot i_c \cdot \frac{dM_{bc}}{d\theta} + i_c \cdot i_a \cdot \frac{dM_{ca}}{d\theta} \quad (5.1)$$

As currents are unipolar the product of any two phase currents must be positive or zero. Thus positive torque can only be produced from periods of *rising* mutual inductance with rotor position.

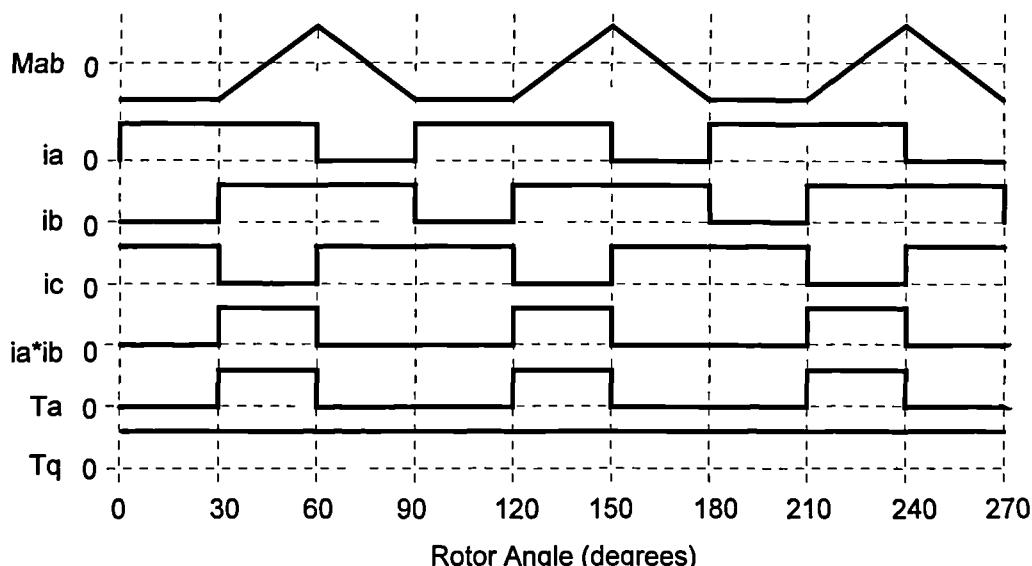


Figure 5.1 Derivation of torque from fully-pitched winding parameters.

Figure 5.1 shows the derivation of torque, based on idealised mutual inductance characteristics (as derived in Chapter 2, section 2.1) and unipolar phase currents. Each phase carries a constant current for 240° and no current for the remaining 120° , with phases displaced by 120° with respect to each other. For each pair of phases the product of phase currents is zero, except for a period of 120° , which for maximum positive torque is synchronised to the period of rising mutual inductance. Each pair of phases produces a torque contribution lasting 120° , which when added to the contributions of other phase pairs, gives constant shaft torque.

Although this derivation of torque is relatively straight forward, it cannot be applied in practice. This is because a real machine is highly non-linear, so assumptions which allow the above derivation are no longer valid. By transforming fully-pitched winding parameters to their equivalent single-tooth winding parameters, analysis can be carried out using easily obtained machine characteristics. Figure 5.2 shows the relationship between the real fully-pitched winding phase currents, and their equivalent single-tooth winding phase currents.

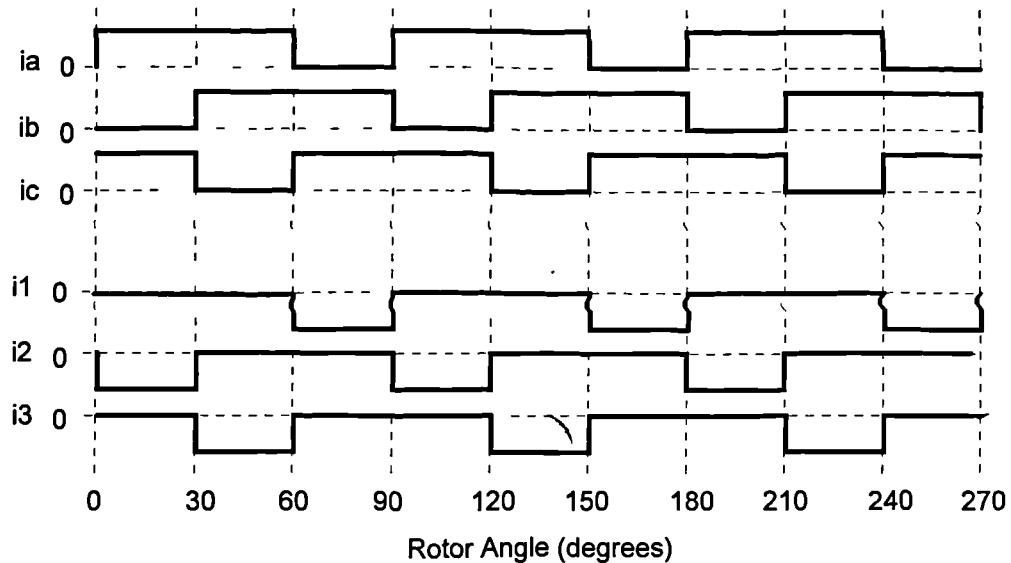


Figure 5.2 Fully-pitched winding currents and equivalent single-tooth winding currents.

The fully-pitched winding and single-tooth winding currents are related by the equation:-

$$\mathbf{i}_{123} = \mathbf{C} \cdot \mathbf{i}_{abc} \quad (5.2)$$

Where the matrix \mathbf{C} is as given in Chapter 2, section 2.2. Note that for positive fully-pitched winding currents, the equivalent single-tooth winding currents are negative. However, the equivalent single-tooth winding parameters are often shown inverted (i.e. positive) in this thesis because interpreting waveforms is much easier, and the actual

direction of phase currents has no effect on the torque produced, for a single-tooth winding switched reluctance machine. From Figure 5.2 it can be seen that for fully-pitched winding phase currents with 240° conduction intervals, the equivalent single-tooth winding phase currents have conduction intervals of 120° .

For an idealised three phase single-tooth winding switched reluctance drive, 120° conduction could be considered normal operation. Torque derivation for this mode of operation is shown in Figure 5.3. Each phase is synchronised so that it is conducting during the period of rising self inductance, this gives a positive torque contribution T_1 , the other torque contributions adding to give constant torque.

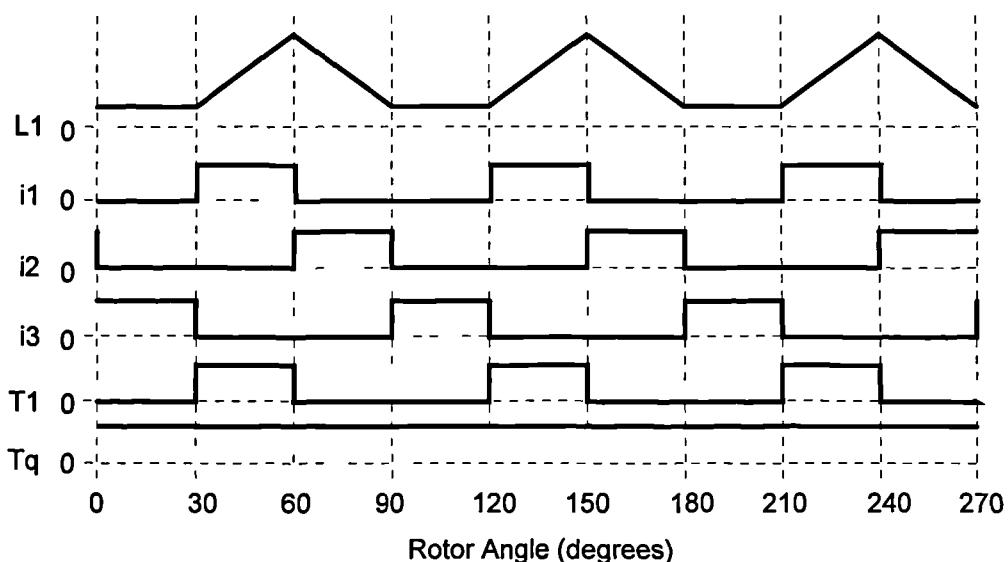


Figure 5.3 Derivation of torque from single-tooth winding parameters.

5.1.2 Phase Current Reference

As with any switched reluctance drive, the torque produced is dependent on the correct synchronisation of phase current to rotor position. This is controlled by altering the phase current demand 'on' and 'off' angles for each phase. With reference to Figure 5.4, it can be seen that the phase current reference i_{ref} is set to the constant demand value I_{dem} between the 'on' and 'off' angles, and to zero at other times. The control angles are measured in electrical degrees, and are given with respect to the aligned position. This is defined as being the position where a reference set of stator and rotor teeth are in alignment. In the case of unipolar excitation, one electrical cycle corresponds to a rotation of one rotor tooth pitch, which for a 12:8 machine is 45° mechanical. The phase 'conduction' angle is the difference between the 'on' and 'off' angles and is nominally 240° for unipolar operation. The current references for all three phases have the same form, but they are displaced from each other by $1/3$ of an electrical cycle.

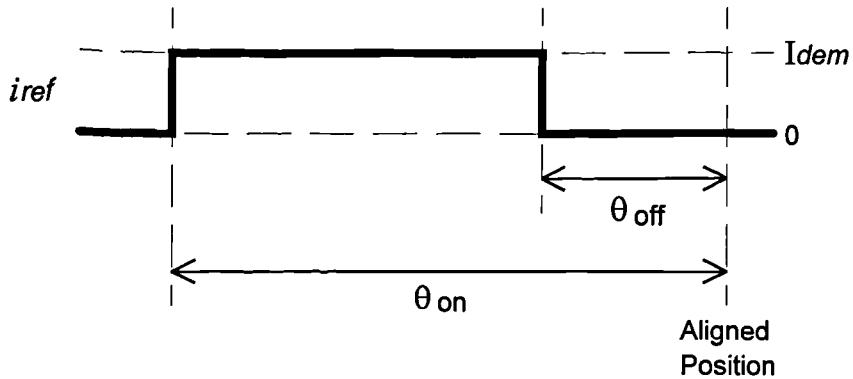


Figure 5.4 Phase current reference waveform showing control angles and current demand level.

5.2 Simulation Studies

To simulate the operation of the experimental fully-pitched winding switched reluctance machine drive and the equivalent single-tooth winding switched reluctance machine drive, simulation programs have been written in the High Level Language 'C', and run in a DOS environment.

These programs use measured flux-linkage and static torque characteristics, to equate current, flux-linkage and torque. Extensive use is made of transformations from fully-pitched to single-tooth winding parameters and vice-versa. This allows for instance, fully-pitched winding phase current to be related to flux. This is achieved by transformation of current to single-tooth winding values, looking up flux-linkage using the equivalent single-tooth winding characteristics, and then transforming the flux-linkage back to fully-pitched winding values. The transformations make use of the matrix C and its inverse, as given in Chapter 2 section 2.2, which is for fully-pitched winding and single-tooth winding machines having the same total number of turns per phase. Unity turns ratio results in an equivalent single-tooth winding machine, which requires twice the current per phase of the fully-pitched winding machine, for the same total MMF. To keep the power converter peak VA the same in each case, the dc-link voltage is halved in the single-tooth winding switched reluctance drive simulation.

The simulation does not take account of the resistive voltage drop across the windings (except for calculation of losses) and assumes either ideal current control or full voltage control.

Machine cross sections thus far have illustrated a 6:4 machine as this is simpler than the 12:8 machine used. Figure 5.4.1 (see inside back cover) is a cross section of the experimental machine showing winding positions and the angle measurement datum, as well as idealised mutual inductance and current profiles. For the control angles given in section 5.2.1 of 260 and 20 electrical degrees, the corresponding mechanical angles (θ) for phase 'b' excitation are $12.5^\circ \pm (n \times 45^\circ)$ and $42.5^\circ \pm (n \times 45^\circ)$ respectively.

5.2.1 Current Control at Low Speed

A closed loop current controller can be modelled as an ideal current source at low speed because the current rise and fall times are small compared to the period. It is thus valid to simulate drive operation at low speed using ideal current sources. In this case the actual speed is of no consideration. Results in this section are for fully-pitched winding phase currents of 10A peak, 240° conduction interval, 260° 'on' and 20° 'off' angles. This gives an equivalent single-tooth winding phase current of 20A peak, and an average torque of 63Nm.

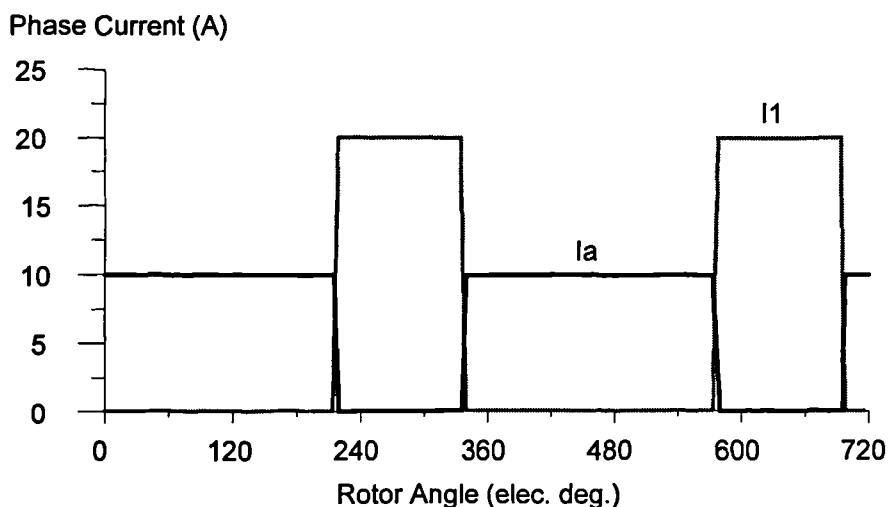


Figure 5.5 Real fully-pitched winding phase current - I_a , and equivalent single-tooth winding phase current - I_1 .

Fully-pitched winding and equivalent single-tooth winding phase currents are shown in Figure 5.5. As an ideal current source is used these waveforms follow the theoretical ideal. The corresponding flux-linkages are shown in Figure 5.6, as would be expected flux-linkage rises in the single-tooth winding as the aligned position is approached. Fully-pitched winding flux-linkage is calculated from the equation:-

$$\Psi_{abc} = \mathbf{C} \cdot \Psi_{123} \quad (5.3)$$

It can be seen that during each third of a cycle, the fully-pitched winding flux-linkage follows the flux-linkage profile in one of the single-tooth winding phases. Comparing Figure 5.5 and Figure 5.6, it is clear that half way through the current conduction interval there is a step change in flux-linkage. In each half of the conduction interval flux-linkage rises steadily with position. During each of these periods electrical energy is being supplied to the machine to produce mechanical output, and correspond to periods of rising mutual inductance between successive pairs of phases. Whilst the fully-pitched winding phase current is zero, the flux linkage is in fact negative and

increasing in magnitude. In terms of fully-pitched winding parameters, this is due to the negative mutual inductance between the non-conducting phase and other conducting phases. Note that the change in flux-linkage at either end of the conduction interval is the same, and is equal to the sum of the aligned and unaligned equivalent single-tooth winding flux-linkages.

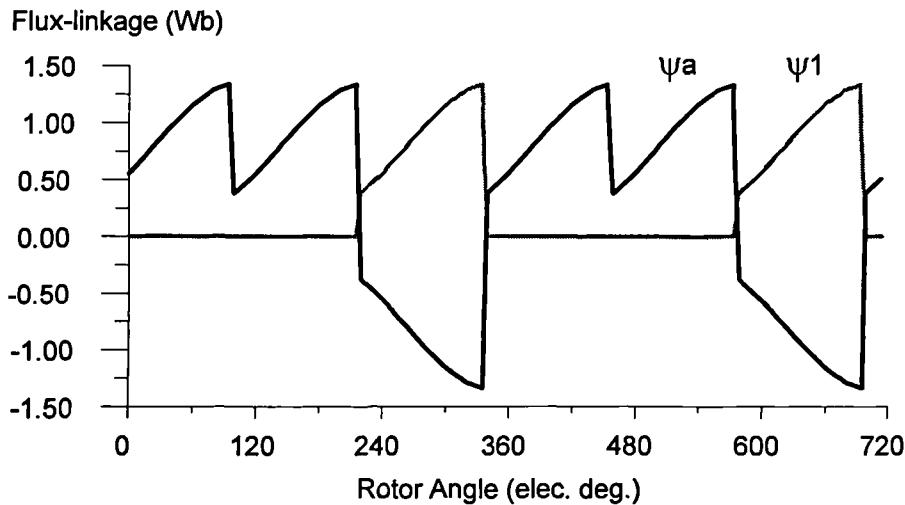


Figure 5.6 Real fully-pitched winding flux-linkage - ψ_a , and equivalent single-tooth winding flux-linkage - ψ_1 .

Torque produced by a single phase can be found from the equivalent single-tooth winding current and machine characteristics, *addition of the three torque contributions* gives the total shaft torque as shown in Figure 5.7. The large peak to peak torque ripple is due to the machine torque/position characteristics when saturated.

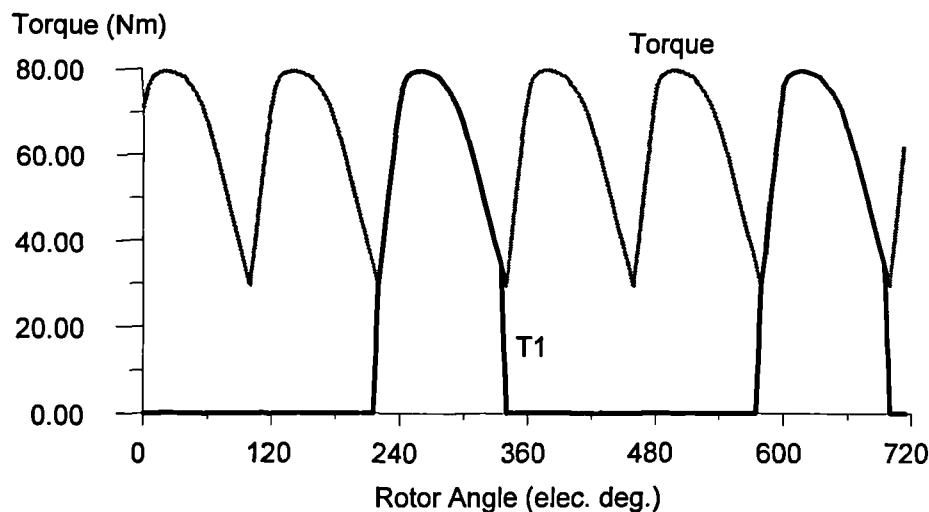


Figure 5.7 Total shaft torque -Torque, and torque contribution of equivalent single-tooth winding phase 1 - T_1 .

Looking at the equivalent single-tooth winding current, flux-linkage and torque waveforms, it is clear that these correspond very closely to what may be expected for a conventional single-tooth winding switched reluctance drive. Indeed the slot MMFs and stator tooth fluxes are the same in the experimental fully-pitched winding switched reluctance machine with 240° conduction, as in the equivalent single-tooth winding switched reluctance machine with 120° conduction.

5.2.2 Voltage Control at High Speed

Voltage control is simulated by application of the full dc-link voltage to a phase for defined periods. This mode of operation corresponds to operation at high speed, using a single pulse voltage mode controller (Sugden *et al* [5.1]), or a closed loop current controller which never reaches the desired current, and hence the output is always saturated. As in the experimental drive resistive voltage drop is small compared to the applied voltage, its effect is neglected in simulation, and the flux-linkage is simply assumed to be equal to the applied volt-time integral.

For a fully-pitched winding switched reluctance drive operating under voltage control, there are significant differences between the equivalent single-tooth winding voltages and those typical of a conventional single-tooth winding switched reluctance drive. This is in contrast to current mode operation, where equivalent currents are the same. These differences are explored in the following figures which compare the voltage mode operation of a fully-pitched winding switched reluctance drive, to that of a single-tooth winding switched reluctance drive. Comparisons are made using single-tooth winding parameters for both drives as these are easily understood.

Simulated results are for operation at 1500rpm, producing an average shaft torque of approximately 20Nm. For the fully-pitched winding switched reluctance drive the dc-link voltage is 500V with the 'on' angle 265° , corresponding figures for the single-tooth winding switched reluctance drive are 250V and 240° .

Considering first the applied voltages; Figure 5.8 shows the real and equivalent single-tooth winding phase voltage for a fully-pitched winding switched reluctance drive, and Figure 5.9 shows the real and equivalent fully-pitched winding voltage for a single-tooth winding switched reluctance drive.

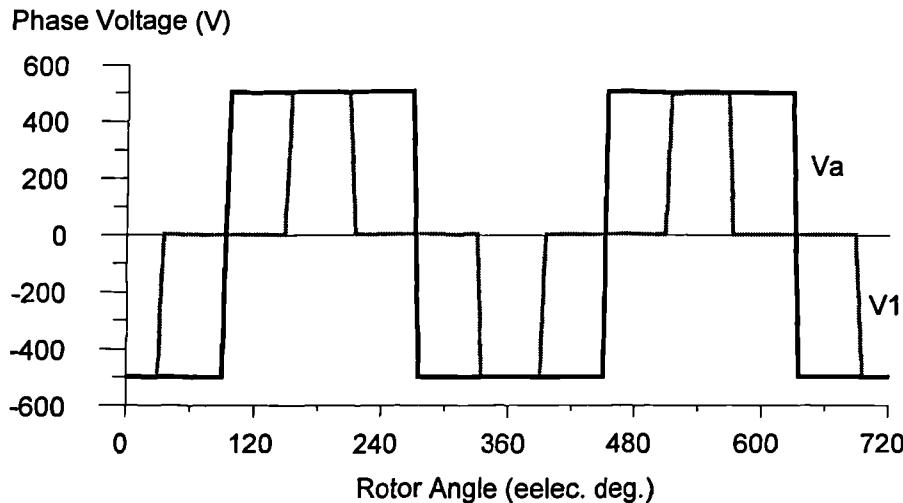


Figure 5.8 Voltage control of the fully-pitched winding switched reluctance machine. Phase voltage - V_a , and equivalent single-tooth winding phase voltage - V_1 .

In each case the applied voltage is a squarewave, with a peak amplitude equal to the dc-link voltage. Examination of the equivalent voltage waveforms shows that in both cases the amplitude is greater than the dc-link voltage, and so the equivalent voltage could not be produced by the corresponding drive.

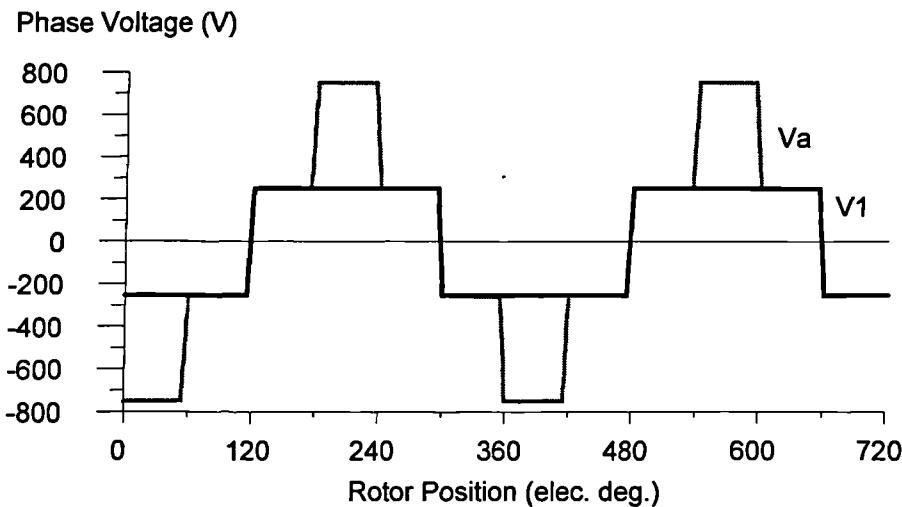


Figure 5.9 Voltage control of the single-tooth winding switched reluctance machine. Phase voltage V_1 , and equivalent fully-pitched winding phase voltage V_a .

The different single-tooth winding phase voltages produce significantly different flux-linkage waveforms, as shown in Figure 5.10. Subscripts 'fpw' and 'stw' refer to waveforms produced by a fully-pitched winding switched reluctance drive, and a single-tooth winding switched reluctance drive respectively. The single-tooth winding switched reluctance drive produces a triangular flux waveform, whilst the fully-pitched winding switched reluctance drive gives an equivalent single-tooth winding flux-linkage that has a 'ramp and pedestal' waveshape.

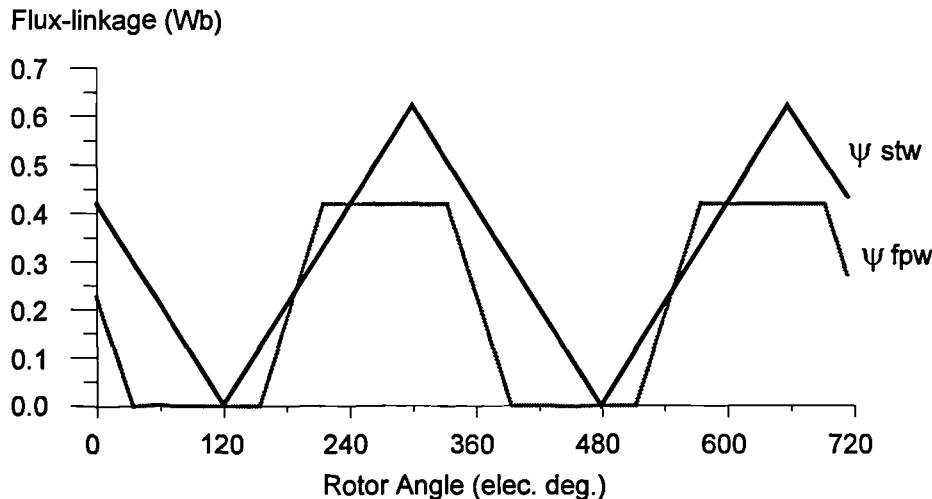


Figure 5.10 Real and equivalent single-tooth winding flux-linkage. Voltage control of the single-tooth winding machine - ψ_{stw} , and voltage control of the fully-pitched winding machine - ψ_{fpw} .

Although the peak flux-linkage is higher for the single-tooth winding switched reluctance drive than the fully-pitched winding switched reluctance drive, the peak current is actually lower, as illustrated in Figure 5.11, and is due to differences in the phasing of flux-linkage relative to rotor position.

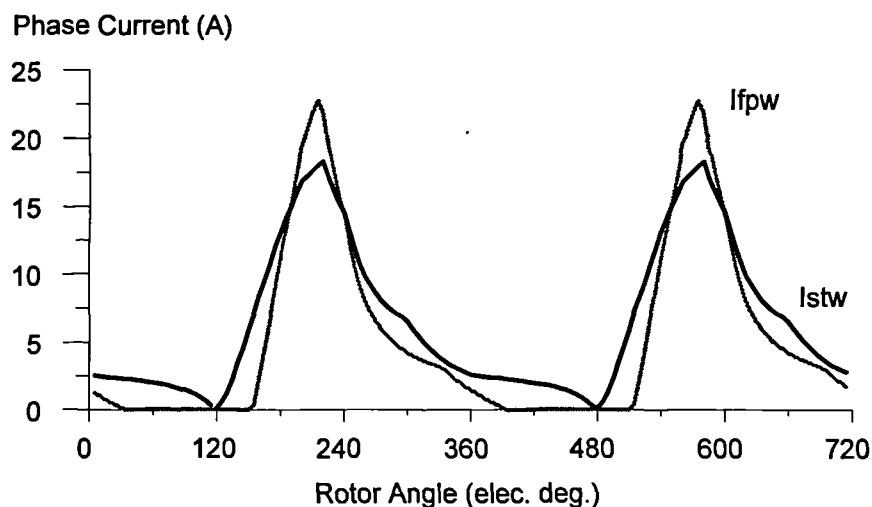


Figure 5.11 Real and equivalent single-tooth winding phase current. Voltage control of the single-tooth winding machine - I_{stw} , and voltage control of the fully-pitched winding machine - I_{fpw} .

As the equivalent single-tooth winding flux-linkage, ψ_{fpw} , is discontinuous, so too is the phase current I_{fpw} , as shown in Figure 5.11. This results in a lower minimum torque as is evident in Figure 5.12. More phase overlap in the single-tooth winding switched reluctance drive flux-linkage, ψ_{stw} , helps reduce torque ripple.

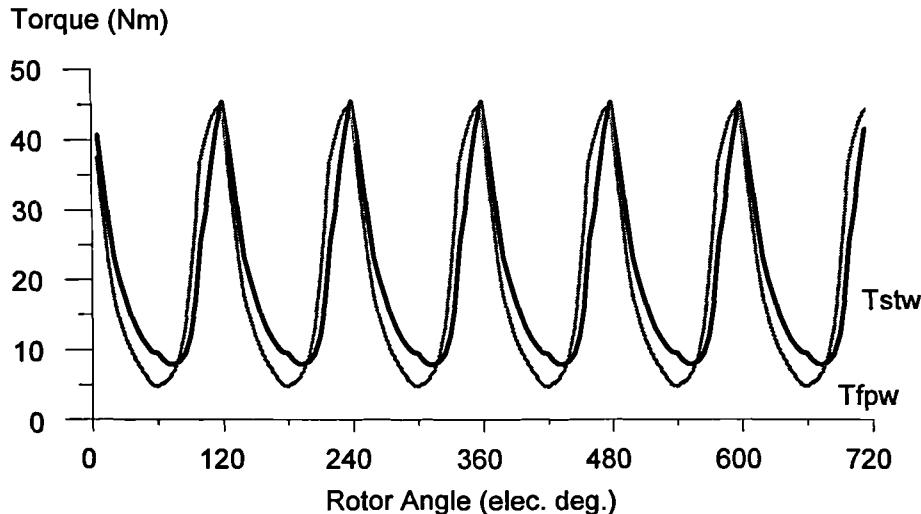


Figure 5.12 Total shaft torque. Voltage control of the single-tooth winding machine - T_{stw} , and voltage control of the fully-pitched winding machine - T_{fpw} .

Examination of the single-tooth winding flux/current loci, helps bring out the differences in operation between the fully-pitched winding and single-tooth winding switched reluctance drives. The period of constant flux-linkage is evident in Figure 5.13, which shows the fully-pitched winding switched reluctance drive locus. It is also apparent that for a given peak current and peak flux-linkage, the locus encloses almost the maximum theoretical area.

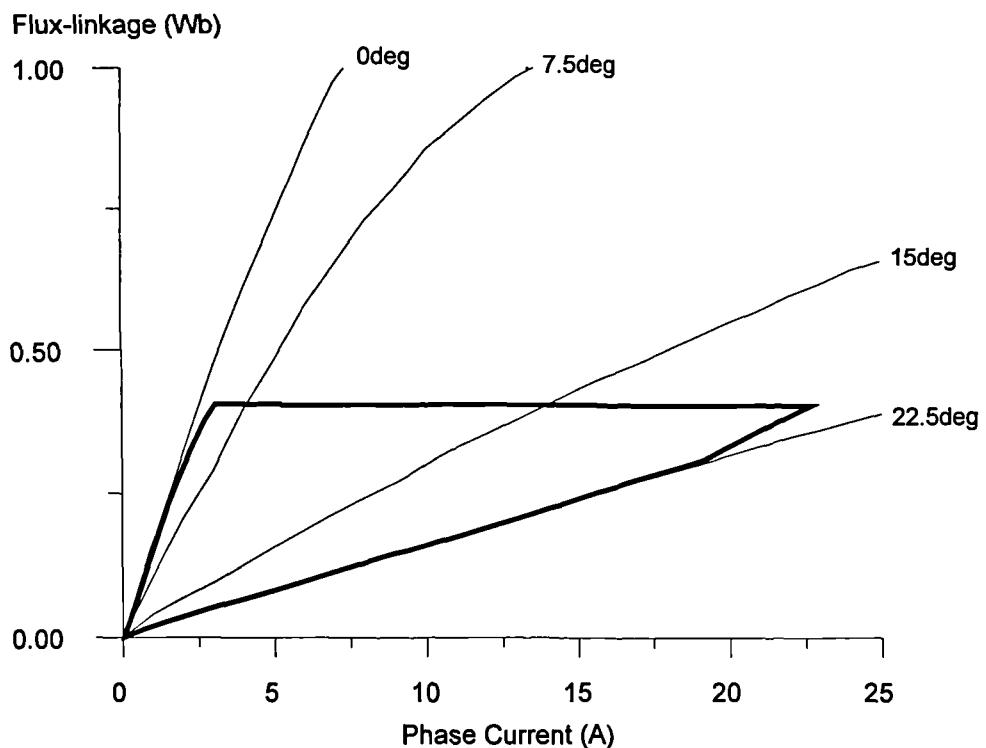


Figure 5.13 Equivalent single-tooth winding flux-linkage/current locus for voltage control of the fully-pitched winding machine.

Comparing Figure 5.13 with Figure 5.14, it is apparent that the higher peak flux-linkage in Figure 5.14 enables the same torque to be produced for a lower peak current. For maximum torque the fully-pitched winding switched reluctance drive has a flux waveform that is at its peak value, from near the unaligned position, to near the aligned position. This leads to a large peak current near the unaligned position. However the single-tooth winding switched reluctance drive flux waveform allows peak flux linkage to be placed nearer to the aligned position, thus reducing the peak current whilst still maintaining output.

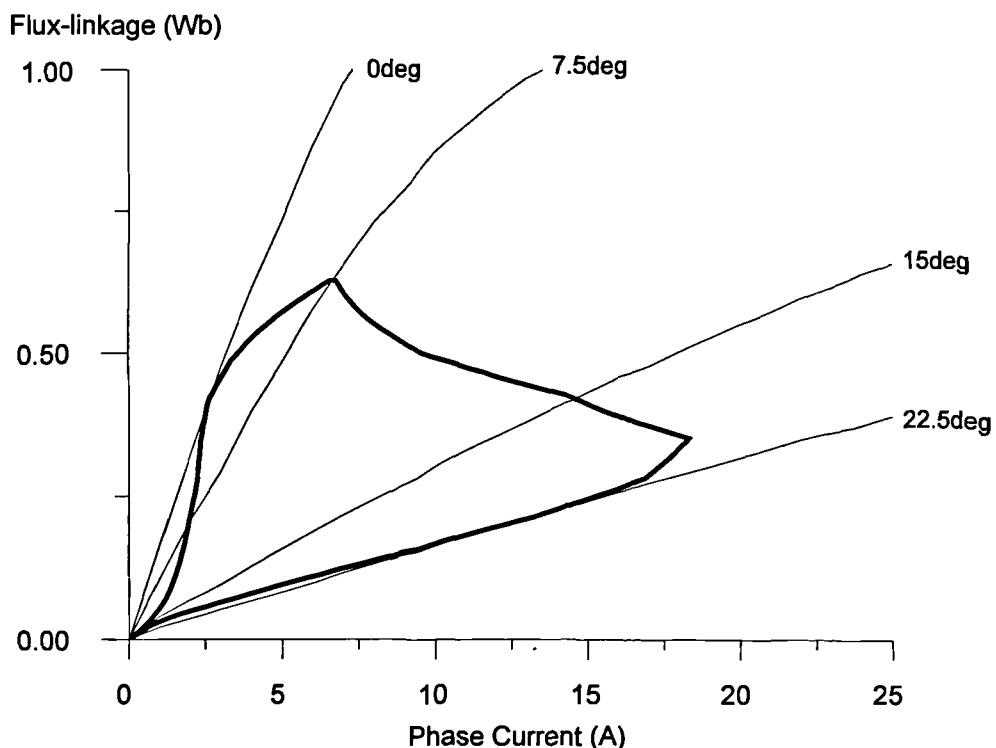


Figure 5.14 Single-tooth winding flux-linkage/current locus for voltage control of the single-tooth winding machine.

Table 5.1 provides a comparison between the operation of a fully-pitched winding and single-tooth winding switched reluctance drive, in terms of both the fully-pitched winding, and single-tooth winding parameters. This shows that although the peak phase currents are greater for the fully-pitched winding switched reluctance drive, the average and RMS values are lower. This is due to the discontinuous nature of the equivalent single-tooth winding current. Machine copper losses are approximately 20% lower in the fully-pitched winding machine under these conditions. However when the converter ratings required to produce the same output are examined, it can be seen that the fully-pitched winding switched reluctance drive requires a rating some 40% higher.

Parameter	Type	Fully-pitched winding switched reluctance drive	Single-tooth winding switched reluctance drive	Units
Real, or equivalent single-tooth winding phase current	peak	22.98	18.3	A
	average	5.23	6.6	A
	RMS	8.49	8.55	A
Real, or equivalent fully-pitched winding phase current	peak	13.02	10.8	A
	average	5.23	6.6	A
	RMS	6.47	7.16	A
Shaft torque	average	19.7	20.3	Nm
Shaft power	average	3.09	3.19	kW
Converter rating ($6 \times V_{dc} \times I_{ph}$)	peak	39.0	27.4	kVA
	RMS	19.4	12.8	kVA
Converter kVA/kW	peak	12.6	8.6	-
Machine copper loss	average	144	174	W

Table 5.1 Comparison of voltage control for fully-pitched winding and single-tooth winding switched reluctance drives.

Note that in Table 5.1 the converter rating is defined as being the product of 3 (phases) \times 2 (switches per phase) \times dc-link voltage \times phase current.

5.3 Experimental Operation of the Fully-Pitched Winding Switched Reluctance Drive

Results in this section were obtained from measurements made on the experimental fully-pitched winding switched reluctance machine drive, and are for unipolar operation. The power converter was configured to have a two transistor chopper per phase (asymmetrical half-bridge), as is common for conventional single-tooth winding

switched reluctance drives. To do this two of the four transistors in each 'H' bridge are disabled, and are thus effectively removed from the circuit. Control software is configured for two processor operation, with the third processor providing real time monitoring only. An adaptive PID controller is used for phase current control, with phase current references of the form shown in Figure 5.1 and Figure 5.4. Further details of this processor configuration and current control method can be found in Chapters 3 and 4, sections 3.1.3 and 4.1.3, and Appendix D.1.

As with any voltage fed machine drive, it is only possible to emulate an ideal current source at zero speed due to the finite supply voltage. As speed increases the phase current takes a progressively greater proportion of the cycle time to reach the desired level, until at high speed the current never reaches the desired current, and operation is fully under voltage control. For much of a drive's speed range operation is neither with full current, nor full voltage control, but voltage control with chopping when the current reaches the desired level, *Idem*. For the experimental drive operating at speeds less than 125rpm the current controller spends most of the time chopping at the desired current level and so the machine can be considered to be under current control. Above 1200rpm the current controller is only able to reach the desired current levels for a very short time each cycle, chopping is almost eliminated, and so the drive can be considered to be in voltage control. Operation under these two sets of conditions is presented in the following sections.

5.3.1 Low Speed Operation

In this section all results and waveforms are for a fixed set of conditions as follows:- motoring at 101rpm with a current demand *Idem* (see Figure 5.4) of 10A and a dc-link voltage of 480V. The average shaft torque is 56Nm.

If a pair of stator and rotor teeth are excited by a constant MMF, then in an idealised machine the torque stays constant until the aligned position is reached. In the experimental machine however, the torque drops off rapidly as the aligned position is approached. Because of this maximum average torque is generated if the driving MMF is reduced before the stator and rotor teeth are fully aligned. For the experimental machine the optimum 'on' and 'off' angles were found to be 262° and 22° respectively (see Figure 5.4). This corresponds to the reduction of the driving MMF to zero 2.75° mechanical before alignment. As can be seen in Figure 5.15, the actual phase currents follow the ideal currents (as shown in Figure 5.1), the main discrepancy being a finite rise and fall time. The bottom trace in Figure 5.15 is one of the equivalent single-tooth winding currents which is calculated in real time from the real phase currents.

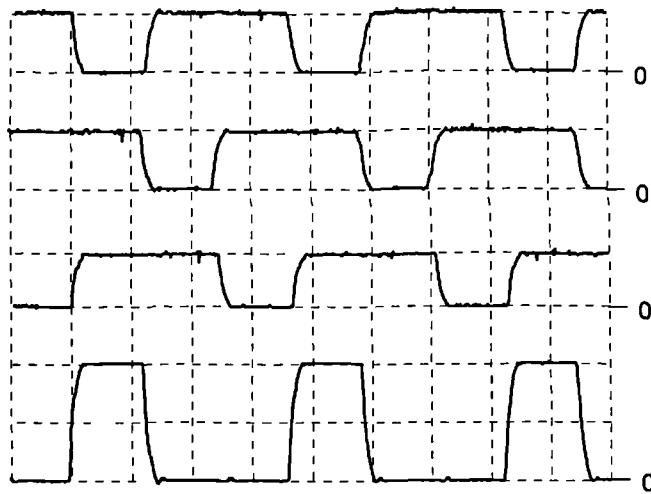


Figure 5.15 Low speed real and equivalent current waveforms. Traces 1, 2 and 3 - I_a, I_b, I_c , 10A/div. Trace 4 - Equivalent single-tooth winding current - I_1 , 10A/div. Horizontal - 20ms/div.

Owing to the mutual coupling between phases, any change in one phase current will tend to change the flux linking all three phases. Hence half way through the conduction period of one phase, as the other two phases commutate there is a large change in flux linkage, as shown in Figure 5.6. Due to the finite controller bandwidth and dc-link voltage, the controller is unable to supply the necessary volt seconds instantly, resulting in a current disturbance. In Figure 5.15 the current disturbance is too small to show up on the scale used, as the controller has been designed to minimise this effect. However it is evident in Figure 5.16, as the current variation around 10A. This figure is the flux-linkage/current locus for the real (fully-pitched winding) current and flux-linkage in one phase of the machine.

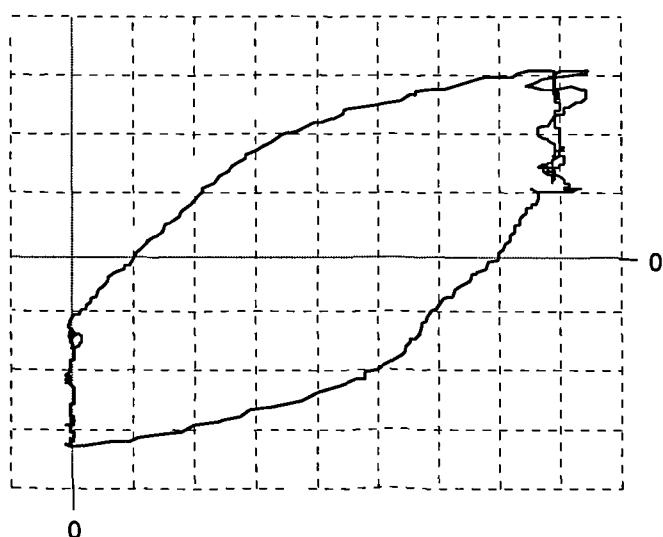


Figure 5.16 Low speed fully-pitched winding ψ, i locus. Horizontal - Phase current I_a , 1.25A/div. Vertical - Flux-linkage ψ_a , 0.4Wb/div.

Although presented here as a two dimensional graph, the flux linking a phase is actually a non-linear function of four variables: the three phase currents, and rotor position. As with a conventional single-tooth winding switched reluctance machine, the area enclosed within the fully-pitched winding flux-linkage/current locus represents the energy converted to mechanical output. Although the average shaft torque can be related to the flux-linkage/current locus area, it is not easy to predict the magnetic circuit utilisation or identify torque producing areas from the fully-pitched winding flux-linkage/current locus. However, by using the equivalent single-tooth winding flux-linkage/current locus it is possible to predict average torque, magnetic circuit utilisation, and identify torque producing regions by mental superposition of the magnetisation curves. Figure 5.17 shows the equivalent single-tooth winding flux-linkage/current locus, which is similar in appearance to the flux-linkage/current locus of a conventional single-tooth winding switched reluctance drive, operating at low speed. Note that the area enclosed within the two loci are equal. Although the actual currents are unipolar, the equivalent single-tooth winding current has a small negative current 'tail'. This is because when the real fully-pitched winding currents are switched, the finite rise and fall times cause a small change in overlap between phases, which when transformed to the equivalent single-tooth winding currents, gives rise to short periods of negative current.

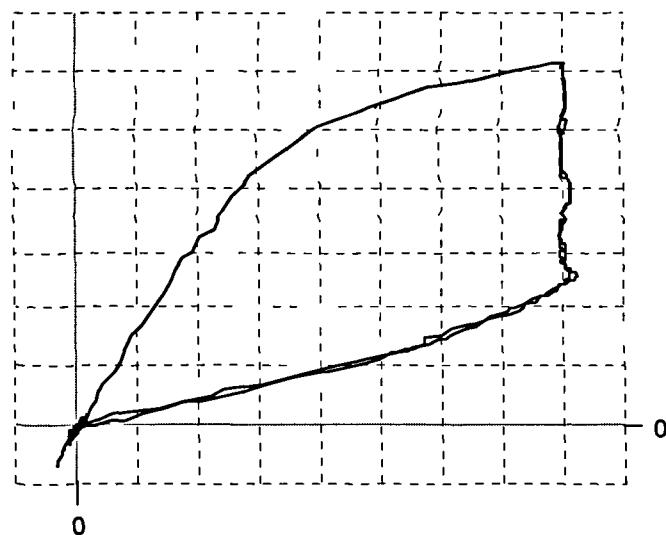


Figure 5.17 Low speed equivalent single-tooth winding ψ, i locus. Horizontal - Phase current I_1 , 2.5A/div. Vertical - Flux-linkage ψ_1 , 0.2Wb/div.

From the equivalent single-tooth winding currents, in conjunction with stored machine characteristics, it is possible to calculate the instantaneous flux-linkage and torque for each equivalent phase, as well as the total torque. Traces in Figure 5.18 show the equivalent single-tooth winding current, flux-linkage and instantaneous torque. The

bottom trace shows the total torque, which is calculated from the sum of the three single-tooth winding torques. All the transformed and derived parameters were calculated in real time by the drive control processors.

As phase currents are constant during their respective conduction intervals and there is no overlap of conduction intervals, the torque produced by each phase follows the machine torque/position characteristics, which determine the average torque and ripple. At higher phase currents, the shaft torque becomes more non-linear with position (see static torque/position characteristics in Appendix A.4) due to saturation effects. The effect of this is that the peak to peak torque ripple increases as a percentage of the average torque with rising phase currents. With a phase current demand (I_{dem}) of 5A there is 59% pk-pk torque ripple, increasing to 94% at 10A. It is not possible to produce overlapping equivalent single-tooth winding current pulses, in order to reduce torque ripple, without increasing the peak fully-pitched winding current, and dramatically altering its waveshape. This aspect is covered in more depth in Chapter 8, which looks at constant torque operation.

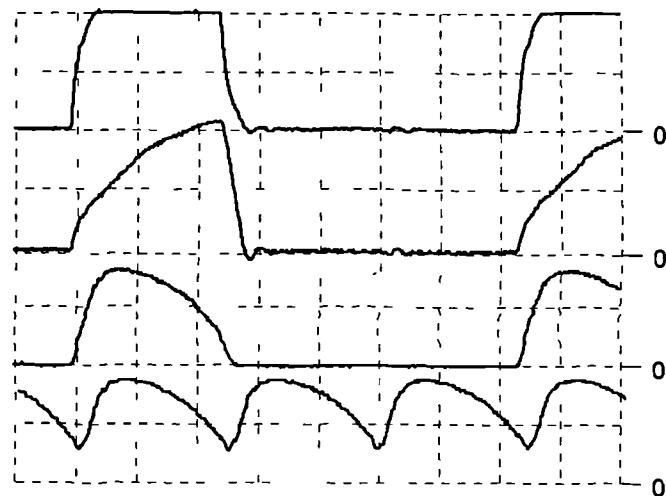


Figure 5.18 Low speed equivalent single-tooth winding current, flux-linkage and torque. Trace 1 - Phase current I_1 , 10A/div. Trace 2 - Flux-linkage ψ_1 , 0.5Wb/div. Trace 3 - Single phase torque T_1 , 50Nm/div. Trace 4 - Total torque, 50Nm/div. Horizontal 10ms/div.

At low speed altering the phase conduction period from the ideal pattern of 240° on 120° off, has a detrimental effect on torque production. This is most easily seen by reference to the equivalent single-tooth winding current and corresponding torque. Figure 5.19 shows operation at low speed, with the conduction angle increased from the ideal of 240° to approximately 260° . The equivalent single-tooth winding current no longer has a only a single pulse 120° long, as can be seen in the trace three of Figure 5.19. The additional current pulse, approximately 180° from the main pulse, produces a

negative torque component. The overall effect on shaft torque is to reduce average torque, and increase torque ripple. Note that although this effect is very easy to see in the equivalent single-tooth winding parameters, it is not immediately obvious from the fully-pitched winding parameters that a small increase in conduction interval will reduce average torque and increase torque ripple. In terms of slot MMF, the extended conduction interval results in pairs of stator and rotor teeth being excited *after* they have passed the aligned position, thus producing a negative torque component.

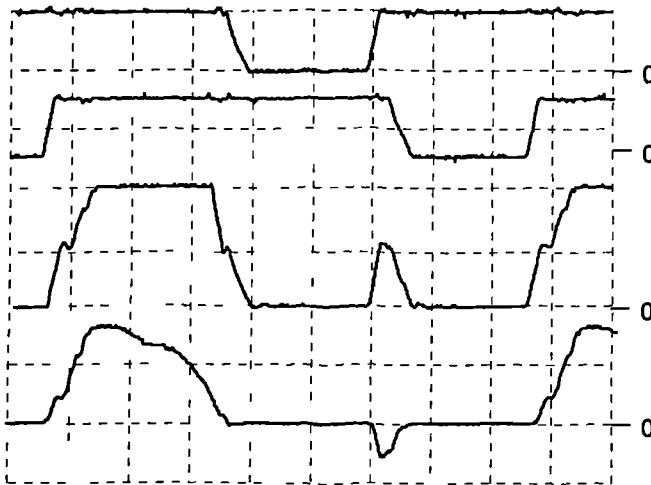


Figure 5.19 The effects of sub-optimum conduction widths at low speed. Trace 1 and 2 - Phase currents I_b and I_c , 10A/div. Trace 3 - Equivalent single-tooth winding current I_1 , 10A/div. Trace 4 - Single phase torque T_1 , 50Nm/div. Horizontal 10ms/div.

5.3.2 High Speed Operation

All the results in this section are for a fixed set of conditions as follows:- motoring at 1225rpm with a current demand I_{dem} of 10A and a dc-link voltage of 480V. The average shaft torque produced is 16Nm.

At low speeds the phase current conduction angle is 240° , but as speed increases the drive moves towards full voltage control with positive volts applied during the 'on' time and negative volts during the 'off' time. In order to balance the applied volt-seconds the conduction angle must be reduced, ultimately to 180° , otherwise a large DC current will flow. As in a conventional SR drive, the conduction sequence must be advanced relative to rotor position, so as to maximise the current in the positive torque producing region. For the results presented here, the 'on' and 'off' angles are 286° and 88° respectively, giving a 'conduction' angle of 198° .

At high speed the inability of the drive to provide the volt-seconds required for constant phase current begins to have a considerable impact on phase current waveshape when

generating significant torque. The phase current takes on a characteristic 'twin peaks' waveshape, as shown in the top three traces of Figure 5.20. After rising from zero to an initial peak value, the phase 'a' current drops slowly as the 'motional emf' exceeds the applied dc-link voltage at this point. The current then drops more quickly, as negative volts are applied to switch off phase 'c' and couple to phase 'a', adding to the 'motional emf'. Next, as phase 'b' is switched on the positive volts applied to that phase couple to phase 'a' aiding the applied voltage, and phase 'a' current starts to rise. Finally at the end of the cycle negative volts are applied, and the current falls.

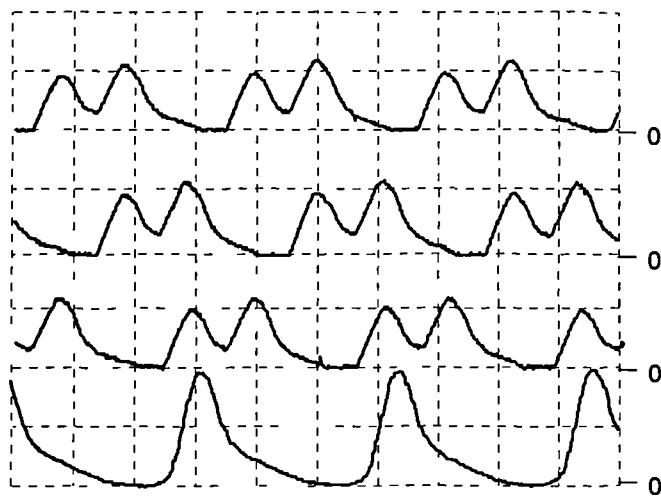


Figure 5.20 High speed real and equivalent current waveforms. Traces 1, 2 and 3 - I_a , I_b , I_c , 10A/div. Trace 4 - Equivalent single-tooth winding current - I_1 , 10A/div. Horizontal - 2ms/div.

The explanation given above can help explain the resulting waveforms, but does not give an insight into what is really happening. Throughout the cycle, the emf induced in a phase is a function of the three phase currents and their rate of change, together with the rate of change of position. This dependency on so many variables, all of which are changing rapidly throughout a cycle, makes understanding the fully-pitched winding currents quite complex. Looking at the equivalent single-tooth winding current, the bottom trace in Figure 5.20, it can be seen that the complex fully-pitched winding current transforms to a familiar looking single-tooth winding current. The single-tooth winding current has a waveshape similar to that of a conventional single-tooth winding switched reluctance drive, operating at high speed. Current rises quickly during the period of low inductance, and then falls driven down first by the 'motional emf' and then also by applied volts as the phase is switched off. As the inductance rises approaching the aligned position the current decay gets slower and slower.

Although the equivalent single-tooth winding current looks similar to that which could be expected for a conventional single-tooth winding switched reluctance drive, there are

significant differences, as explained in section 5.2.2 of this chapter. Figure 5.21 shows several real and equivalent single-tooth winding parameters. The top trace shows that there is a small amount of chopping in the actual phase voltage, as the current exceeds the 10A demand level, before the negative volts are applied to reduce the current to zero. As the current reaches zero there is also chopping, reducing the voltage. Trace four shows the equivalent single-tooth winding voltage, this has relatively short positive and negative voltage pulses, with periods of near zero volts between. The resulting flux-linkage waveform, which was calculated from the applied voltage, is shown in the bottom trace. From the single-tooth winding flux-linkage it is easy to see how the single-tooth winding current waveform arises; with current falling at constant flux-linkage as the aligned position is approached. Only now, by transforming the single-tooth winding current back to fully-pitched winding values, can the fully-pitched winding current waveform be properly explained.

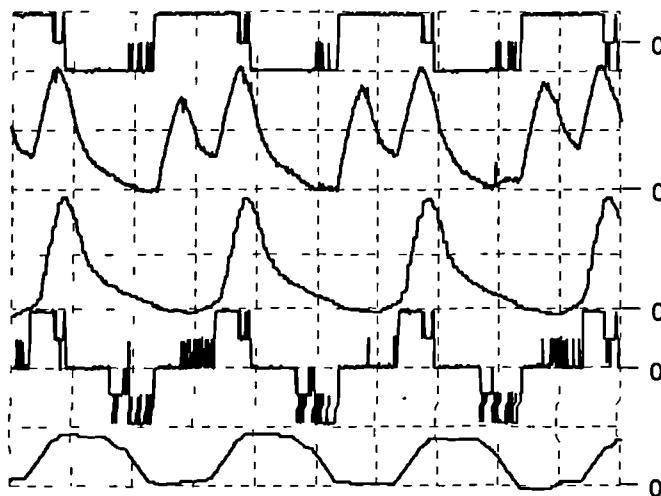


Figure 5.21 High speed real and equivalent current, voltage and flux-linkage waveforms. Trace 1 - Phase voltage, 1kV/div. Trace 2 - Phase current, 5A/div. Trace 3 - Equivalent current, 10A/div. Trace 4 - Equivalent voltage, 500V/div. Trace 5 - Equivalent flux-linkage, 0.5Wb/div. Horizontal - 2ms/div.

Figure 5.22 shows the real fully-pitched winding flux-linkage/current locus. It is very difficult to gain any insight into machine operation from this, other than to note that the enclosed area is proportional to mechanical output.

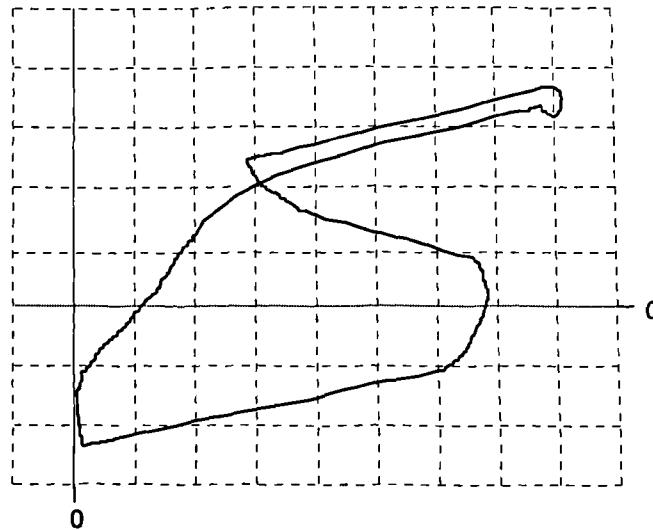


Figure 5.22 High speed fully-pitched winding ψ, i locus. Horizontal - Phase current I_a , 1.25A/div. Vertical - Flux-linkage ψ_a , 0.2Wb/div.

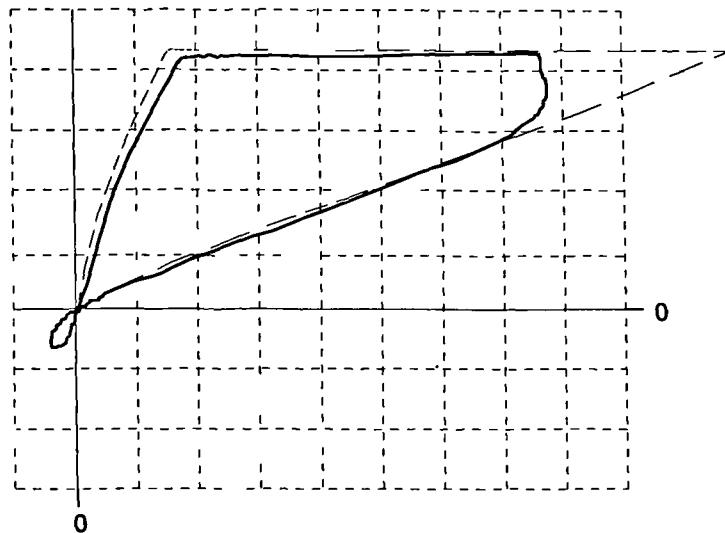


Figure 5.23 High speed equivalent single-tooth winding ψ, i locus. Horizontal - Phase current I_1 , 2.5A/div. Vertical - Flux-linkage ψ_1 , 0.1Wb/div. Dashed trace - Simulated locus.

As described in section 5.2.2 of this chapter, the equivalent single-tooth winding flux-linkage/current locus is significantly different for a fully-pitched winding switched reluctance drive under voltage control. The experimental equivalent single-tooth winding flux-linkage/current locus is shown in Figure 5.23. Also shown in this figure, as the dashed trace, is the simulated flux-linkage/current locus. This is for full voltage control with the same 'on' angle and speed as the measured results. It was simulated as described in section 5.2.2 of this chapter. Differences exist between the loci because of the short time spent chopping at the peak current, evident in the much reduced peak current, and finite resistive voltage drop, seen as the slight widening of the gap between

traces at constant flux. These are both aspects of the real drive which are not modelled in the simulation.

5.3.3 Power Converter Rating

The flux-linkage/current loci shown in Figure 5.23 brings out important points about the power converter rating of a fully-pitched winding switched reluctance drive in voltage control. It is evident that the peak current in the simulated locus is very much higher than the actual peak current, but the areas enclosed within the two loci are not very different. In the simulation, full positive voltage is applied for 180° with no current limiting, but in the real drive the controller applies negative volts to the fully-pitched winding when the current exceeds the demand level. This limits the fully-pitched winding current and the equivalent single-tooth winding peak current.

measured equivalent
 The area enclosed by the *measured equivalent* flux-linkage/current locus was determined to be 4.16J, this corresponds to an average torque of 15.9Nm. This compares exceptionally well with the measured torque of 16Nm. The area within the simulated locus was calculated to be 4.93J, and corresponds to an average torque of 18.8Nm. These figures, as well as ones for the simulated and measured phase currents are shown in Table 5.2.

Parameter	Experimental Drive	Simulated Drive	Units
Area enclosed within flux-linkage/current locus	4.16	4.93	J
Peak single-tooth winding current	19.4	27.0	A
Peak fully-pitched winding current	10.2	15.5	A
Ratio of torque (from locus area)	1.18		-
Ratio of peak fully-pitched winding currents	1.52		-

Table 5.2 Comparison of experimental and simulated results for operation at 1225rpm.

From these figures it can be seen that current limiting in the real drive reduces torque by a factor of 1.18, and the peak fully-pitched winding current by a factor of 1.52, comparing figures for the measured and simulated results. Thus by current limiting the actual drive reduces the kVA/kW required by a factor of 1.29. In section 5.2.2 of this chapter figures were given in Table 5.1 comparing simulated fully-pitched winding and

single-tooth winding switched reluctance drives. These results showed that the fully-pitched winding switched reluctance drive needed a peak converter rating 1.42 times higher, when producing the same torque. Figures in Table 5.3 compare converter ratings for full voltage control at 1500rpm, with the 'advance' angle set for maximum torque or minimum converter rating.

Parameter	Advance angle set for optimum torque	Advance angle set for optimum kVA/kW
Fully-pitched winding switched reluctance drive kVA/kW	12.8	11.25
Single-tooth winding switched reluctance drive kVA/kW	8.5	8.0
Ratio of kVA/kW	1.51	1.41

Table 5.3 Comparison of converter ratings for fully-pitched winding and single-tooth winding switched reluctance drives.

These simulated figures suggest that a fully-pitched winding switched reluctance drive would require a converter with a peak rating at least 41% higher. However, comparison of simulated results with the actual drive, showed that by limiting the peak current the converter kVA/kW rating was reduced 29%. Taking this factor into account, then for the same shaft power, the fully-pitched winding switched reluctance drive would require a converter rating only 1.09 times that of the single-tooth winding switched reluctance drive. Note that if speed is increased to the point where there is no chopping at the peak current, then simulated figures for the fully-pitched and single-tooth winding switched reluctance drives are 11.1kVA/kW and 9.56kVA/kW respectively, a ratio of 1.16.

5.4 Torque/Speed Characteristics

The torque/speed characteristics presented in this section are measured results for the experimental fully-pitched winding switched reluctance drive with unipolar excitation. Phase current is controlled by the adaptive PID controller previously described. The phase current reference has a constant peak current throughout the speed range so there is no abrupt change of mode from current chopping, to single pulse operation (Sugden *et al* [5.1]). At each operating point the 'on' and 'off' angles are adjusted to give the maximum torque output.

5.4.1 Operation with Series Connected Phase Coils

The phase windings of the fully-pitched winding switched reluctance machine are composed of two coils. With the coils connected in series, the fully-pitched winding switched reluctance machine has the same total number of turns per phase as the experimental single-tooth winding switched reluctance machine. To produce the same slot MMFs for a given phase current and unipolar excitation, the fully-pitched winding machine needs half as much current per phase as the single-tooth winding machine. This is the configuration used throughout most of this thesis.

Torque and shaft power characteristics given in the following figures are for a dc-link voltage of 480V nominal whilst motoring, and with phase current demands I_{dem} of 5A and 10A (see Figure 5.4).

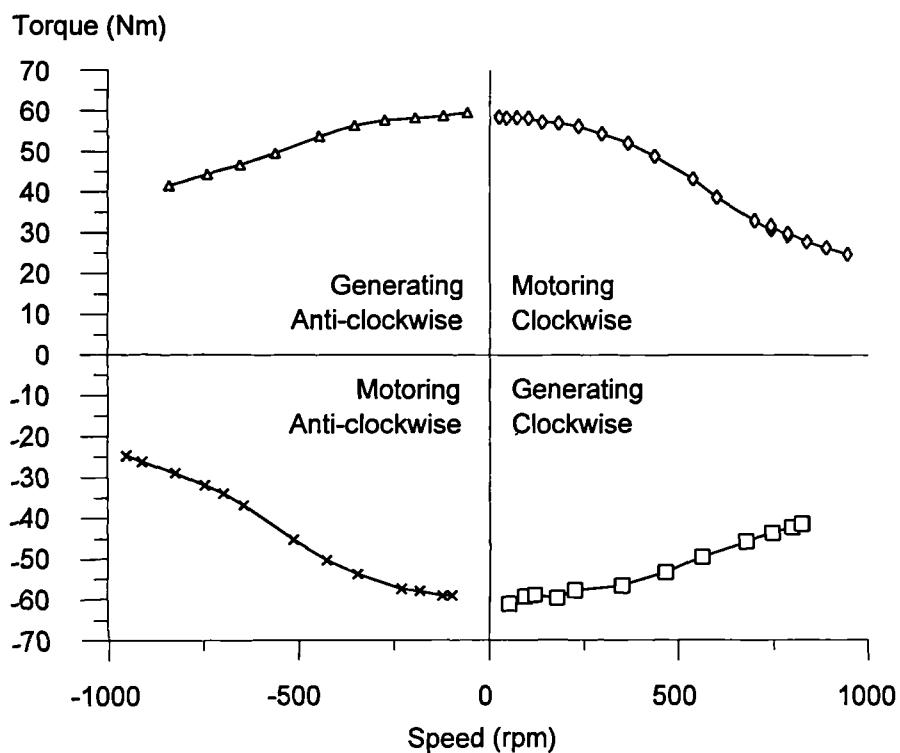


Figure 5.24 Four quadrant torque/speed envelope for $I_{dem} = 10A$.

All four quadrants of operation are shown in Figure 5.24 and Figure 5.25. At low speed there is a region of approximately constant torque, of same amplitude in all four quadrants, where the converter is operating as a near perfect current source. The torque then falls with increasing speed as the actual currents deviate further from their reference values. Most obvious in Figure 5.25, is that the torque in generating quadrants falls off more slowly than when motoring, and so the shaft power continues to increase. This is not because the machine inherently has a higher base speed for generating in comparison to motoring, but is due to the controller and power converter. When

generating, the energy returned to the converter increases the dc-link voltage until the dc-link dump circuit comes into operation, this happens at approximately 515V with an average dc-link voltage whilst generating of 510V. This increase in dc-link voltage, compared to the 480V nominal whilst motoring, gives greater torque at high speed. The second reason, and one which becomes more significant as speed increases, is current controller bandwidth. When generating the controller must first apply positive volts to establish some phase current, the 'motional emf' then increases rapidly and with it the phase current. The controller must then reverse the applied voltage, to extract energy from the machine. As the current rises very rapidly at high speed the current controller is unable to react fast enough, and the current overshoots the desired level. In fact the controller will lose control of the current when the 'motional emf' exceeds the dc-link voltage. This causes the peak phase current whilst generating at high speed to be higher than the demand value, and hence more torque is produced.

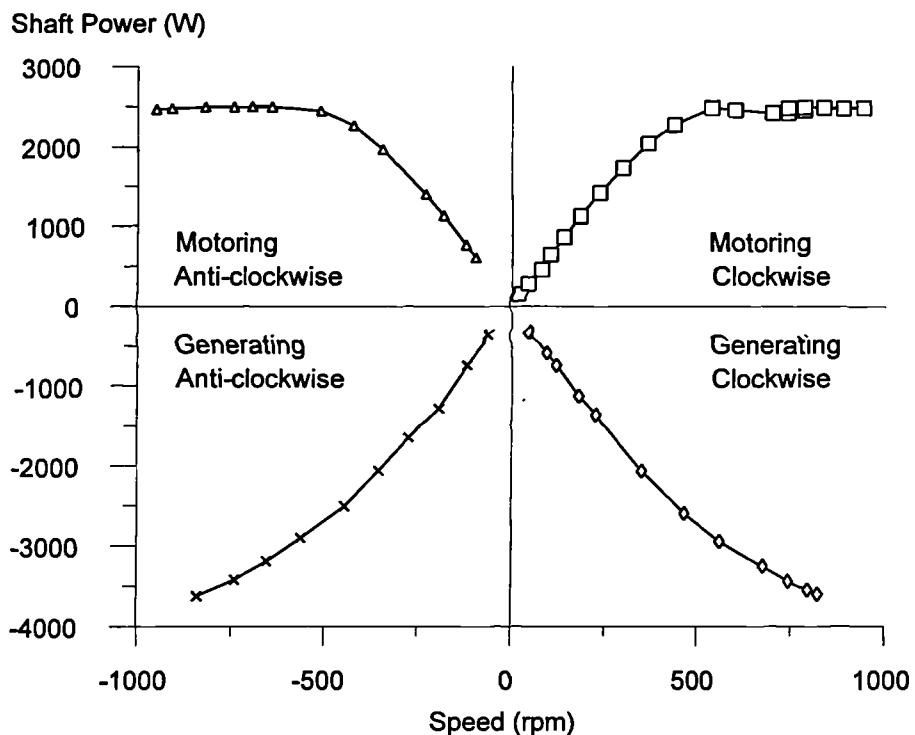


Figure 5.25 Four quadrant shaft power envelope for $I_{dem} = 10A$.

From Figure 5.25 it is clear that the base speed when motoring is approximately 600rpm. The motoring power profile is similar to many other types of electric drive; with power rising linearly with speed up to base speed, and then remaining approximately constant from that point on. For a non-saturating switched reluctance machine, the torque is proportional to the phase current squared, but in a real machine magnetic circuit saturation reduces this substantially at higher currents. This effect can

be seen by comparing the low speed torque produced with current demands I_{dem} of 5A and 10A, as shown in Figure 5.26.

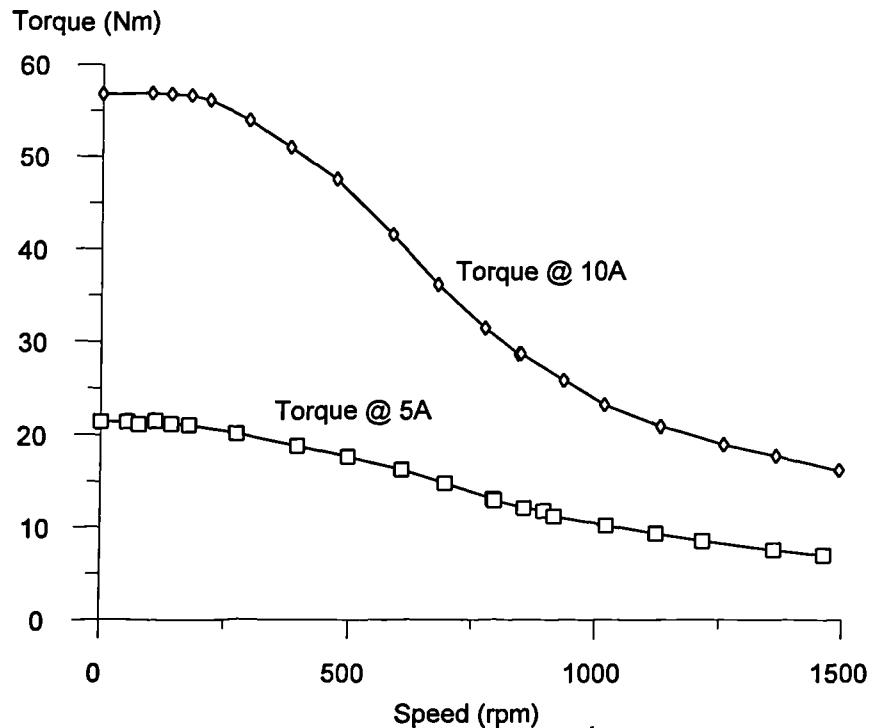


Figure 5.26 Torque/speed envelope for $I_{dem} = 5A$ and 10A.

Figure 5.27 shows that the base speed is the same at both phase current demand levels.

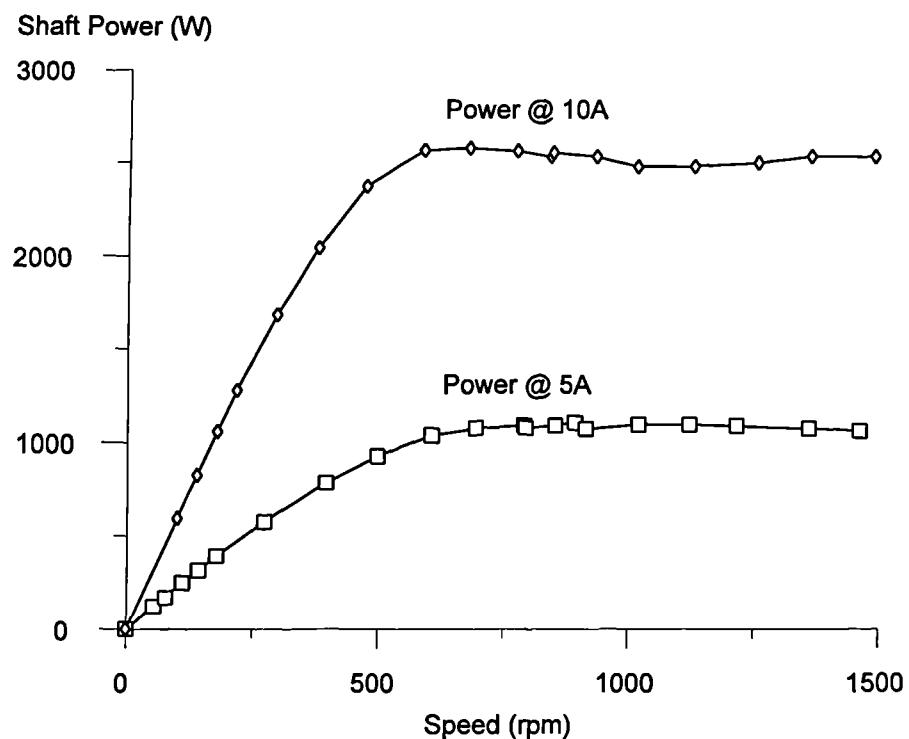


Figure 5.27 Shaft power envelope for $I_{dem} = 5A$ and 10A.

5.4.2 Operation with Parallel Connected Phase Coils

Series connection of coils in the fully-pitched winding switched reluctance machine gives the same number of turns per phase as that in the experimental single-tooth winding switched reluctance machine. This requires twice the current per phase in the single-tooth winding machine, to produce the same slot MMFs. For equal converter VA rating, the dc-link voltage in the single-tooth winding switched reluctance drive should be half that of the fully-pitched winding switched reluctance drive, as explained in section 5.2 of this chapter. In order to make a more direct comparison, the fully-pitched winding machine was re-configured with parallel connection of the two fully-pitched winding phase coils. This is so the same current per phase gives the same slot MMFs, in the fully-pitched winding and single-tooth winding machines. The converters can therefore have the same current rating and dc-link voltage, for the same VA rating.

The torque characteristics for parallel connected coils were measured with a dc-link voltage of 420V nominal and phase current demand I_{dem} of 20A (see Figure 5.4). The low speed torque (56Nm) was found to be the same as for series connected coils and a phase current demand of 10A. This is as would be expected since the slot MMFs' are identical in each case. Examination of the shaft power characteristics shows that the base speed for parallel connected coils is approximately 1000rpm at a dc-link voltage of 420V. This would give a base speed of 1150rpm at 480V, which as theory suggests is approximately twice the base speed for series connection of 600rpm.

5.5 Operation of the Single-tooth Winding Switched Reluctance Drive

In order to make valid comparisons between the performance of a fully-pitched winding switched reluctance drive and a single-tooth winding switched reluctance drive, the experimental converter and controller were connected to the experimental single-tooth winding switched reluctance machine.

5.5.1 Experimental waveforms

In this section measured results for unipolar operation of the experimental drive and single-tooth winding switched reluctance machine are presented. For these results the power converter has a full 'H' bridge per phase, although the current demand is unipolar. This is why the phase current is able to go negative for short periods during each cycle, as the phase current undershoots the demand value. Control software is configured for two processor operation, with the third processor providing real time monitoring only. The phase current controller is an adaptive PID controller, with phase current references

of the form shown in Figure 5.4. Further details of this processor configuration and current control method can be found in Chapters 3 and 4, sections 3.1.3 and 4.1.3, and Appendix D.1.

The standard test configuration, for unipolar operation of the fully-pitched winding switched reluctance drive, is with a current demand I_{dem} of 10A and 240° conduction. The corresponding excitation for the single-tooth winding switched reluctance drive is with I_{dem} set to 20A, and a nominal conduction angle at low speed of 120°. For the single-tooth winding switched reluctance machine results a dc-link voltage of 420V was used. The top trace in Figure 5.28 shows low speed operation with 120° conduction. By increasing the conduction angle to 180°, as shown in trace two of Figure 5.28, the average torque can be increased and torque ripple reduced. In each case the actual currents follow the reference closely.



Figure 5.28 Phase current at low and medium speed. All traces are 10A/div. Trace 1 - 120° conduction, 60Nm, 178rpm, 10ms/div. Trace 2 - 180° conduction, 64Nm, 188rpm, 10ms/div. Trace 3 - 52Nm, 756rpm, 2ms/div.

Trace three in Figure 5.28 is for operation at the higher speed of 756rm, at this speed the current is starting to deviate significantly from the reference. The long current tail approaching the high inductance aligned position is clearly evident. At this speed the 'motional emf' and dc-link voltage are almost equal, hence the fairly constant current plateau at a level less than the demand value.

At the much higher speed of 1350rpm the 'motional emf' is greater than the dc-link voltage over much of the conduction interval, and the current takes on a characteristic shape as illustrated in Figure 5.29. Current rises rapidly near the unaligned position, and then falls quickly as the 'motional emf' rises rapidly, then as the applied volts reverse the current drops again, with a long tail approaching the aligned position.

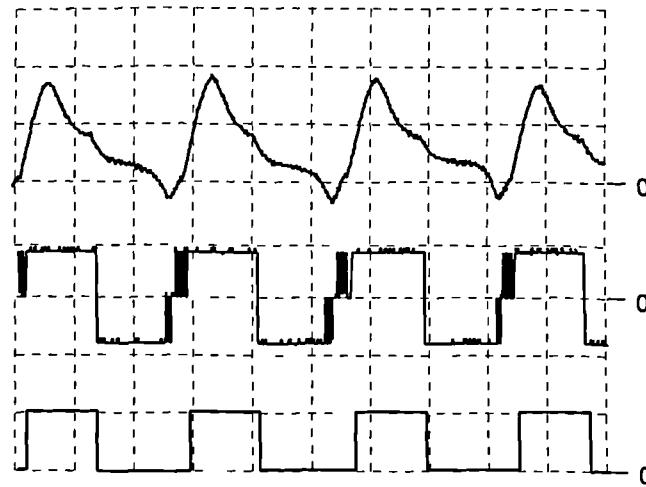


Figure 5.29 High speed operation, 1350rpm, 28Nm. Trace 1 - Phase current, 10A/div. Trace 2 - Phase voltage 500V/div. Trace 3 - Current demand, 20A/div. Horizontal 2ms/div.

5.5.2 Torque/speed Characteristics

Measured torque/speed characteristics presented in this section are for the experimental single-tooth winding switched reluctance drive. Phase current is controlled by the adaptive PID controller previously described. The phase current reference has a constant peak current throughout the speed range, and at each operating point the 'on' and 'off' angles are adjusted to give the maximum torque output. The dc-link voltage is 420V.

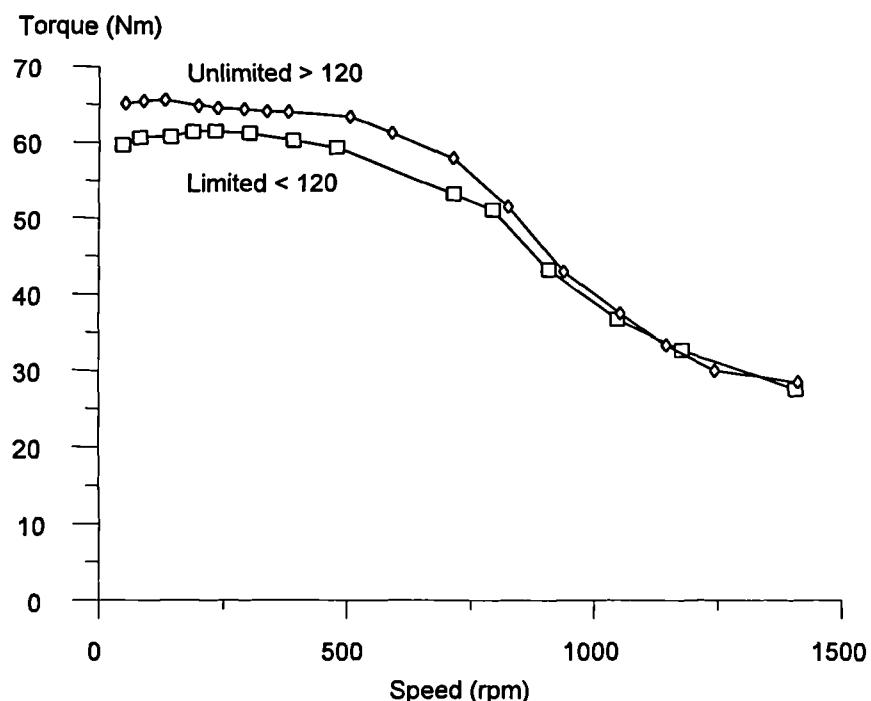


Figure 5.30 Torque/speed envelopes for $I_{dem} = 20A$, conduction widths unlimited and limited to 120° .

Figure 5.30 gives the torque/speed envelope for a current demand of 20A. The 'limited' curve has the conduction angle limited to $\leq 120^\circ$ at low speed. This mode of operation corresponds to unipolar operation of the fully-pitched winding switched reluctance drive, with 240° conduction. For the 'unlimited' curve, the conduction angle increases to 180° at low speed in order to maximise torque output. It should be noted that by increasing the conduction angle from 120° to 180° , the torque goes up from 60Nm to 65Nm (an increase of 8%), but machine copper losses go up from 318W to 478W (an increase of 50%). Thus the torque per unit loss goes down by 38%.

5.5.3 Comparison with Fully-Pitched Winding SR Machine

Figure 5.31 compares torque/speed envelopes for the experimental controller and converter, with the fully-pitched winding and single-tooth winding switched reluctance machines. The fully-pitched winding switched reluctance machine is connected with phase coils in parallel. This is so that at low speed with equal peak phase currents, 240° conduction in the fully-pitched winding machine, and 120° conduction in the single-tooth winding machine, the slot MMFs will be the same. In each case the current demand is set to 20A and the dc-link voltage is 420V.

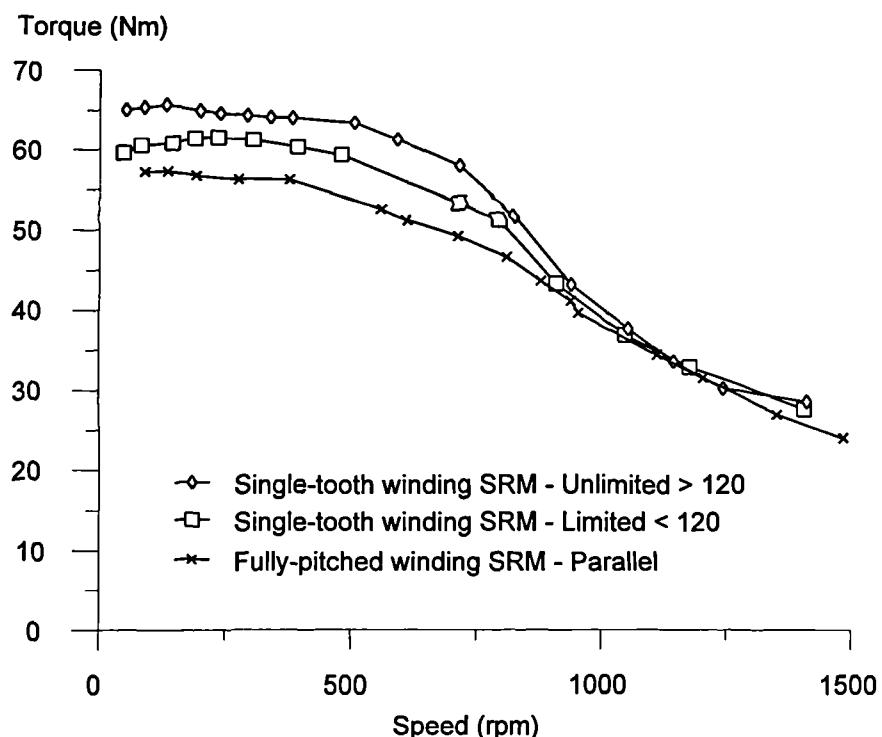


Figure 5.31 Torque/speed envelopes for the experimental fully-pitched winding and single-tooth winding switched reluctance drives.

In theory the single-tooth winding switched reluctance machine, with conduction limited to 120° , should give the same low speed torque as the fully-pitched winding

switched reluctance machine since the slot MMFs are the same. In fact, the single-tooth winding switched reluctance machine gives about 5% more torque at low speed, but as speed increases the two machines give almost equal torque. As expected, the single-tooth winding switched reluctance machine gives more torque when the conduction angle is increased to 180° at low speed.

The difference in low speed torque between the fully-pitched winding and single-tooth winding machines is not easily accounted for. Power converter and measurement circuits are the same in each case, and the mechanical dimensions of the machines are the same to within measurement tolerances of better than 0.05mm. Phase resistance measurements suggest that the turns ratio is as expected. Thus if machine mechanical and electrical dimensions are the same, the differences in torque must be due to secondary effects. It was postulated by Mecrow [5.2] that differences may be due to end-winding effects, although no evidence was given to support this. Note that Mecrow found that the fully-pitched winding switched reluctance machine gave more torque than the single-tooth winding switched reluctance machine in static torque tests. This is the opposite of low speed torque characteristics measured by the author. This makes explaining differences between sets of results even more difficult. It is worth noting however that a different single-tooth winding switched reluctance machine was used by Mecrow, to that used by the author, although the same make and model.

5.6 Summary and Conclusions

With unipolar phase currents positive torque can only be produced from periods of rising mutual inductance, and to produce continuous torque each phase should conduct for two thirds of a cycle, overlapping other phases by one third of a cycle.

Simulation has been used to predict electrical waveforms and torque, at low speed with an ideal current source, and at high speed with an ideal voltage source. Extensive use of transformations has been made to calculate equivalent single-tooth winding parameters from fully-pitched winding parameters, in both simulation, and to aid the understanding of operating waveforms. At low speed simulated and measured waveforms correspond very closely, and at high speed the small differences are due to current limiting by the real controller - an effect which is not simulated.

If a comparison is made on the basis of slot MMF and stator tooth flux between a fully-pitched winding machine drive and a conventional single-tooth winding SR drive, then it is found that the degree of correlation depends upon the operating mode. At low speed, when operating under current control, the tooth fluxes and slot MMFs are very similar, and so it could be said that the underlying mode of operation is the same.

However, at high speed when operating under voltage control, the tooth fluxes and slot MMFs differ, and so the underlying modes of operation are different. In terms of terminal current and voltage, there is a great difference at all speeds between the fully-pitched winding and single-tooth winding machines, as would be expected from their different modes of torque production.

With a constant peak current demand the torque/speed envelope of the experimental fully-pitched winding switched reluctance drive is much like many other electrical machine drives with shaft power rising up to a base speed, and then remaining constant thereafter. Comparisons made with the experimental single-tooth winding switched reluctance drive show that the torque/speed envelopes are quite similar with the control strategy employed. In terms of power converter rating, simulated and experimental figures suggest that for the same shaft power a fully-pitched winding switched reluctance drive would require a converter rating between 9% and 16% higher, depending on conditions.

Chapter 6

Bipolar Operation

6.1 Theory of Operation

6.1.1 The Method of Torque Production

Bipolar operation of a fully-pitched winding switched reluctance machine drive allows much greater scope in the choice of current waveform than for unipolar operation, as each phase can contribute to the output torque throughout a cycle. For a three phase fully-pitched winding switched reluctance machine the torque equation can be expressed as:-

$$T = i_a \cdot i_b \cdot \frac{dM_{ab}}{d\theta} + i_b \cdot i_c \cdot \frac{dM_{bc}}{d\theta} + i_c \cdot i_a \cdot \frac{dM_{ca}}{d\theta} \quad (6.1)$$

Assuming phase self inductances are constant and the magnetic circuit is linear. Positive torque is produced if each phase current pair has the same sign whilst the mutual inductance is rising, or have different signs whilst the mutual inductance is falling. Thus in contrast to unipolar operation, which only makes use of rising mutual inductance periods, bipolar operation can make use of both rising and falling mutual inductance.

The simplest mode of operation with bipolar phase currents is that with two phases excited at any one time. As shown in Figure 6.1, each phase conducts for 120° during each half cycle with current reversing every half cycle. Each phase is displaced by 120° from the others. Phase currents are synchronised to rotor position such that two phases are conducting whilst the mutual inductance between them is falling. Torque contributions from each pair of phases adds to give constant shaft torque.

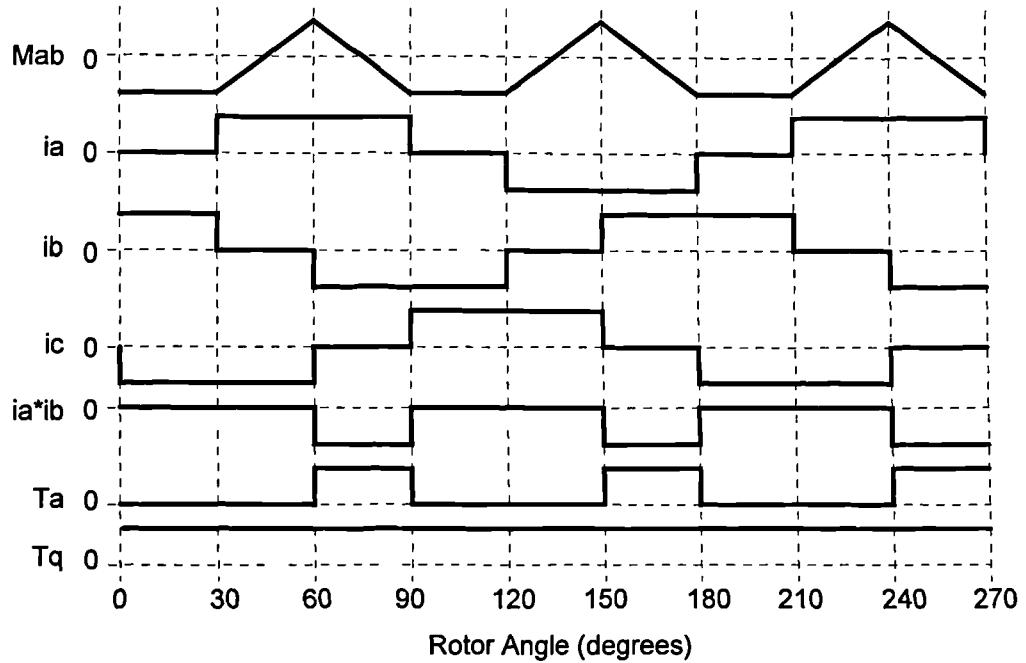


Figure 6.1 Derivation of torque from fully-pitched winding parameters with bipolar two phase on excitation.

Calculation of the equivalent single-tooth winding phase currents reveals that they have the same waveshape as the fully-pitched winding currents. Torque derivation from these currents is illustrated in Figure 6.2, and is relatively straight forward. Note that each single-tooth winding current pulse lasts for a rotation of $2/3$ of a rotor tooth pitch, and so in reality will also produce some negative torque. In the simplified derivation shown, the constant unaligned inductance region does not predict this.

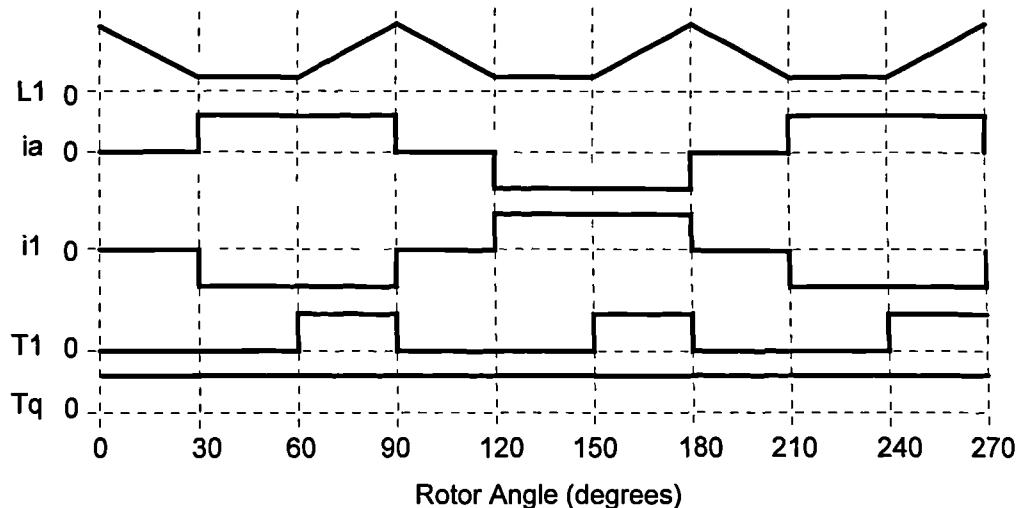


Figure 6.2 Derivation of torque from single-tooth winding parameters with bipolar two phase on excitation.

By exciting each phase with a bipolar squarewave, as shown in Figure 6.3, it is possible to use all phases simultaneously to produce positive torque, in a three phase fully-pitched winding switched reluctance machine. Phases conduct for 180° during each half cycle, with the current reversing every half cycle, and phases are displaced from each other by 120°. Phase currents are synchronised to the rotor position such that whilst the mutual inductance between a pair of phases is rising they conduct current in the same direction, and whilst the mutual inductance is falling they conduct current in opposite directions. Positive torque is thus produced from periods of rising and falling mutual inductance, with overlapping torque contributions giving constant shaft torque.

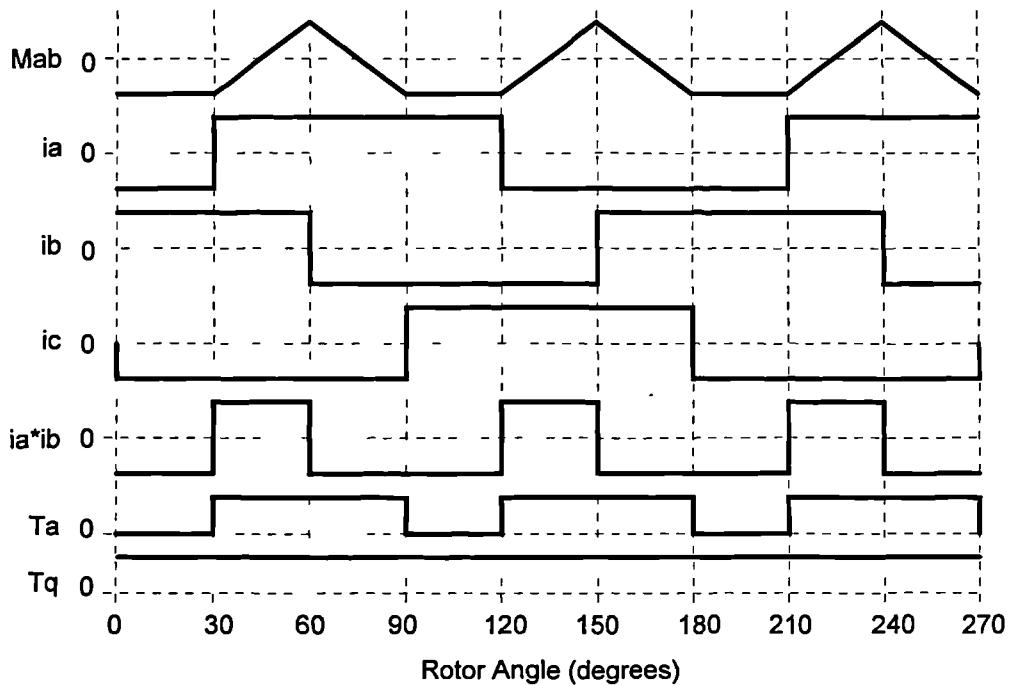


Figure 6.3 Derivation of torque from fully-pitched winding parameters with bipolar three phase on excitation.

The equivalent single-tooth winding currents have a stepped waveshape, as shown in Figure 6.4, with the higher step three times the amplitude of the lower step. The stepped and continuous nature of the equivalent single-tooth winding current makes torque calculation more difficult. As current is continuous there must be times when negative torque is being produced. In the simplified derivation of Figure 6.4 this is minimised by aligning the current peak with rising inductance. For the idealised case torque is proportional to current squared, and so the positive torque pulse is nine times greater than the negative pulse. In reality there is no long period of constant inductance and saturation reduces the torque at high currents substantially. Thus the negative torque contributions are increased and the positive contributions reduced, so overall the actual torque will be less than this simple model predicts.

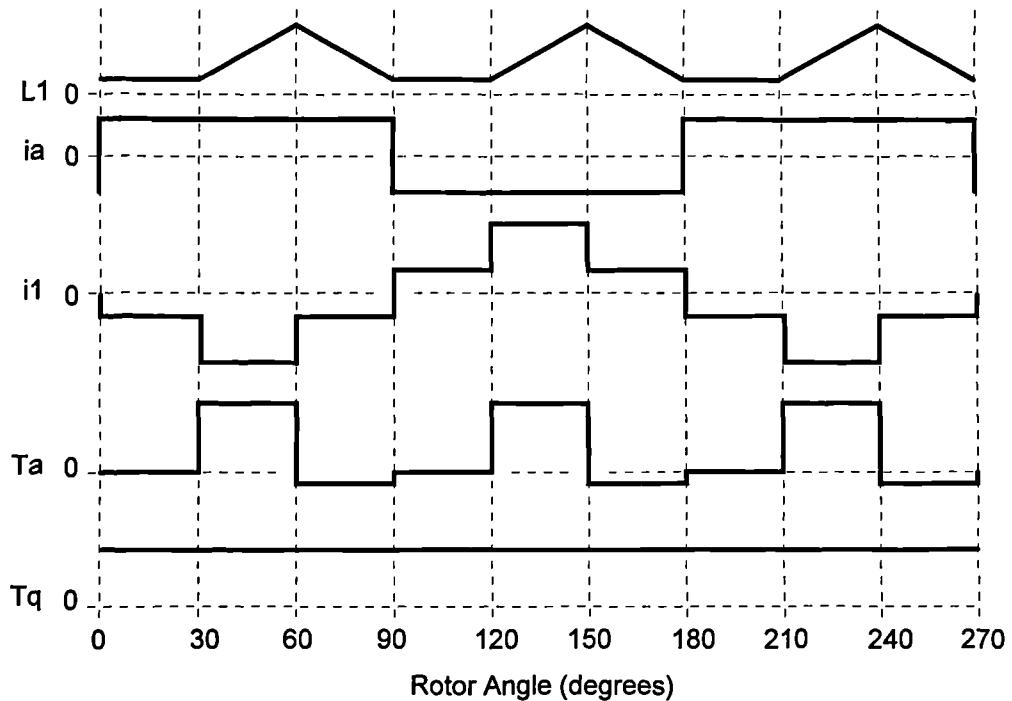


Figure 6.4 Derivation of torque from single-tooth winding parameters with bipolar three phase on excitation.

6.1.2 Phase Current Reference

The phase current reference and control angles for bipolar operation are defined in a similar way to those used for unipolar operation. For waveforms such as bipolar two phase on conduction, there is a defined conduction width of θ_{cw} during each half cycle, with an overall rotor advance angle θ_{on} . The conduction angle has no meaning for non trapezoidal currents and so only the advance angle is defined.

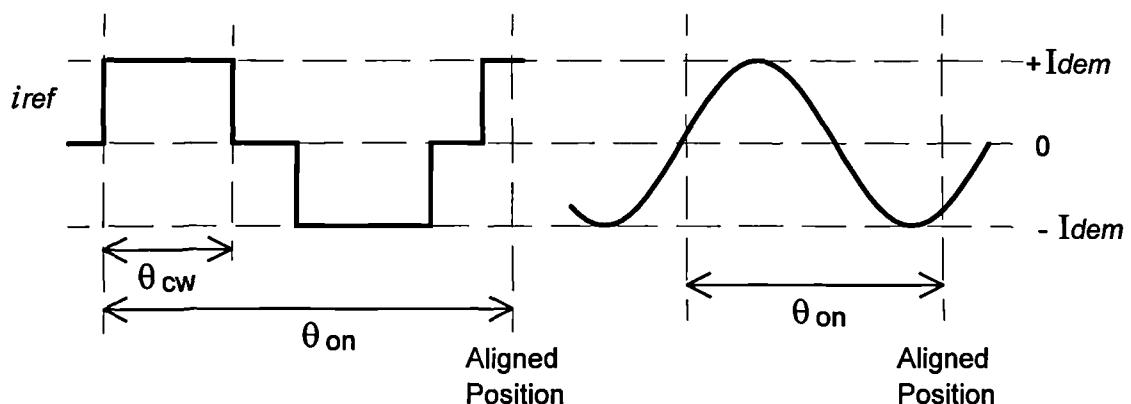


Figure 6.5 Phase current reference waveform showing control angles and current demand level.

With reference to Figure 6.5, it can be seen that the phase current amplitude is defined by I_{dem} . The control angles are measured in electrical degrees, and are given with respect to the aligned position. This is defined as being the position where a reference set of stator and rotor teeth are in alignment. One electrical cycle corresponds to a rotation of two rotor tooth pitches, which for a 12:8 machine is 90° mechanical. Current references for each phase are the same, but displaced by $1/3$ of an electrical cycle from each other.

6.2 Simulation Studies

Simulation programs used for unipolar operation, and described in Chapter 5 section 5.2, can also be used to simulate bipolar operation with a variety of waveforms. Operation with three types of current waveform will be investigated in this section: Bipolar two phase on, 120° conduction - this will be referred to as 'trapezoidal' operation. Bipolar three phase on, 180° conduction - 'squarewave' operation. And with sinusoidal phase currents - 'sinusoidal' operation. In each case low speed operation is assumed, so ideal current sources are used to model the phase current controller. Peak phase current is $8.0A$, and the phase angle θ_{on} is selected so as to give maximum average torque.

6.2.1 Trapezoidal Operation

As can be seen from Figure 6.6 the fully-pitched winding and equivalent single-tooth winding phase currents have the same waveshape. The conduction period during each half cycle is 120° , corresponding to a rotation of $2/3$ of a rotor tooth pitch. As excitation of one pair of stator and rotor teeth can produce positive torque for at most $1/2$ a tooth pitch, there must be a period of negative torque production. This is evident in the bottom trace of Figure 6.6. Note that to maximise torque most of the negative torque producing period is near the low inductance unaligned position, as the torque per ampere is lower at this position. The negative torque contributions, although lowering the total torque, have the beneficial effect of reducing the torque ripple. From the fully-pitched winding flux-linkage waveform, ψ_a , it is evident that the applied voltage must change sign ten times per cycle, which has implications for high speed operation.

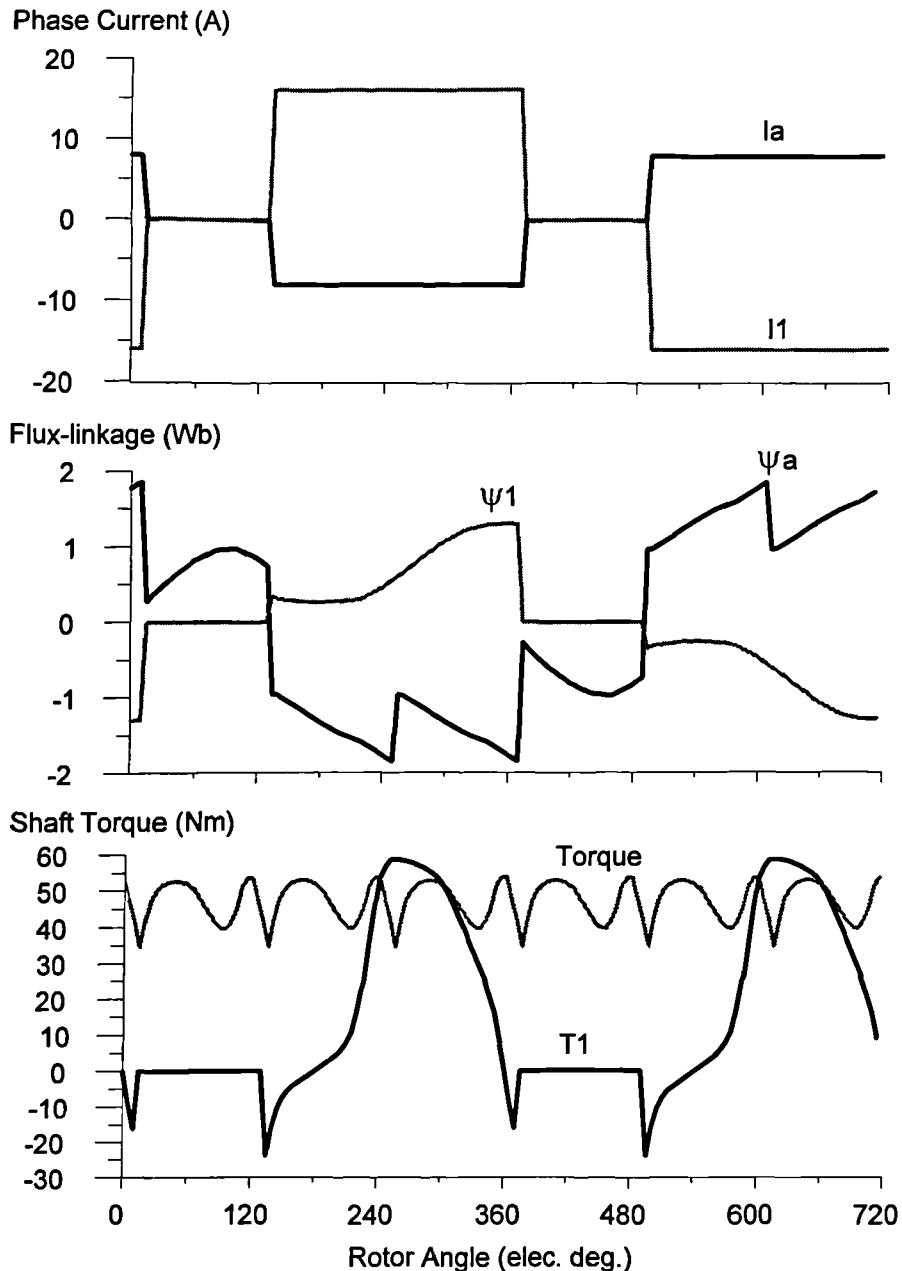


Figure 6.6 Simulated real and equivalent single-tooth winding parameters for trapezoidal phase currents.

6.2.2 Squarewave Operation

By extending the conduction angle from 120° to 180° during each half cycle, the resulting current waveform is a bipolar squarewave. The equivalent single-tooth winding phase current, shown in Figure 6.7, is continuous and has a stepped waveshape. With 8A fully-pitched winding currents, the corresponding single-tooth winding current is stepped with amplitudes of 8A and 24A. Although the single-tooth winding current has a complex waveshape, the fully-pitched winding flux-linkage waveform is less complex than for trapezoidal operation. From the equivalent single-tooth winding flux-

linkage ψ_1 and torque T_1 , shown in Figure 6.7, it is evident the flux-linkage is relatively constant, and the torque low, for just less than one third of a half cycle. This leaves two further segments, each about one third of a half cycle long, during which there are large positive and negative torque contributions. Peak positive and negative torque values are $+100.5\text{Nm}$ and -17.3Nm , resulting in a peak total torque of 83.3Nm . Note that although the negative torque contribution reduces the peak torque (as in Figure 6.6) the shape of the individual torque contributions does not reduce the torque ripple, indeed it increases it.

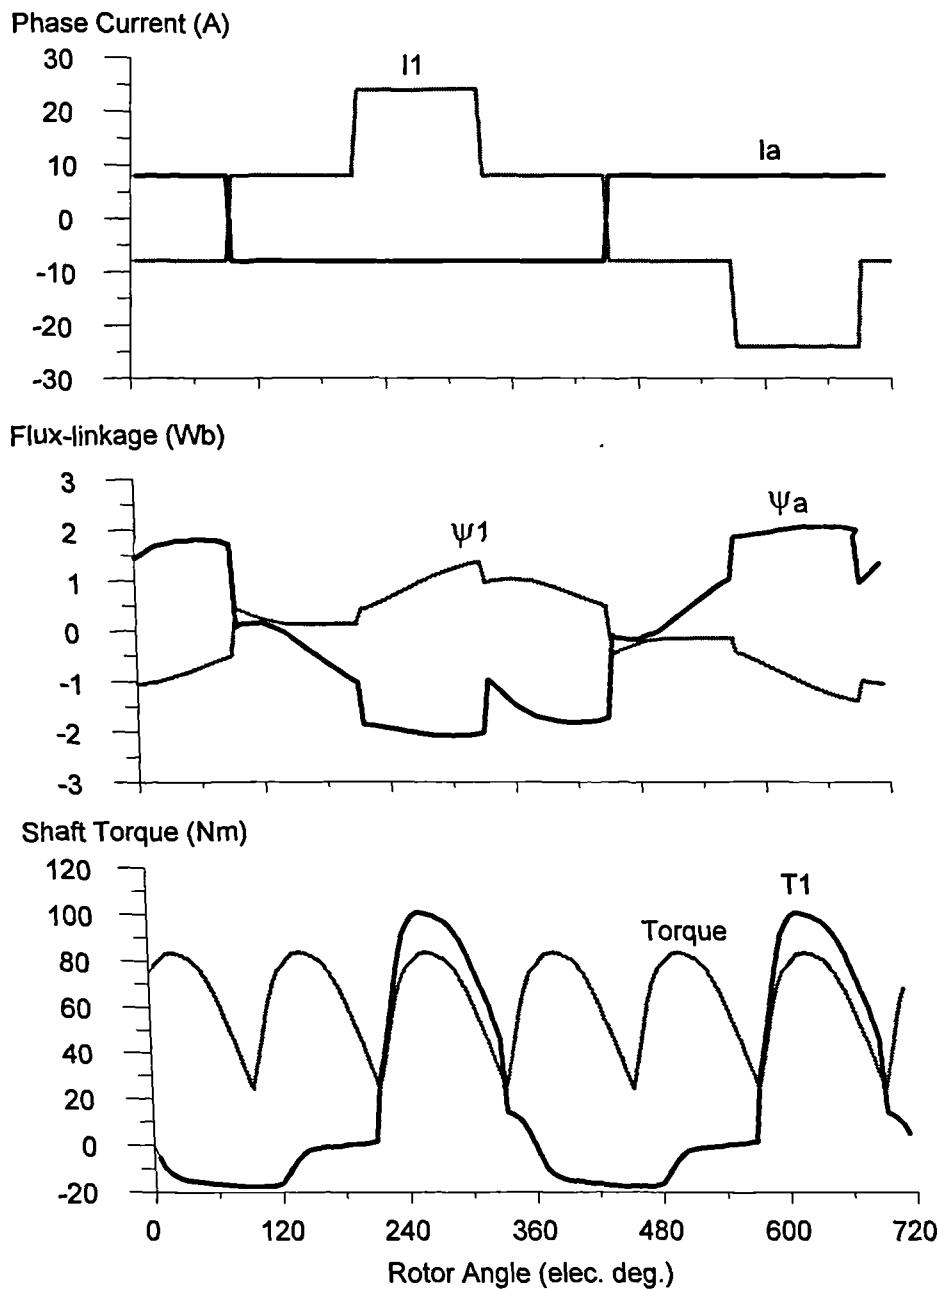


Figure 6.7 Simulated real and equivalent single-tooth winding parameters for squarewave phase currents.

6.2.3 Sinusoidal Operation

Figure 6.8 illustrates various parameters for operation with sinusoidal phase currents. As could be easily predicted, the equivalent single-tooth winding phase currents are also sinusoidal. The smooth and continuous nature of the phase currents leads to a flux-linkage profile with changes much more slowly than for previous bipolar modes. In theory this will enable the current controller to track the current reference to a higher speed without significant errors. The individual torque contributions add such that the peak torque and torque ripple are reduced, as with trapezoidal operation.

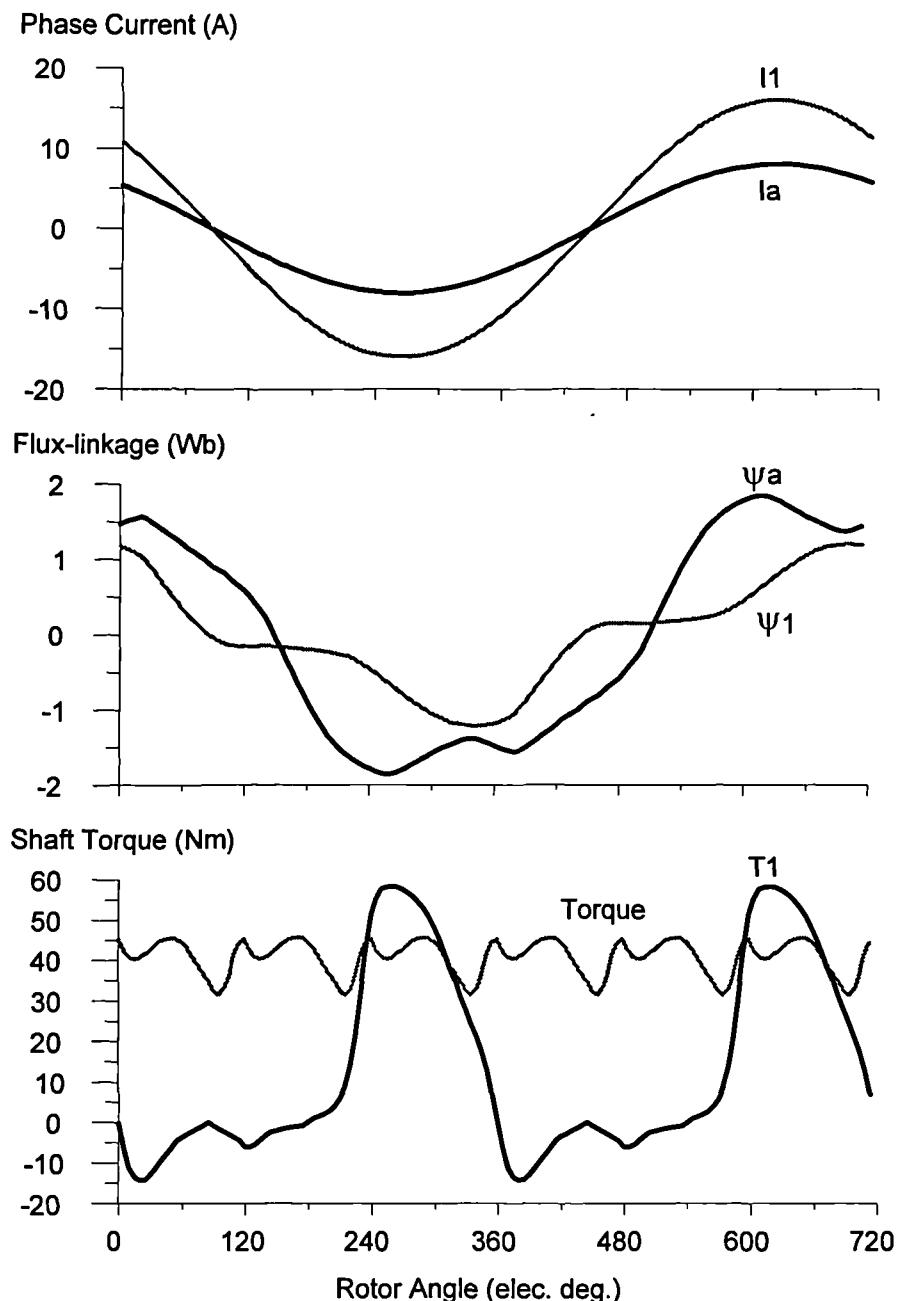


Figure 6.8 Simulated real and equivalent single-tooth winding parameters for sinusoidal phase currents.

6.2.4 Comparison of Simulated Low Speed Performance

From the preceding simulations it is evident that the bipolar modes of operation investigated have slot MMFs which correspond to having single-tooth windings excited for more than half a rotor tooth pitch in duration. Correct synchronisation of MMF to rotor position is required in order to minimise the negative torque contributions which must arise with these long excitation periods. The manner in which these opposing torque contributions sum together makes a considerable difference, to both the average torque and torque ripple. Table 6.1 compares the torque performance of the three bipolar modes investigated thus far, and unipolar operation.

Waveform	Peak Current (A)	RMS Current (A)	Average Torque (Nm)	Torque Ripple (%)	Nm per RMS Ampere
Unipolar	8	6.53	47.1	77	7.21
Trapezoidal	8	6.53	47.1	40	7.21
Squarewave	8	8	63.8	93	7.98
Sinusoidal	8	5.65	40.6	35	7.19
Unipolar	9.8	8	61.5	79	7.68
Trapezoidal	9.8	8	61.6	35	7.70
Squarewave	8	8	63.8	93	7.98
Sinusoidal	11.3	8	63.5	31	7.93

Table 6.1 Comparison of torque performance for simulated unipolar and bipolar phase currents.

Two sets of results are given, one for phase currents of 8A peak, and one for phase currents of 8A RMS. Looking first at the top half of Table 6.1, which is for 8A peak, it can be seen that the torque ripple for trapezoidal and sinusoidal operation is approximately half (or less) than that of the other two modes. Torque per RMS ampere of phase current is almost the same for all modes, except squarewave, which is significantly greater. Comparing figures for equal RMS currents of 8A (the bottom half of Table 6.1) it is clear that for equal machine copper losses the average torque in each case is almost the same, and hence so too is the torque per ampere. Squarewave

operation still has a small advantage in this respect. At the higher average torque level torque ripple for trapezoidal and sinusoidal operation have both decreased, whilst that for unipolar operation has increased. To conclude, for equal machine copper loss the average torque is almost independent of the mode of operation, but the torque ripple for trapezoidal and sinusoidal operation is very much lower than either unipolar or squarewave operation.

6.3 Drive Operation with Separately Connected Phases

In this section measured results for bipolar operation of the experimental drive and fully-pitched winding switched reluctance machine will be presented. For these results the power converter is configured to have a full 'H' bridge per phase, with each phase separately connected. Control software is configured for two processor operation, with the third processor providing real time monitoring only. The phase current controller is an adaptive PID controller, with phase current references as shown in Figure 6.5. Further details of this processor configuration and current control method can be found in Chapters 3 and 4, sections 3.1.3 and 4.1.3, and Appendix D.1.

Operation with three current reference waveforms will be considered for most of the results; trapezoidal, squarewave, and sinusoidal conduction. Low speed operation where the current controller is able to track the reference for most of the period, and high speed operation where the drive is mostly in voltage control, will be presented.

6.3.1 Low Speed Operation

Measured waveforms in this section are for the experimental fully-pitched winding switched reluctance drive, using bipolar current references and a dc-link voltage of 420V, and were obtained with the switched reluctance machine motoring.

Trapezoidal Currents

The waveforms in this section are for phase current demands (I_{dem}) of 10A peak, with the conduction angle in each half cycle set to 120°. At the optimum advance angle average torque is 56.7Nm at a speed of 93rpm.

Experimental waveforms are shown in Figure 6.9, from which it is clear that the actual phase currents track the reference with little error. Trace two in Figure 6.9, although very similar to the simulated waveform in Figure 6.6, has a small disturbance in the middle of each conduction period. This is because the slot MMFs used to calculate the equivalent single-tooth winding currents are not constant as the actual currents switch

due to their finite rise and fall times. The effect of this current disturbance is evident in the third trace of Figure 6.9, which shows the torque contribution calculated from the single-tooth winding current. Although the current disturbance adds harmonic content to the total torque, it does not increase the peak to peak ripple, which is relatively low at 40% peak to peak.

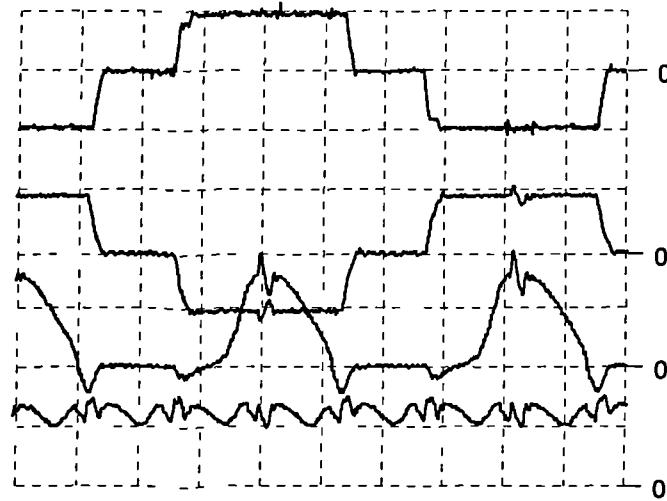


Figure 6.9 Low speed current and torque for trapezoidal conduction. Trace 1 - Phase current I_a , 10A/div. Trace 2 - Equivalent single-tooth winding current I_1 , 20A/div. Trace 3 - Single phase torque T_1 , 50Nm/div. Trace 4 - Total torque 50Nm/div. Horizontal 20ms/div.

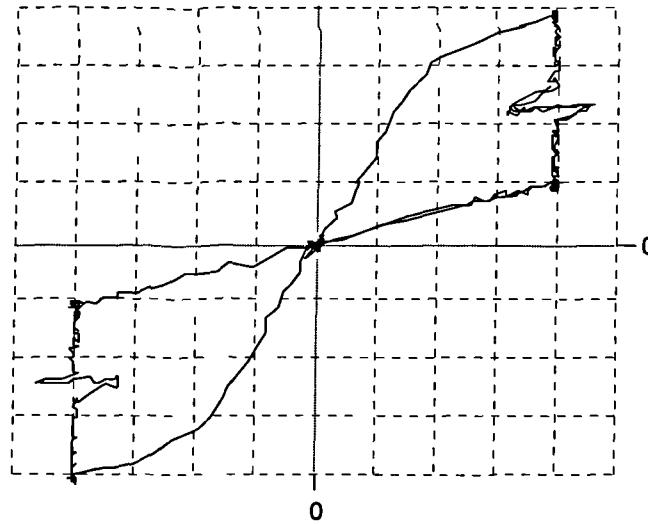


Figure 6.10 Low speed equivalent single-tooth winding flux-linkage/current locus for trapezoidal conduction. Horizontal - Current 5A/div. Vertical - Flux-linkage 0.25Wb/div.

Figure 6.10 shows the equivalent single-tooth winding flux-linkage/current locus. Each half of the locus has a shape very similar to that for unipolar operation with 240°

conduction at low speed. Note that the total area enclosed is approximately twice that for unipolar operation as there are two lobes. However there are only four cycles per revolution for bipolar operation, compared to eight for unipolar operation so the rate of energy conversion is similar in each case.

Squarewave Currents

The waveforms in this section are for phase current demands (I_{dem}) of 10A peak, with the conduction angle in each half cycle set to 180° so the phase current is a squarewave. At the optimum advance angle average torque is 72.8Nm at a speed of 112rpm.

The top two traces in Figure 6.11 show the actual phase current, and the equivalent single-tooth winding current calculated by the controller. With real phase currents of 10A peak, the equivalent single-tooth winding current is stepped with amplitudes of 10A and 30A. As can be seen from the third trace in Figure 6.11, the torque contribution for an equivalent current of 30A peaks at a substantial 130Nm. Negative torque contributions, due to the continuous phase current, reduce the peak total torque to approximately 115Nm. Torque ripple for this mode of operation is high, 93% peak to peak, much greater than for operation with trapezoidal conduction.



Figure 6.11 Low speed current and torque for squarewave conduction. Trace 1 - Phase current I_a , 10A/div. Trace 2 - Equivalent single-tooth winding current I_1 , 20A/div. Trace 3 - Single phase torque T_1 , 50Nm/div. Trace 4 - Total torque 50Nm/div. Horizontal 20ms/div.

The equivalent single-tooth winding flux-linkage/current locus, shown in Figure 6.12, has a 'dumbbell' shape, which is quite different from other modes of operation. Note that most of the area enclosed by the locus (and hence torque produced) resides in an area where the machine is saturated.

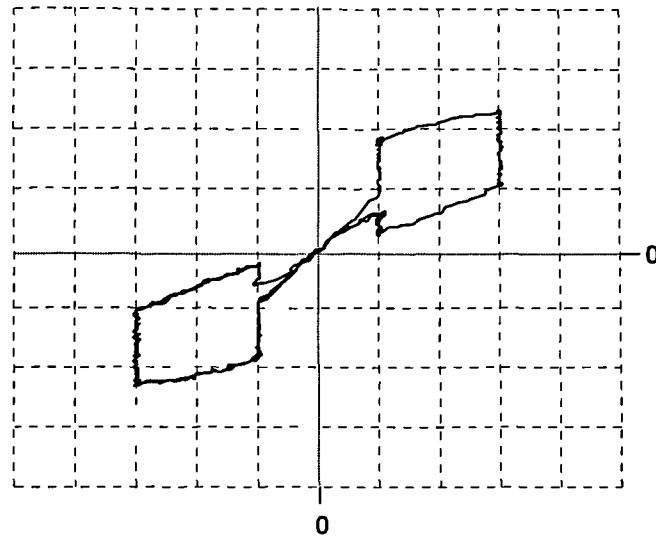


Figure 6.12 Low speed equivalent single-tooth winding flux-linkage/current locus for squarewave conduction. Horizontal - Current 10A/div. Vertical - Flux-linkage 0.5Wb/div.

Sinusoidal Currents

The waveforms in this section are for phase current demands (*Idem*) of 14.1A peak, an RMS current of 10A. At the optimum advance angle average torque is 70.5Nm at a speed of 108rpm.

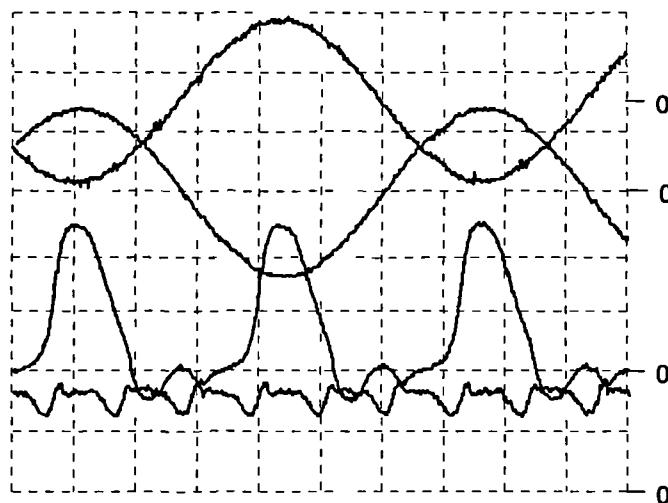


Figure 6.13 Low speed current and torque for sinusoidal phase currents. Trace 1 - Phase current I_a , 10A/div. Trace 2 - Equivalent single-tooth winding current I_1 , 20A/div. Trace 3 - Single phase torque T_1 , 50Nm/div. Trace 4 - Total torque 50Nm/div. Horizontal 20ms/div.

Actual and equivalent single-tooth winding phase currents are shown in the top two traces of Figure 6.13. The low di/dt allows the current controller to track the reference

closely at all times. With a phase current of 10A RMS the equivalent single-tooth winding current is 20A RMS, 28.3A peak. As with squarewave phase currents this large current produces substantial positive torque contributions. Sinusoidal operation has an advantage over square wave operation, in that there is reduced excitation of each stator tooth during periods when it produces negative torque. This is one of the factors that results in sinusoidal operation having a peak to peak torque ripple one third of that for squarewave operation. For the waveforms shown in Figure 6.13, torque ripple is 37% peak to peak.

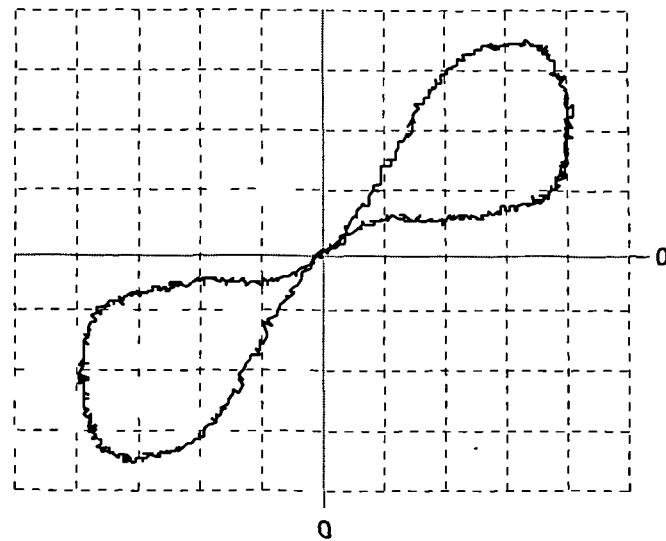


Figure 6.14 Low speed equivalent single-tooth winding flux-linkage/current locus for sinusoidal phase current. Horizontal - Current 5A/div. Vertical - Flux-linkage 0.25Wb/div.

Figure 6.14 is the equivalent single-tooth winding flux-linkage/current locus for sinusoidal phase currents of 10A peak. From the shape of the locus it is evident that for a given peak current the area enclosed within the locus is less than the theoretical maximum, and that of trapezoidal operation. The low rate of change of flux for teeth approaching the aligned position, absence of harmonics in the phase current waveform and reduced torque ripple, results in substantially reduced acoustic noise compared to many other modes of operation.

6.3.2 High Speed Operation

In general the high speed torque produced by bipolar modes of operation is much less than for unipolar operation, given equal phase current reference magnitudes. This is not a fundamental feature of bipolar operation, but is due mainly to a combination of the current reference waveform and PID current controller.

Trapezoidal and Squarewave Operation

In order to maximise torque output at high speed, the 120° conduction width of trapezoidal excitation needs to be extended towards 180° . The mode of operation has therefore become identical to squarewave operation. At a sufficiently high speed the actual phase current never reaches the demand value and so the applied voltage follows the current reference squarewave. This mode of operation is illustrated in Figure 6.15, which shows operation at 1250rpm with an average shaft torque of 12.3Nm. Trace one shows the phase voltage and trace two the corresponding equivalent single-tooth winding flux-linkage. These waveforms have the same shape as for unipolar operation under voltage control, but with two important differences that have consequences for current waveforms and torque production. Firstly, the frequency of the applied voltage is half that of unipolar operation, so a single cycle lasts for a rotation of two rotor tooth pitches. Secondly, the flux waveform is symmetrically offset about zero, rather than being unipolar. As the flux-linkage is continuous, so too is the equivalent single-tooth winding phase current shown in trace three of Figure 6.15. Each flux half cycle lasts for a whole rotor tooth rotation resulting in several undesirable features:- a large peak current near the unaligned position, high RMS current, and large negative torque contributions.



Figure 6.15 High speed current, voltage and flux for bipolar conduction, ≈ 12 Nm. Trace 1 - Phase voltage V_a , 1kV/div. Trace 2 - Equivalent single-tooth winding flux ψ_1 , 1.25Wb/div. Trace 3 - Equivalent single-tooth winding current I_1 , 20A/div. Trace 4 - Phase current I_a , 10A/div. Horizontal 2.5ms/div.

If the equivalent single-tooth winding flux-linkage was forced to be like that of unipolar operation, but alternating sign each half cycle, then it should be possible to produce as much torque at high speed as with unipolar operation. Figure 6.16 shows how the fully-pitched winding phase voltage required to achieve this is derived.

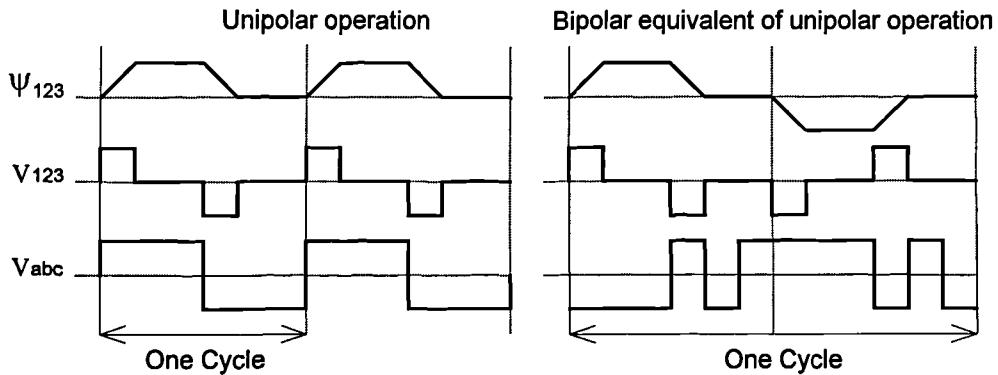


Figure 6.16 Derivation of bipolar voltage control waveform from unipolar voltage control.

With a squarewave reference the controller only produces the required voltage waveform under conditions that give very low torque (i.e. near zero). By replacing the bipolar current reference shown Figure 6.5, with a *current* reference having the desired *voltage* waveshape, as illustrated in Figure 6.16, then at high speed the actual phase voltage will follow the desired shape.

This mode of operation is shown in Figure 6.17, producing 15.6Nm at 1250rpm. It is evident from traces two and three, which show equivalent single-tooth winding current and flux-linkage, that the modulus of these waveforms is the same as for unipolar operation under voltage control. Equivalent single-tooth winding flux-linkage is now zero for 1/3 of the time, which reduces RMS phase current from 7.15A to 5.13A, and also reduces the negative torque contributions, increasing shaft torque from 12.3Nm to 15.6Nm.

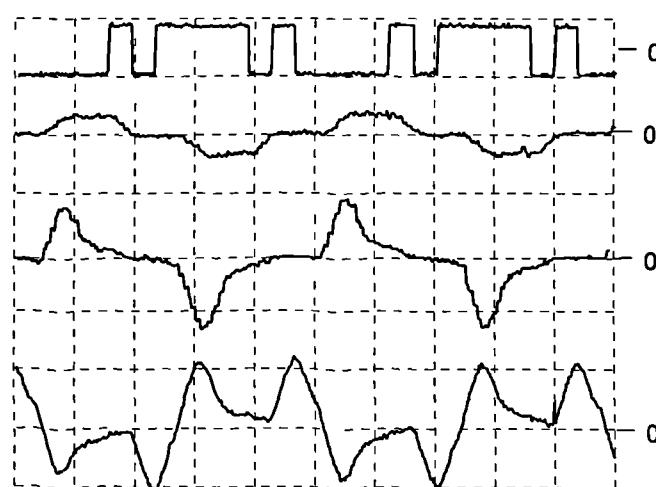


Figure 6.17 High speed current, voltage and flux for bipolar conduction, ≈ 15 Nm. Trace 1 - Phase voltage V_a , 1kV/div. Trace 2 - Equivalent single-tooth winding flux ψ_1 , 1.25Wb/div. Trace 3 - Equivalent single-tooth winding current I_1 , 20A/div. Trace 4 - Phase current I_a , 10A/div. Horizontal 2.5ms/div.

Average shaft torque is approximately the same as for unipolar operation with this excitation pattern. Thus by application of the appropriate voltage it is possible for bipolar operation to match the output of unipolar operation at high speed. However, this requires a different reference waveshape at high speed from that used at low speed, which complicates control.

Sinusoidal Operation

Although the flux-linkage waveform for sinusoidal operation is not stepped, there are periods when the rate of change of flux is much greater, as can be seen from Figure 6.8. As speed increases the actual flux is no longer able to keep up with the reference at these times, and hence the current falls. An unfortunate side effect of sinusoidal operation is that the main torque producing periods coincide with these periods and the current shortfall greatly reduces the torque produced.

6.3.3 Operation with Novel Current Waveforms

The idea of calculating the required fully-pitched winding waveforms from desired single-tooth winding waveforms, as used in the previous section to derive voltages, can also be applied to current waveforms. For a conventional single-tooth winding switched reluctance machine the phase current direction does not affect torque production. Operation would therefore be the same if alternate current pulses had different polarities. Figure 6.18 shows current reference waveforms for a three phase single-tooth winding switched reluctance machine with the sign alternating each rotor tooth pitch, I_{123} , and the equivalent fully-pitched winding phase current I_{abc} . Each cycle now lasts for two rotor tooth pitches, and so the current pulses last for 60° (electrical) during each half cycle.

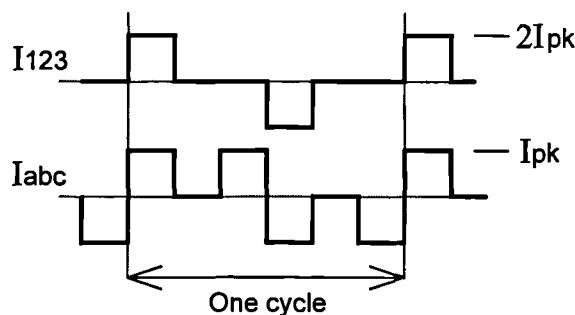


Figure 6.18 Waveform derived from bipolar 120° single-tooth winding current.

This waveform has the interesting property that the single-tooth winding conduction interval can be made greater than $1/3$ of a tooth pitch in duration, without increasing the peak fully-pitched winding current. This is in contrast to unipolar operation, where the

single-tooth winding conduction width cannot be extended without doubling the peak fully-pitched winding peak current. Furthermore, the fully-pitched winding duty cycle is 2/3, so the average and RMS phase currents are the same as for unipolar operation.

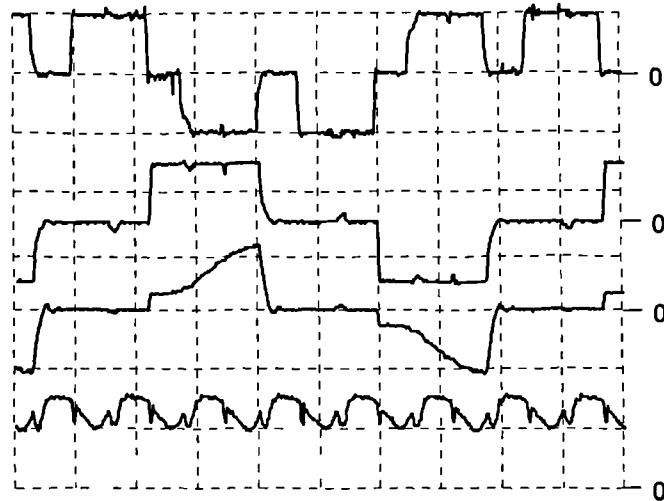


Figure 6.19 Low speed current, flux and torque, for bipolar 180° conduction equivalent single-tooth winding current. Trace 1 - Phase current I_a , 10A/div. Trace 2 - Equivalent single-tooth winding current I_1 , 20A/div. Trace 3 - Equivalent single-tooth winding flux ψ_1 , 1.25Wb/div. Trace 4 - Total torque 50Nm/div. Horizontal 20ms/div.

In Figure 6.19 operation with the single-tooth winding current adjusted for a conduction period of 1/2 a tooth pitch is illustrated. Speed is 97rpm with a peak phase current of 10A producing 60.4Nm. As the equivalent single-tooth winding current pulses have a duration of 1/2 a tooth pitch during each half cycle, torque contributions from successive phases overlap, increasing the average torque and reducing torque ripple.

6.4 Torque/Speed Characteristics

The torque/speed characteristics presented in this section are measured results for the experimental fully-pitched winding switched reluctance drive, with various bipolar excitation sequences. Phase current is controlled by the adaptive PID controller previously described. The phase current reference has a fixed basic shape and peak current throughout the speed range. At each operating point the 'on', and where appropriate 'conduction width', angles are adjusted to give the maximum torque output.

Figure 6.20 shows the torque/speed characteristics for bipolar operation with trapezoidal, squarewave, and sinusoidal phase currents, unipolar operation also shown for comparison. In each case the peak phase current is 10A. At low speed squarewave conduction gives significantly more torque, and sinusoidal operation less torque, than the other two modes. However, this can to some extent be accounted for by differences

in average and RMS phase current. For all the bipolar modes torque falls much more rapidly with speed than for unipolar operation, which is due to the mode of operation rather than a fundamental rule of bipolar operation (see section 6.3.2).

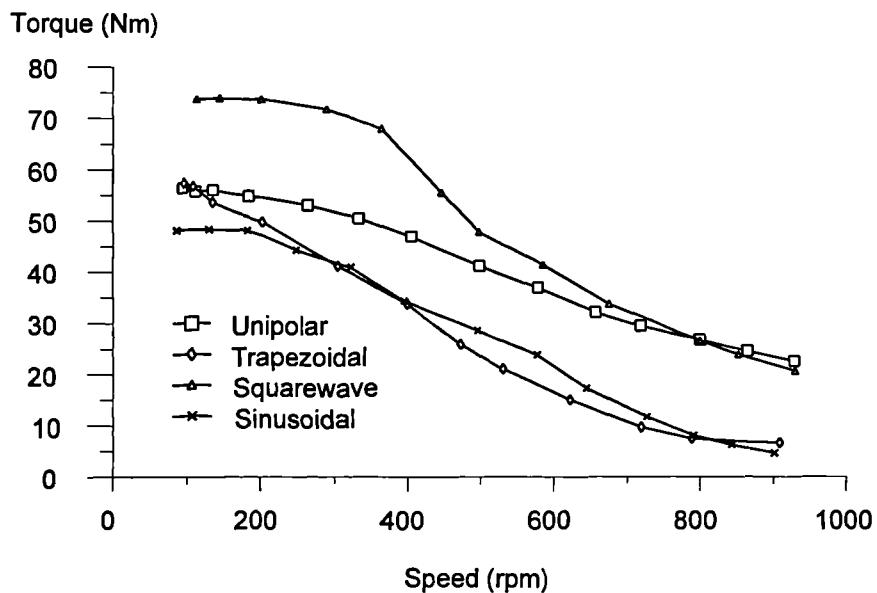


Figure 6.20 Torque/speed characteristics for 10A peak phase currents.

With equal RMS phase currents the low speed average torque is approximately the same for each waveform, as shown in Figure 6.21. With equal RMS phase currents, at low speed the machine copper losses are the same for each mode, and so the torque per unit loss in each case is very similar. This is despite of the very different phase current waveshapes and torque contributions from each pair of stator and rotor teeth.

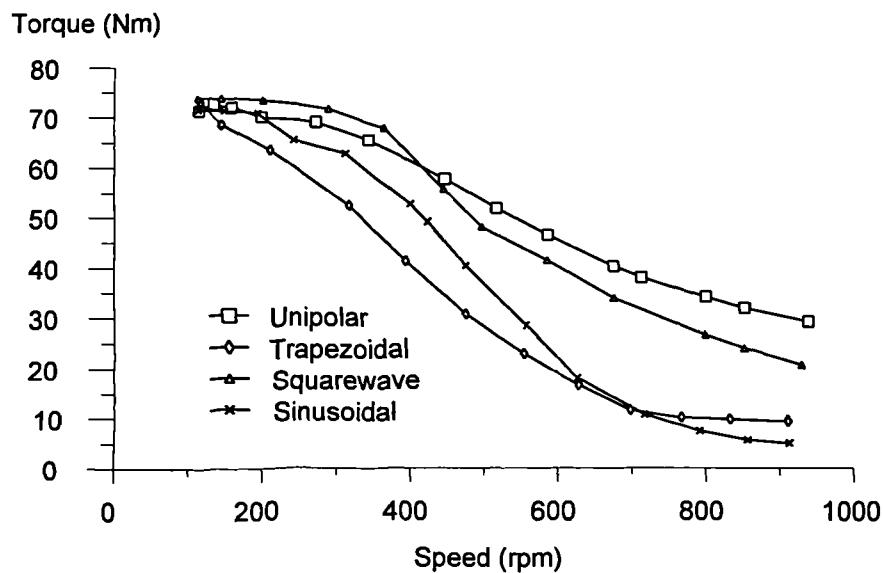


Figure 6.21 Torque/speed characteristics for 10A RMS phase currents.

Figure 6.22 shows the shaft power characteristics for operation with 10A RMS phase currents. Note that whereas unipolar operation tends towards constant power at high speed, all the bipolar modes fall with speed after an initial peak. For trapezoidal and sinusoidal operation the fall off of shaft power at high speed is quite dramatic.

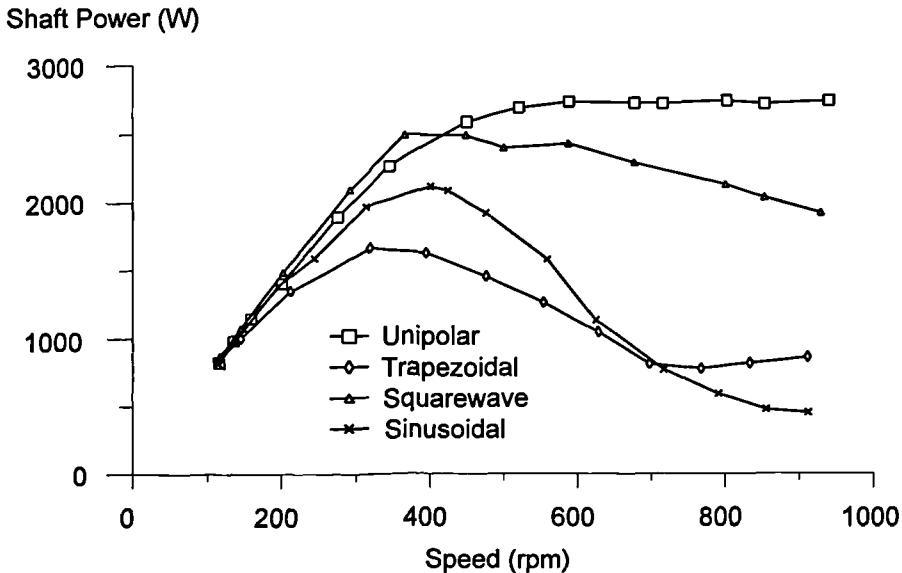


Figure 6.22 Shaft power characteristics for 10A RMS phase currents.

It should be noted that the poor high speed torque performance of the bipolar modes shown is due mainly to the current reference and controller used. By changing the current reference waveshape at high speed it is possible to produce as much shaft power as unipolar operation, as established in section 6.3.2. However, all bipolar modes require twice as many switching devices as unipolar operation, which implies large penalty in terms of converter kVA/kW.

6.5 Star Connected Operation

Results in this section are for operation with the three phases of the experimental fully-pitched winding machine star connected. A three wire configuration is used with the star point floating. This arrangement forces the sum of the three phase currents to be zero which means that it is not possible to have squarewave phase currents, hence only operation with trapezoidal and sinusoidal phase currents is considered here.

6.5.1 Experimental Drive Operation

Trapezoidal operation with star connected phases is very similar to that with a 'H' bridge per phase at low speed. Figure 6.23 shows operation with 10A peak phase currents at 93rpm producing 55.1Nm. Comparison of current and torque waveforms with those in

Figure 6.9 shows they are broadly the same. There are however small differences, these show up most in the second and third traces which show equivalent single-tooth winding current and torque.

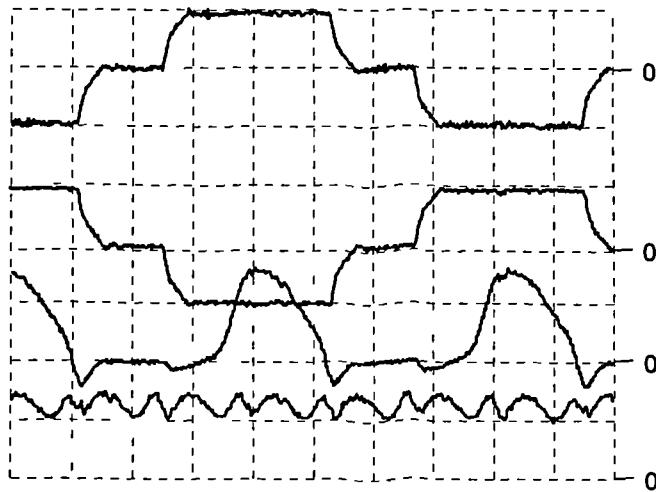


Figure 6.23 Low speed current and torque for star connected trapezoidal conduction. Trace 1 - Phase current I_a , 10A/div. Trace 2 - Equivalent single-tooth winding current I_1 , 20A/div. Trace 3 - Single phase torque T_1 , 50Nm/div. Trace 4 - Total torque 50Nm/div. Horizontal 20ms/div.

Looking first at the equivalent single-tooth winding current in Figure 6.23, during each half cycle the current remains constant throughout the conduction interval. By comparison, in Figure 6.9 at the middle of each conduction interval there is a small discontinuity and corresponding torque perturbation. This discontinuity occurs because the conduction angle is not exactly 120° due to current rise and fall times. When star connected the sum of the three actual phase currents is forced to be zero, this equalises current rise and fall times thus producing a perturbation free equivalent single-tooth winding current. The overall effect of this can be seen by comparing the total shaft torque in Figure 6.23 and Figure 6.9. With star connection the peak to peak ripple is only reduced slightly, to 35%, but there is a notable reduction of torque ripple harmonics.

Low speed operation with sinusoidal phase currents is the same whether star connected or with a 'H' bridge per phase. However it is interesting to note what happens at high speed. Figure 6.24 shows operation at 730rpm with sinusoidal phase current references of 10A RMS. Clearly the actual phase current is very distorted and much lower in amplitude than the reference. Converter terminal output voltage, measured with respect to the dc-link negative voltage rail, is a squarewave. This gives a line to line voltage as illustrated in trace two. Trace three is the star point voltage, measured with respect to the dc-link negative voltage rail. With a dc-link voltage of 420V the star point voltage

goes approximately 175V above and below the dc-link rails. The resulting phase voltage shown in trace four, has a peak to peak amplitude of almost three times the dc-link voltage.

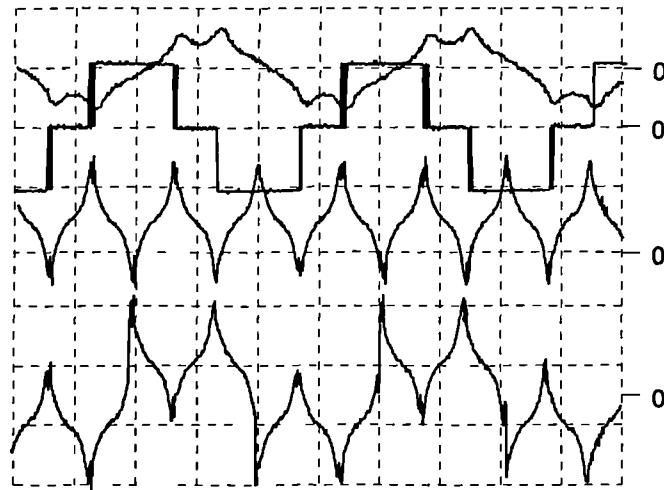


Figure 6.24 High speed current and voltage for star connection and sinusoidal phase currents. Trace 1 - Phase current I_a , 10A/div. Trace 2 - Line to line voltage 400V/div. Trace 3 - Star point voltage, 400V/div. Trace 4 - Phase voltage, 400V/div. Horizontal 5ms/div.

6.5.2 Torque/speed Characteristics

The torque/speed characteristics shown in Figure 6.25 are for 10A peak trapezoidal conduction with star connected phases, and 10A RMS sinusoidal phase currents with star and 'H' bridge connection.

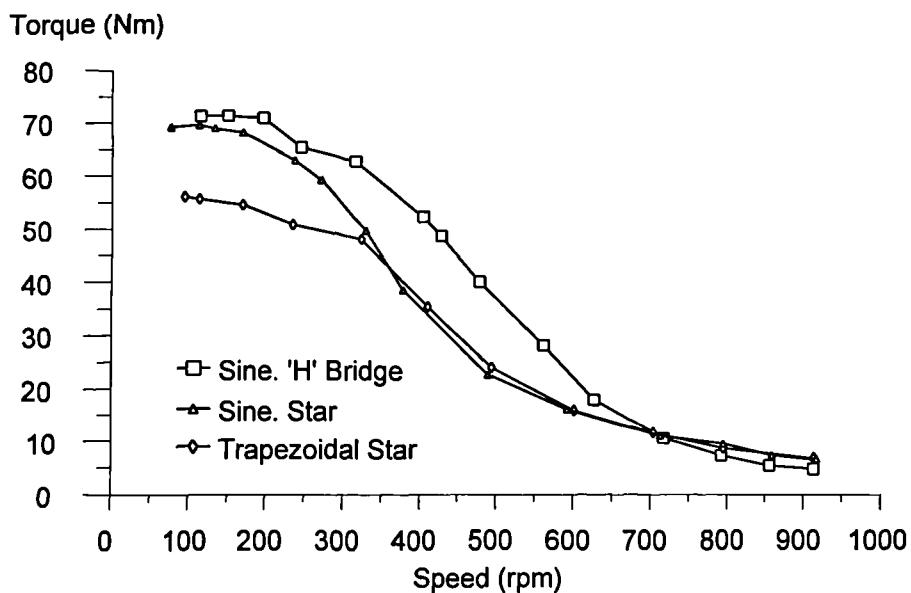


Figure 6.25 Torque/speed characteristics for star connected phases.

By comparison of the two sinusoidal characteristics it can be seen that the low speed torque is almost the same, but the torque drops off more rapidly with speed when star connected, until at high speed star connection gives slightly more torque. The reason torque drops off more rapidly with speed when star connected is that each phase voltage cannot be individually controlled, and as the phase currents become distorted the star connection imposes conditions not present with 'H' bridge connection. At high speed it is postulated that the higher torque with star connection is due to the phase voltage exceeding the dc-link voltage, a condition clearly not possible with 'H' bridge connection.

6.6 Delta Connected Operation

With the three phases connected in a delta configuration individual phase currents can no longer be controlled as only line currents are measured by the control system. With a line to line voltage that results in sinusoidal line currents, the phase currents have a large third harmonic component which results in a severely distorted phase current, as shown in Figure 6.26. Harmonic analysis of the phase current waveform shows that the third harmonic has an amplitude approximately 0.75 times that of the fundamental. It is not possible to have sinusoidal phase currents in the delta connected experimental fully-pitched winding switched reluctance machine, because the phase voltages required violate the fundamental rule of delta connection, that the sum of the phase voltages must be zero. With line currents of 10A peak the average shaft torque at speeds up to 200rpm is approximately 45Nm.

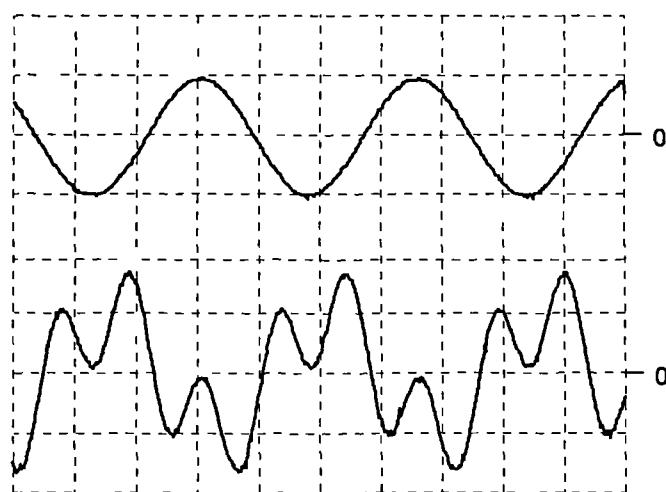


Figure 6.26 Low speed operation with delta connected phases and sinusoidal line currents. Trace 1 - Line current, 10A/div. Trace 2 - Phase current 10A/div. Horizontal 20ms/div.

6.7 Summary and Conclusions

Analysis based on idealised inductance profiles has shown that with trapezoidal conduction sequences positive torque is produced from periods of falling mutual inductance only, and for squarewave excitation from periods of rising and falling mutual inductance. With bipolar excitation one electrical cycle corresponds to a rotation of two rotor tooth pitches.

Simulation studies of operation with trapezoidal, squarewave and sinusoidal phase currents has shown that pairs of stator and rotor teeth are excited for a rotation of more than half a rotor tooth pitch at a time. Thus at times pairs of teeth are producing negative torque contributions, a feature which cannot be seen from the ideal inductance profiles.

Experimental results for low speed operation correspond very well with simulated results, and confirm that there are significant differences in torque ripple. Sinusoidal and trapezoidal excitation have approximately half the torque ripple of squarewave operation.

Torque/speed characteristics show that at low speed the torque per RMS ampere of phase current is almost the same for trapezoidal, squarewave, sinusoidal and unipolar phase currents. At high speed all three bipolar modes produce substantially less torque than unipolar operation. However, it has been shown that this is due to the current reference and controller. *With a suitable reference waveform bipolar operation can equal unipolar power output, but this requires a more complex controller as the fundamental reference waveshape must change as speed increases.*

It has been shown that operation with star connected phases is possible with trapezoidal and sinusoidal phase currents, however delta connection gives very unsatisfactory results due to a large third harmonic circulating current.

Although in theory bipolar squarewave operation gives substantially more torque for the same peak phase current as unipolar operation, in practice this advantage does not materialise. At low speed the torque per RMS ampere is the same, so there is no advantage in machine losses, and at high speed torque capability is no greater. Thus in terms of power converter kVA/kW rating, bipolar operation is at a significant disadvantage because twice as many power switches are required, unless the machine is star connected.

Chapter 7

Flux Control of Switched Reluctance Machines

7.1 Theory of operation

As phase current is often the controlled variable in electrical machine drives it seems natural to use closed loop current controllers. However, the highly non-linear nature of switched reluctance machines makes the optimisation of closed loop current controllers difficult. For voltage source converters the directly controlled variable applied to the machine is voltage, with the current being a function of the flux-linkage and rotor position. Furthermore, if the resistive voltage drop is neglected, flux-linkage is linearly proportional to the applied volt-seconds. These considerations suggest that it may be more appropriate to design a controller based on flux-linkage rather than current.

7.1.1 Phase Current Control Methods

In most machine drives phase current is the controlled variable as shaft torque is proportional to phase current or phase current squared, and current transducers are readily available. However the phase winding of a switched reluctance machine presents a very non-linear load to a current controller because the effective inductance of the phase varies, with both rotor position and phase current. This has adverse implications for the two most popular types of current controller used in switched reluctance drives; hysteresis or bang-bang controllers, and PI or PID controllers with fixed frequency pulse width modulated (PWM) output. Hysteresis controllers are simple to implement, unconditionally stable and give a fast response to step changes. However, with a constant hysteresis band, the switching frequency varies inversely with load time constant, making the trade-off of switching losses and current error band more difficult to optimise. Limits on maximum switching frequency and system delays, both place limits on the minimum current ripple which can be achieved. Although constant switching frequency PID controllers make calculations of switching loss easier, it is very difficult to optimise their time domain response. For a simple non-adaptive controller, the gains must be set such that the system is stable over the full range of system time constants. Whatever the point at which the response is optimised, the response at other points will tend to be over or under damped. A further complication

with PI and PID controllers is that some form of anti-windup must be incorporated to limit the integral term when the output saturates. If this is not incorporated the response to demand step changes will tend to overshoot.

When using voltage source converters to drive electrical machines the overall controlled variable is usually phase current, although this is actually achieved by controlling the machine phase voltage and flux-linkage. Thus for the machine, the directly controlled variable is flux-linkage rather than current. For machines of a few kW or more, the voltage drop due to phase resistance is negligible, and flux-linkage is simply equal to the applied volt-time integral. The resulting phase current is dependant on the machine type, rotor angular velocity and the instantaneous rotor position. Whereas phase current may be a non-linear function of flux-linkage and time, the flux-linkage is a linear function of voltage and time. Thus as flux-linkage may be controlled directly, and from a control standpoint is not dependant on load non-linearity and time constants, this suggests that superior control performance could be obtained by making flux-linkage the controlled variable.

In order to use a flux-linkage controller to achieve the desired phase currents, some form of transformation must be used to obtain the flux-linkage reference from the phase current reference. For the highly non-linear flux-linkage/current characteristics of a switched reluctance machine, *lookup tables holding machine characteristics are the most appropriate method*.

Previous work by Lovatt [7.1] comparing a hysteresis current controller, PID current controller and flux-linkage controller, showed that the flux controller performed better than the other types, particularly when compared over a wide dynamic range. The experimental system was implemented using analogue electronics and used a stationary non-linear inductor as a load. Flux-linkage feedback was achieved by measuring inductor current and estimating the flux-linkage from a three point piece-wise linear approximation to the inductor characteristics. This too was implemented using analogue circuitry.

Work presented in this chapter extends the principles used by Lovatt by applying flux-linkage control to a rotating switched reluctance machine, rather than a fixed inductor. Lookup tables are used to estimate flux-linkage from measured phase current and rotor position, and an advanced digital controller is used to optimise the response.

7.1.2 A Dead Beat Flux Controller

An ideal controller could be considered to be one that has infinite bandwidth, no delay, zero steady state error, no overshoot and maintains this response irrespective of changes in system parameters. A close approximation to this can be obtained from an auto-tuning discrete time dead-beat controller, with a sufficiently high sample rate. This controller has a bandwidth of half the sample frequency, a delay of one sample interval, and gives a unity gain step response free from overshoot or ringing.

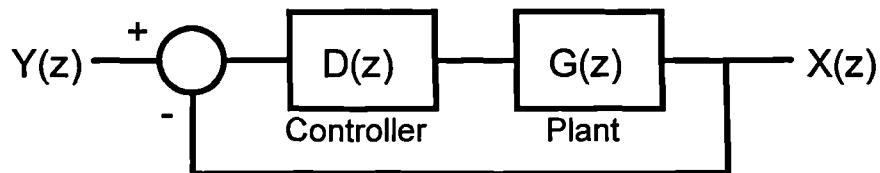


Figure 7.1 Unity feedback system model.

For the system model shown in Figure 7.1 to obtain a dead-beat response the system must satisfy the following equation:-

$$H(z) = \frac{X(z)}{Y(z)} = \frac{D(z)G(z)}{1 + D(z)G(z)} = \frac{1}{z} \quad (7.1)$$

Hence the controller required is given by:-

$$D(z) = \frac{1}{z-1} \cdot \frac{1}{G(z)} \quad (7.2)$$

In this case the plant is an inductive circuit with resistance R , carrying a current i , and with an applied voltage V . Flux-linkage is given by the equation:-

$$\psi(t) = \int_0^t (V - iR) dt \quad (7.3)$$

If the resistive voltage drop is negligible in comparison to the applied voltage, the discrete time plant transfer function for a sample interval of T is:-

$$G(z) = \frac{\psi(z)}{V(z)} = \frac{T}{z-1} \quad (7.4)$$

Hence, for dead-beat response the control transfer function required is as follows:-

$$D(z) = \frac{1}{z-1} \cdot \frac{1}{G(z)} = \frac{1}{z-1} \cdot \frac{z-1}{T} = \frac{1}{T} \quad (7.5)$$

Thus the controller simply has to multiply the flux error by the sample frequency to obtain a dead-beat system response. By preceding the controller with a one step advance for the reference value, the resulting overall system response has no delay because the reference advance compensates for the system delay.

This is the basis for the switched reluctance drive flux controller implemented and shows that theoretically the controller required has a very simple transfer function. However the practical implementation has to compensate for the non-ideal nature of various parts of the system. The main departures from the ideal system described above are that the flux-linkage is not measured directly, and there is a delay in the system from the sampling of variables to the application of voltage to the motor phase.

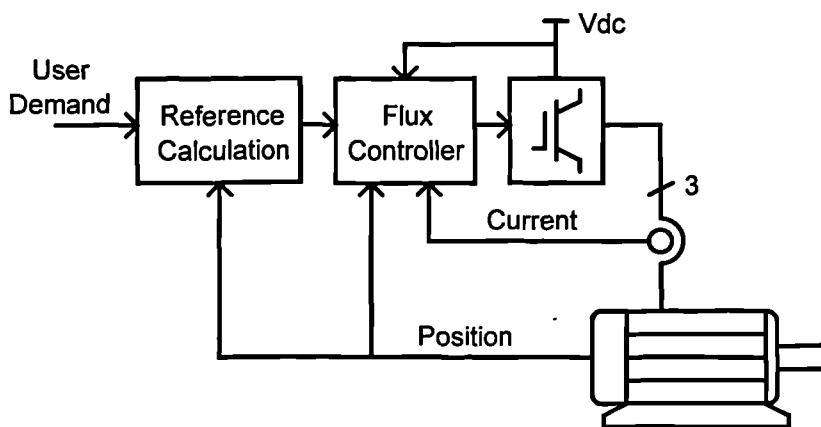


Figure 7.2 Block diagram of the motor flux control system.

Figure 7.2 shows a block diagram of the switched reluctance motor flux control system. In addition to the user demand value the only measured parameters are the motor phase currents, rotor position and the dc-link voltage. Individual phase flux-linkage reference values are calculated from the flux-linkage reference, or current reference and rotor position. These demand values are then passed to the flux controller, which calculates the required power converter switching times in order to achieve the desired motor flux. These processes will now be described in more detail.

Implementation of the switched reluctance drive flux-linkage controller is shown as a transfer function block diagram in Figure 7.3. Operation can be summarised as follows:- A flux-linkage reference value is calculated from the reference waveform, and rotor position with the one step speed position advance added. The controller then compares the flux reference value with the actual flux which is calculated from current and position. From this the required converter duty cycle δ is determined. Voltage applied to the machine windings then gives rise to the current $i(t)$.

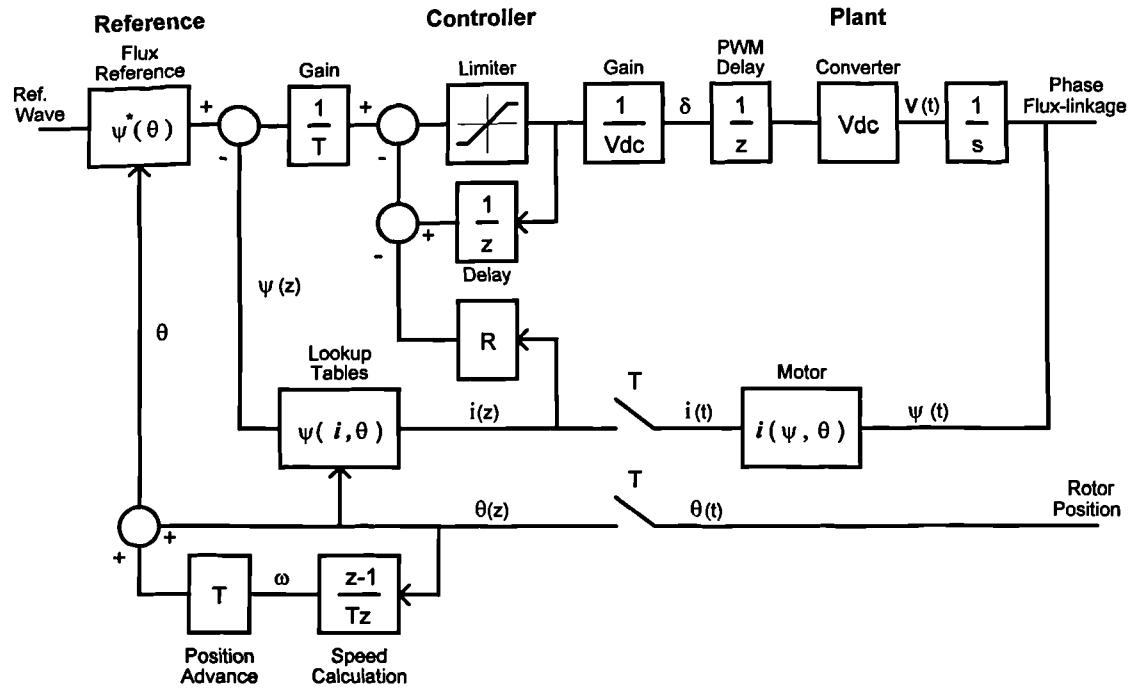


Figure 7.3 Motor flux controller transfer function block diagram.

Flux-linkage reference values for each phase are obtained from lookup tables, or from equations that describe the reference waveshape. Overall system delays are minimised by calculating the flux reference one step ahead. Rotor speed ω is calculated from the rate of change of rotor position, and used to predict the rotor position at the next sample time using the following equation:-

$$\theta_{adv} = \theta_n + \omega \cdot T \quad (7.6)$$

Where T is the sample interval and θ is the rotor position. Function $F(\theta)$ may be implemented as an equation or lookup table, and gives the flux reference as a function of predicted position thus:-

$$\psi_n^* = F(\theta_{adv}) \quad (7.7)$$

Actual phase flux-linkage may be derived directly from integration of the terminal voltage, or indirectly from the phase current and machine characteristics. The voltage integration method has several factors which make it difficult to implement particularly at low speed, and with temperature variations of phase resistance, hence an indirect method is employed. Phase current and rotor position are measured with a sample interval of T , and then the machine flux-linkage/current characteristics are used to determine the flux-linkage. This method requires a high resolution rotor position sensor and does not allow for any differences between static and dynamic flux-linkage

characteristics. The experimental system uses a 10 bit absolute position sensor and measured values for the machine flux-linkage/current characteristics (see Appendix A).

For dead-beat response it was shown in equation 7.5 that the controller needed a transfer function of $1/T$ so the required output voltage is given by:-

$$V_n^* = \frac{\Psi_n^* - \Psi_n}{T} \quad (7.8)$$

For the switching scheme implemented the average output voltage over one period is proportional to the duty cycle δ multiplied by the dc-link voltage (neglecting dead time effects). Duty cycle is thus calculated from:-

$$\delta_n = \frac{V_n^*}{V_{dc}} \quad (7.9)$$

Where V_{dc} is the dc-link voltage. Voltage drop due to the phase winding resistance is compensated for by adding the appropriate voltage to the required output voltage. As shown in Figure 7.3 there is a one sample delay due to the pulse width modulator between the calculation of the required output voltage, and the voltage actually being applied to the phase windings. At the start of the n 'th sample interval current and rotor position are sampled and the actual flux $\psi(n)$ calculated. During this sample interval the applied voltage is $V(n-1)$, thus at the start of the next sample interval during which $V(n)$ is actually applied, the flux will be $\psi(n)$ plus $V(n-1)T$. This is compensated for by subtracting the previous output voltage from the current value. Note that as shown in Figure 7.3, the output voltage is limited to the dc-link voltage before the one sample delay used for previous sample output compensation. A beneficial side effect of this scheme is that if the present output is limited by saturation, this can be compensated for in the following output voltage. This further improves system response under large signal conditions.

In Figure 7.3 the motor is shown as a transfer function block relating phase current to flux-linkage and rotor position. The flux-linkage lookup tables perform the inverse of this function and so the two blocks combined can be considered as a unity gain block. This is shown in Figure 7.4 as the unity gain negative feedback loop. This figure is a simplified model of the flux-linkage control system shown in Figure 7.3. The converter is represented using the discrete time transfer function given in equation 7.4, output limits have been removed, as has the IR compensation, and the speed advance applied to the flux reference angle is modelled as a one step reference advance.

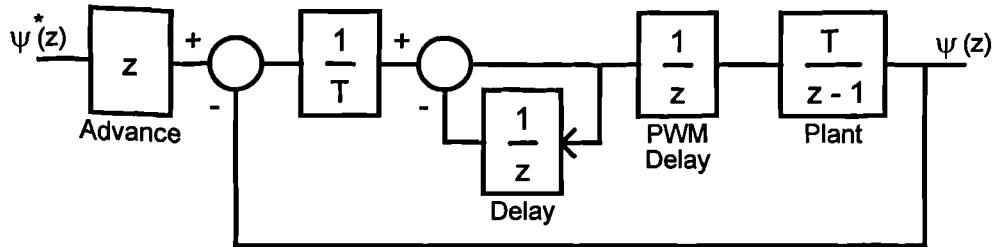


Figure 7.4 Simplified discrete time model of motor flux control system.

For this simplified model the transfer function of the forward path within the feedback loop is:-

$$DG(z) = \frac{1}{T} \cdot \frac{z}{z+1} \cdot \frac{1}{z} \cdot \frac{T}{z-1} = \frac{1}{z^2 - 1} \quad (7.10)$$

Closing the loop and incorporating the reference advance the transfer function of the system is:-

$$T(z) = z \cdot \frac{1}{(z^2 - 1) + 1} = \frac{1}{z} \quad (7.11)$$

Thus overall transfer function of the simplified system is dead-beat. Computationally the flux-controller is relatively simple, involving table lookup and a series of additions and multiplications. No trigonometric or exponential functions are required, furthermore no tuning of gain parameters is required other than measuring the dc-link voltage.

7.1.3 Implementation of a Flux-Linkage Controller for Single-Tooth Winding Machines

For a single-tooth winding switched reluctance machine drive, the flux-linkage controller is based on the dead-beat flux control system shown in Figure 7.3. All the control blocks shown in the controller block diagram Figure 7.5 are as in Figure 7.3, except for the rotor speed calculation. In the experimental system a capture timer on the TMS320C14 is used to measure speed as detailed in Appendix C and D. This gives much greater resolution and accuracy than differentiating rotor position. In addition to reference waveform selection, control system inputs are the actual phase currents $i(z)$, rotor position $\theta(z)$ and speed $\omega(z)$, with output $\delta(z)$ the power converter duty cycle that determines the phase voltage.

In Figure 7.5 the flux reference is actually derived from a current reference waveform using the machine flux-linkage/current characteristics. This method makes it easy to program the reference for a simple current waveshape such as unipolar 240° conduction.

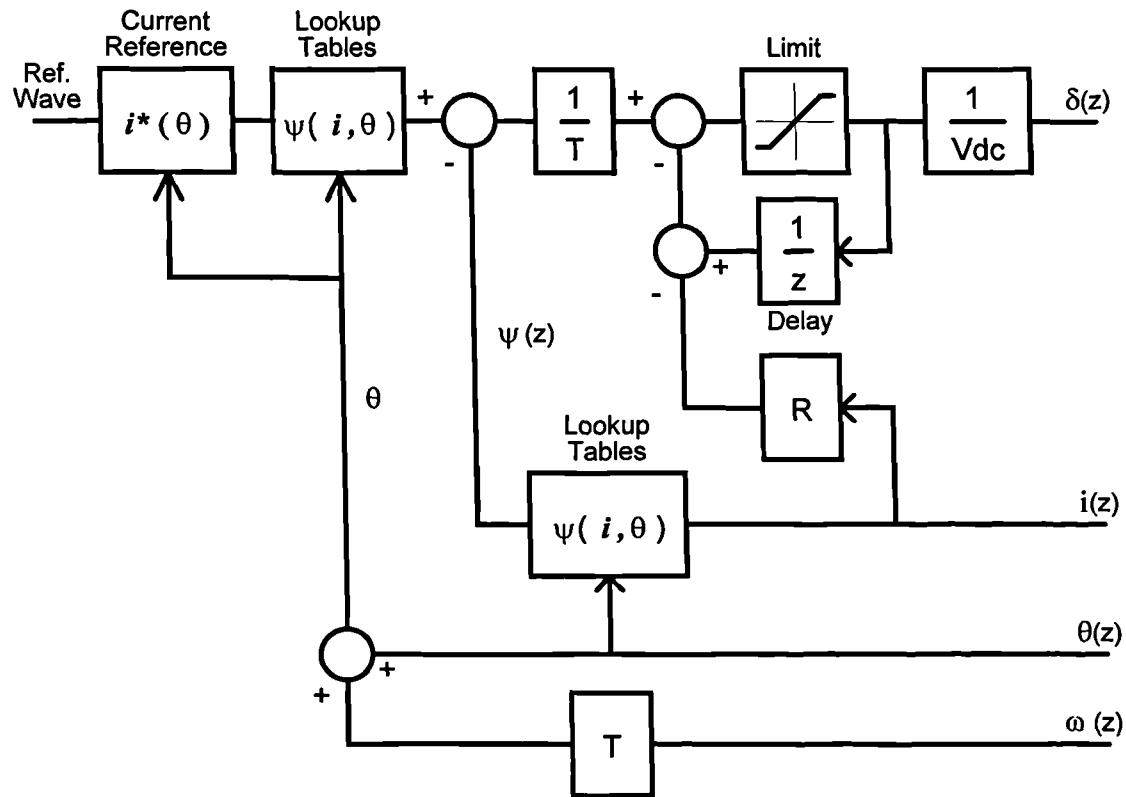


Figure 7.5 Flux-linkage controller for a single-tooth winding switched reluctance machine using current reference waveforms.

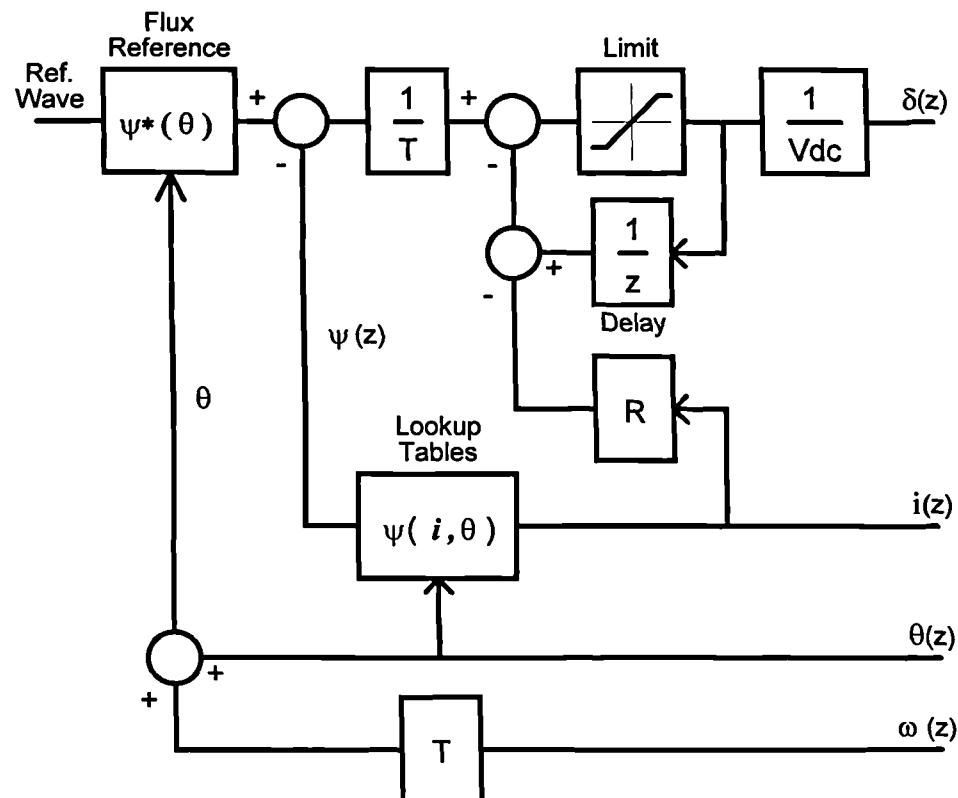


Figure 7.6 Flux-linkage controller for a single-tooth winding switched reluctance machine using flux reference waveforms.

For more complex waveshapes that have been calculated to meet specific performance criteria it is just as easy to provide a flux rather than current reference. As shown in Figure 7.6, this allows the current to flux-linkage block to be removed from the reference calculation thus reducing the amount of processing time required, an important factor in real time implementation.

7.1.4 Implementation of a Flux-Linkage Controller for Fully-Pitched Winding Machines

Figure 7.7 is a block diagram of the fully-pitched winding switched reluctance machine flux-linkage controller. As with the single-tooth winding flux-linkage controller, in addition to reference waveform selection, control system inputs are the actual phase currents $i(z)$, rotor position $\theta(z)$ and speed $\omega(z)$, with output $\delta(z)$ the power converter duty cycle.

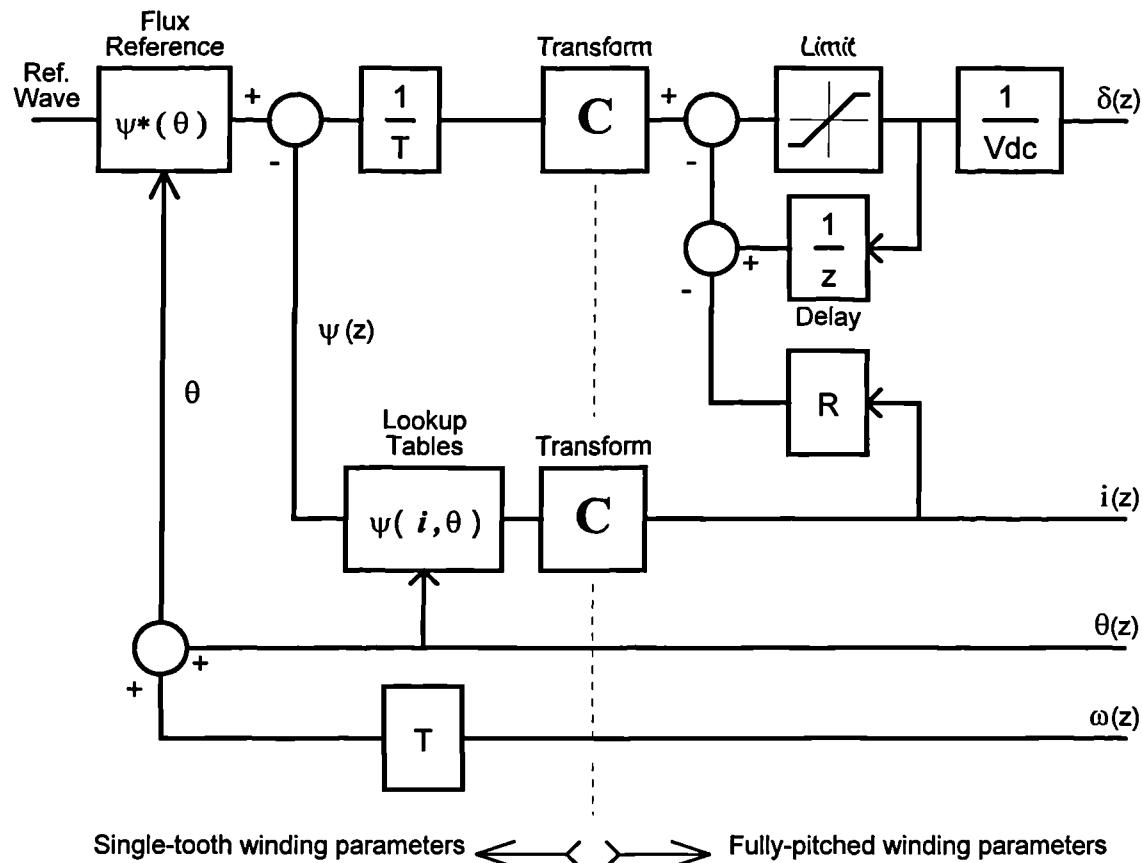


Figure 7.7 Flux-linkage controller for a fully-pitched winding switched reluctance machine using single-tooth winding flux reference waveforms.

As calculation of flux-linkage from current is done in the de-coupled equivalent single-tooth winding phases, actual current measurements must be transformed using the matrix C . Having calculated the single-tooth winding flux error and the voltage required

to correct the error, the voltage is then transformed back to fully-pitched winding values for application to the machine phases. Note that the output voltage limiting is done for the actual voltages as it is these which are limited by the power converter. Resistive voltage drop is also compensated for using the actual phase currents.

In Figure 7.7 the reference waveform is a single-tooth winding flux-linkage reference, this eliminates the need for conversion from current to flux-linkage and the transformation from fully-pitched winding to single-tooth winding parameters.

7.1.5 The Effect of Position Errors on Flux-Linkage Control

The performance of all switched reluctance drives depends critically on correct synchronisation of excitation and rotor position, when using a flux-linkage controller this can be even more important.

With a current controller the current waveshape is not affected by errors in synchronisation to rotor position. The resulting shaft torque depends then only on the torque position characteristics of the machine. For the experimental machine with a relatively flat peak to the torque characteristics, small position errors do not affect the torque dramatically.

If a flux-linkage controller has its flux reference value calculated from a current reference waveshape and rotor position in real time, then as with the current controller the resulting phase current profile is not affected by rotor position errors. Once again torque output is quite tolerant of small position errors.

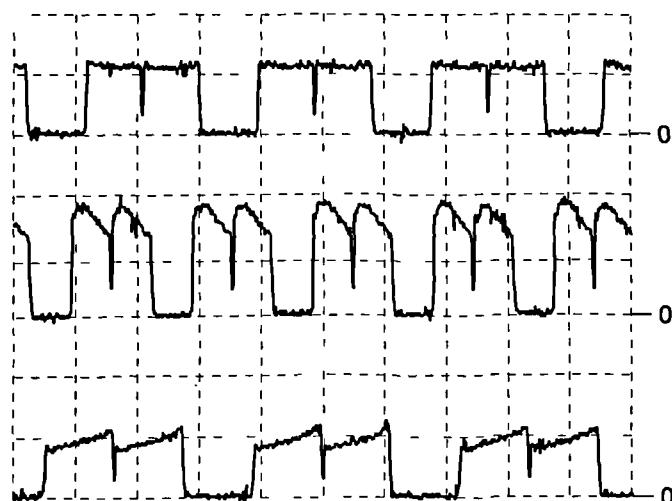


Figure 7.8 Flux-linkage reference position offset. Phase current I_a for nominal 240° unipolar conduction. Trace 1 - Nominal phase advance. Trace 2 - Phase advance 7.8° . Trace 3 - Phase advance -8.3° . Vertical 5A/div. Horizontal 50ms/div.

By using a flux-linkage rather than current reference waveform rotor position errors become much more significant. Flux references are calculated for a known synchronisation to rotor position so that the resulting phase current and torque can be found. If the actual rotor synchronisation does not correspond to that used to derive the reference then the resulting phase current and shaft torque are both affected. Figure 7.8 shows the effect of position errors either side of the nominal angle, for a flux reference calculated to give a unipolar current of constant amplitude and duration of 240°.

The effects of rotor position errors on current profile are most easily understood with reference to a single-tooth winding machine. By moving the flux reference towards the low inductance unaligned position the current for a given flux-linkage is greater. Conversely, moving the flux reference towards the higher inductance aligned position the current goes down. The same basic principle applies to a fully-pitched winding machine, so advancing the flux reference (Figure 7.8 trace 2) causes the current to increase, and retarding the flux reduces the current (trace 3). Rotor position errors can have a large impact on current waveshape as well as amplitude, as shown in Figure 7.9. Trace 1 shows the sinusoidal current that results from correct synchronisation of the flux reference to rotor position. As the reference is advanced towards the unaligned position (traces 2 and 3) the third harmonic current content rises considerably. If there is a large position error towards the unaligned position the current required to reach the desired flux-linkage can exceed the power converter rating, a condition against which the drive must be protected.

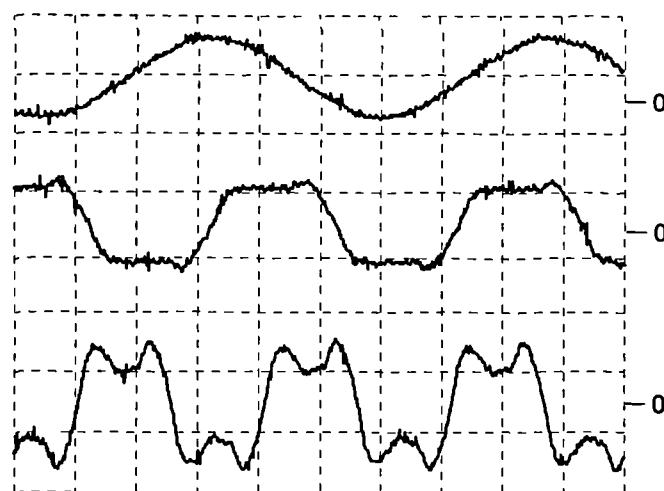


Figure 7.9 Flux-linkage reference position offset. Phase current I_a for nominal sinusoidal conduction. Trace 1 - Nominal phase advance. Trace 2 - Phase advance 5.8°. Trace 3 - Phase advance 11.1°. Vertical 10A/div. Horizontal 50ms/div.

Two factors contribute to the effect on shaft torque of flux reference position errors; the change in current waveshape, and the synchronisation of current to the torque profile. If

the reference is advanced towards the unaligned position then the current will tend to increase, however the current is now at a position with a lower torque per ampere. These two effects cancel each other out to some extent and the overall effect is dependant on the reference waveshape and machine characteristics. Moving the flux reference towards the aligned position will result in a fall in current and usually a fall in the torque per ampere, the torque thus falls rapidly with position error. The combination of these factors makes shaft torque much more sensitive to position errors with flux-linkage control than current control.

7.2 Experimental Results

7.2.1 Comparison of Flux-Linkage and Current Controller Responses

Four different phase current controllers were implemented to compare the performance of conventional current controllers and the experimental flux-linkage controller. Each controller was implemented digitally using hardware and software as described in Chapter 3 and Appendix D. The load in each case was a conventional single-tooth winding switched reluctance machine. Controller sample rate and converter switching frequency was 10kHz, with a single cycle delay from calculation of duty cycle to the application of phase voltage. Phase current reference was 10A peak with a conduction interval of 140° electrical per phase. For the flux-linkage controller flux reference values were calculated from the current reference using lookup table data. Average shaft torque is approximately 60Nm at 150rpm in each case. Measured phase current response and applied phase voltage for each of the controllers is illustrated in Figure 7.10.

A two level hysteresis controller was used to obtain the response shown in Figure 7.10(a). This controller has two possible output states: full positive volts and full negative volts. Current hysteresis band was set at 0.25A although the actual current ripple is much greater than this over most of the cycle due to system delays. If we consider the point in time when the actual phase current ramps outside the hysteresis band, ideally correcting volts would be applied at this instant, however delays in the current measurement system, controller and modulator, delay the detection of current going outside the hysteresis band and the subsequent application of correcting voltage. During this delay the current continues to ramp further outside the hysteresis band and so the actual current ripple is greater than the hysteresis band. Note how the current ripple varies with rotor position and phase current level. This is most obvious after the end of the conduction interval whilst the current reference is at zero. Falling phase inductance as the rotor moves away from the aligned position causes the current ripple to increase.

The performance of a two level hysteresis controller can be improved by allowing the applied phase voltage to be zero whilst phase current is within the hysteresis band. This results in a three level hysteresis controller. As can be seen in Figure 7.10(b) the current ripple and number of switching transitions are both reduced in comparison to the two level hysteresis controller. This is because whilst the applied voltage is zero current only falls due to the resistive voltage drop and 'motional emf'. The rate of change of current is thus much reduced and the current therefore takes longer to go outside the hysteresis band.

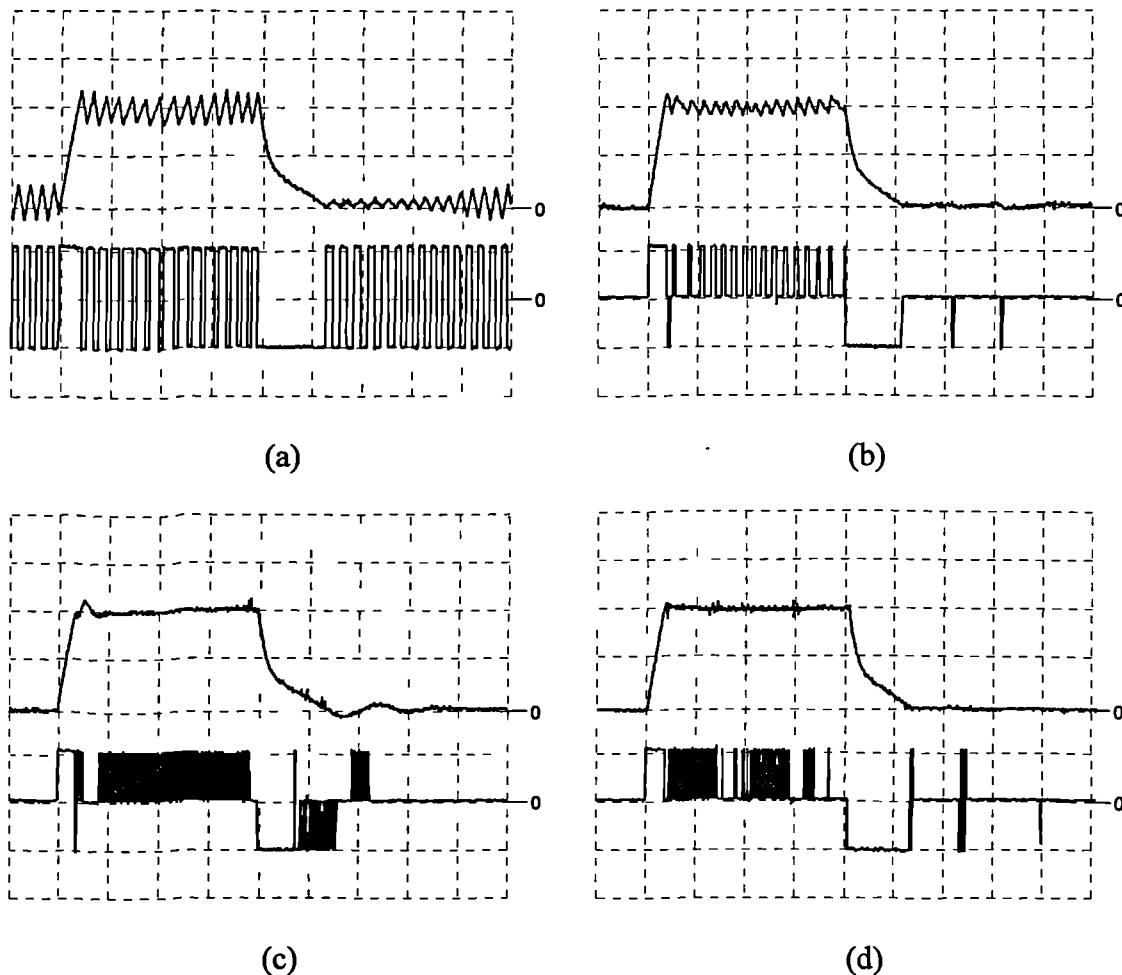


Figure 7.10 Comparison of measured phase current controller responses. (a) - Two level hysteresis current controller. (b) - Three level hysteresis current controller. (c) - PID current controller. (d) - Flux-linkage controller. In each case the traces are: Upper - Phase current 5A/div. Lower - Phase voltage 400V/div. Horizontal 5ms/div.

The PID current controller used to produce the response in Figure 7.10(c) has integral anti-windup features and is adaptive, as described in Chapter 4. Even in this optimised configuration it can be seen that as the set point is reached there is a damped oscillation in phase current. Note that the oscillation decays more slowly at the end of the

conduction interval when the rotor is near the aligned position and phase inductance is greater. Once at the set point the steady state current ripple is much less than for either of the hysteresis controllers.

Figure 7.10(d) shows the flux-linkage controller phase current response. This demonstrates the excellent response that can be achieved using a flux-linkage controller. Upon reaching the set point there is no obvious overshoot or damped oscillation in the phase current. Steady state current ripple as with the PID controller is very much less than either of the hysteresis controllers.

7.2.2 Flux-Linkage Control at Low Speed

All the results presented in this section are for the experimental fully-pitched winding switched reluctance machine and flux-linkage controller software. Extensive use is made of single-tooth winding parameters by the flux-linkage controller and in presenting experimental results. Two of the transformation equations derived in Chapter 2 section 2.2 will be restated here as they aid the understanding of experimental waveforms. Currents have the following relationship:-

$$I_a = -\frac{1}{2}(I_2 + I_3) \quad (7.12)$$

And flux-linkages are related by:-

$$\Psi_a = \Psi_1 - (\Psi_2 + \Psi_3) \quad (7.13)$$

Note that in oscillograms the equivalent single-tooth winding parameters are usually shown inverted so that both fully-pitched winding and single-tooth winding values are positive.

For each set of results the flux-linkage reference is calculated to give unipolar phase current with an amplitude of 10A and duration of 240°. The corresponding equivalent single-tooth winding current is 20A peak with a duration of 120°.

Operating at 90rpm with an average shaft torque of approximately 50Nm, the actual and equivalent single-tooth winding currents are shown in Figure 7.11, also shown in this figure is the equivalent single-tooth winding flux-linkage demand. It is evident that the flux controller gives a good current response with rapid transitions, no overshoot and constant current during the conduction intervals.



Figure 7.11 Flux-linkage controller operation at 90rpm 50Nm, showing real and equivalent single-tooth winding parameters. Trace 1 - Phase current I_a , 5A/div. Trace 2 - Equivalent current I_1 , 10A/div. Trace 3 - Equivalent flux demand ψ_1^* , 1Wb/div. Horizontal 5ms/div.

For the same conditions Figure 7.12 shows the actual phase current, flux-linkage and flux-linkage demands. With reference to the equivalent single-tooth winding flux-linkage in Figure 7.12, equation 7.12 and equation 7.13, it can be seen how the fully-pitched winding flux reference is calculated. As the flux reference makes a step change the actual flux ramps to the new value over a short portion of the period.

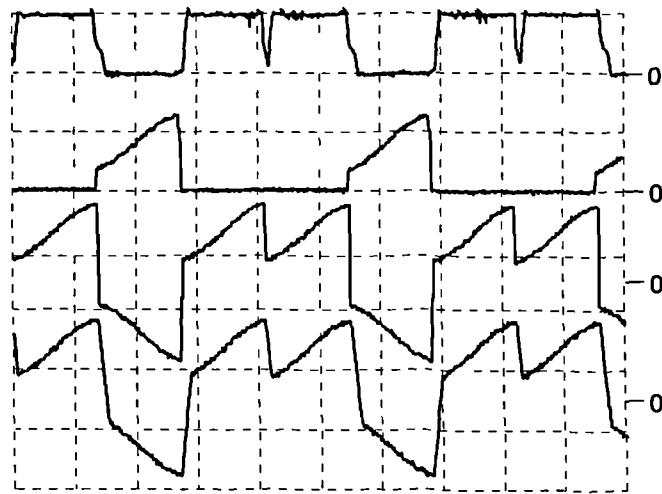


Figure 7.12 Flux-linkage controller operation at 90rpm 50Nm, showing real current and flux-linkage. Trace 1 - Phase current I_a , 10A/div. Trace 3 - Equivalent flux-linkage demand ψ_1^* , 1Wb/div. Trace 3 - Flux demand ψ_a^* , 1Wb/div. Trace 4 - Actual phase flux-linkage ψ_a , 1Wb/div. Horizontal 5ms/div.

At the higher speed of 200rpm but with the same flux demand as previously, it can be seen from traces two and three in Figure 7.13 that the actual flux-linkage takes a greater

proportion of the period than in Figure 7.12 to reach the new flux value following a step change in the reference. Also shown in Figure 7.13 is the phase voltage. Following a reference step change maximum output voltage is applied and remains at this level until the flux reaches the new reference value. It can be seen that the actual flux-linkage ramps in a linear manner whilst the constant voltage is applied.

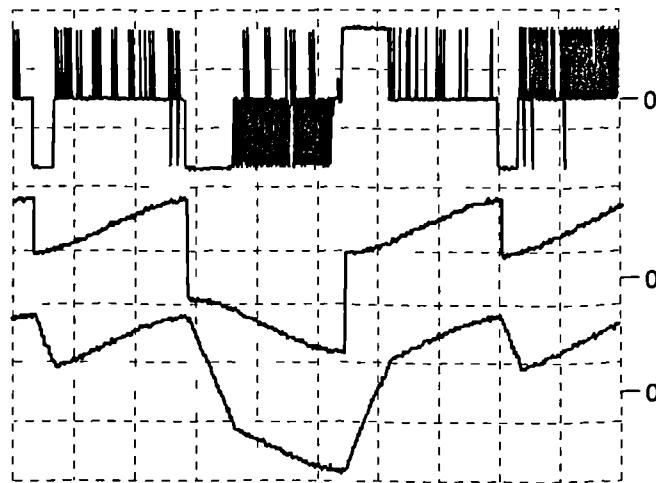


Figure 7.13 Flux-linkage controller operation at 200rpm 50Nm, showing phase voltage and flux-linkage. Trace 1 - Phase voltage V_a , 400V/div. Trace 2 - Flux demand ψ_a^* , 1Wb/div. Trace 3 - Actual phase flux-linkage ψ_a , 1Wb/div. Horizontal 5ms/div.

There is one very obvious difference between the current response of a PID current controller and the experimental flux-linkage controller which deserves some explanation: the large dip in phase current as other phases commutate. From the previous figures it is clear that there is a step change in actual flux-linkage demand in the middle of the phase conduction interval as well as at either end. The flux-linkage controller responds to these step changes in demand by applying maximum voltage of the appropriate polarity to each phase. Analysis of the equivalent single-tooth winding parameters shows that by simultaneously applying maximum voltage to all three actual phases, voltages cancel each other out delaying the current response for a period determined by the difference in aligned and unaligned flux-linkage, and it is this which gives rise to the dip in current.

By contrast the current reference fed to a PID current controller has no step change in the middle of the conduction interval and only the phases that are commutating have a large voltage applied immediately. In this case the phase current controller only applies a negative voltage in response to current overshoot as the other phases commutate, and as a result there is little disturbance in the phase current.

7.2.3 Flux-Linkage Control at High Speed

When using a flux-linkage controller to control phase currents at low speed reference tracking is good, and the current waveforms are similar to those obtained with a PID current controller, with the exception of effects due to phase commutation as described in the previous section. As speed increases the actual flux takes longer to ramp to the new value following a step change in reference. This is evident if Figure 7.12 and Figure 7.13 are compared. The limit of control comes when the actual flux just reaches the demand value following one reference step change, before the next step change occurs. This condition is illustrated in Figure 7.14. The limiting point in this case is the positive peak in the flux-linkage, which is only just reached before the reference steps down. Under these conditions the phase voltage and current are very similar to those obtained with a PID current controller: the current has taken on the 'twin peaks' waveshape characteristic of higher speed operation and shaft torque is approximately 25Nm.

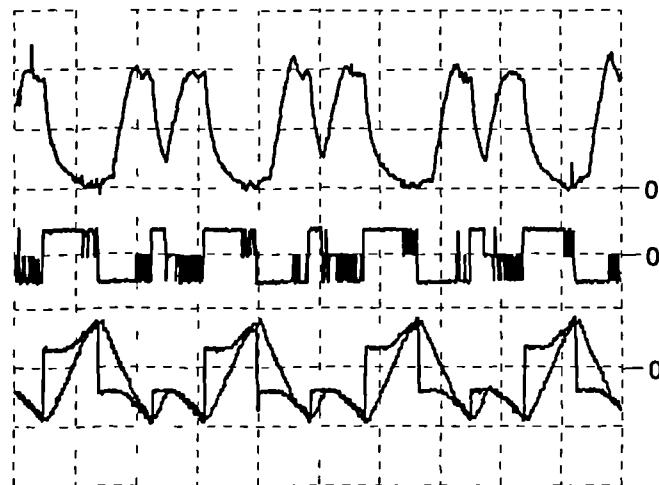


Figure 7.14 Flux-linkage controller operation on the limit of control at 560rpm. Trace 1 - Phase current I_a , 5A/div. Trace 2 - Phase voltage V_a , 1kV/div. Traces 3 and 4 - Flux-linkage demand ψ_a^* and actual flux-linkage ψ_a , 1Wb/div. Horizontal 5ms/div.

As speed is increased further there is insufficient time for the actual flux to cover the peak to peak range required. The actual flux is then no longer able to reach the reference in all segments of the flux reference. Figure 7.15 shows this condition with flux unable to reach the positive peak flux value. Flux-linkage is not now fully controlled, and under these conditions operation differs markedly from PID current controller operation. When the flux-linkage controller is unable to reach the peak to peak flux reference value the actual peak to peak flux-linkage is defined by the speed and dc-link voltage. As the average flux reference value is negative this tends to offset the actual flux towards the peak negative value so, as in Figure 7.15, the actual flux reaches the peak negative value

but not the peak positive value. This leads to gross current waveform distortion and a dramatic fall in shaft torque. By contrast a PID current controller has a positive average reference value so the current tends to have a positive offset when speed increases to the point where the peak to peak value can no longer be achieved. The resulting current waveform distortion is less severe and torque output falls gradually with speed.

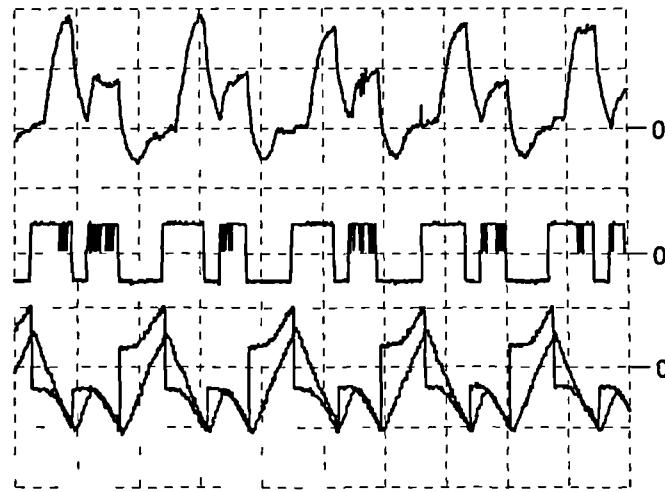


Figure 7.15 Flux-linkage controller operation beyond the limit of control at 680rpm. Trace 1 - Phase current I_a , 5A/div. Trace 2 - Phase voltage V_a , 1kV/div. Traces 3 and 4 - Flux-linkage demand ψ_a^* and actual flux-linkage ψ_a , 1Wb/div. Horizontal 5ms/div.

7.3 Advanced Switched Reluctance Drives Using Flux-Linkage Control

7.3.1 Calculation of Flux References for High Speed Operation

From the previous results for operation at speeds on the limit of flux control and beyond, it may appear that flux-linkage control is not suitable for high speed operation. However the poor torque performance is due solely to the inappropriate flux reference value which was derived directly from the current reference. By correct selection of the flux reference values operation can be extended to high speeds with several advantages over current control. However, a potential disadvantage is that flux reference calculation is more complex than required for current controller operation.

For a single phase carrying current i with resistance R and flux-linkage ψ , the applied voltage V is given by:-

$$V = iR + \frac{\partial\psi}{\partial t} = iR + \frac{\partial\psi}{\partial\theta} \cdot \omega + \frac{\partial\psi}{\partial i} \cdot \frac{di}{dt} \quad (7.14)$$

And so the rate of change of current is:-

$$\frac{di}{dt} = \frac{V - iR - \frac{\partial\psi}{\partial\theta} \cdot \omega}{\frac{\partial\psi}{\partial i}} \quad (7.15)$$

Thus the maximum rate of change of current is a non-linear function of phase current and position. Assuming that the current slew rate is limited by the available voltage, rather than controller bandwidth, calculating phase current references that the current controller will be able to track is made difficult by the non-linear variation of maximum current slew rate. In contrast, if a flux controller is used then the maximum rate of change of flux is:-

$$\frac{\partial\psi}{\partial t} = V - iR \quad (7.16)$$

If the resistive voltage drop is small compared with the dc-link voltage then the equation may be simplified by making a fixed allowance for the resistive voltage drop equal to the maximum value expected. The maximum rate of change of flux-linkage is then simply the dc-link voltage less a constant:-

$$\frac{\partial\psi}{\partial t} = V_{dc} - I_{max} R_{max} \quad (7.17)$$

Thus the maximum rate of change of flux-linkage is independent of phase current and rotor position, simplifying calculation of flux-linkage reference values. Most switched reluctance machine drives use controllers that supply phase currents or voltages that are functions of rotor position rather than time: therefore a more useful measure of flux-linkage rate of change would be with respect to rotor position, rather than time. Neglecting the resistive voltage term in equation 7.16 we obtain:-

$$V = \frac{\partial\psi}{\partial t} = \frac{\partial\psi}{\partial\theta} \cdot \frac{d\theta}{dt} = \frac{\partial\psi}{\partial\theta} \cdot \omega \quad (7.18)$$

and so:-

$$\frac{\partial\psi}{\partial\theta} = \frac{V}{\omega} \quad (7.19)$$

Because there is a simple fixed limit on the rate at which the flux reference may change, for a given speed and dc-link voltage, the limits of the flux controller are easily determined, in contrast to the equivalent case for a current controller.

Rather than having a flux reference waveshape that it is only possible to track at low speed, by limiting the rate of change of flux-linkage to be less than the maximum possible value for a given speed, the flux reference can always be tracked and performance is predictable. This is most easily achieved by composing the flux reference from a series of linear flux ramps, these are chosen to give the desired output without exceeding the maximum rate of change. Figure 7.16 shows operation with a five segment flux reference. Waveforms displayed in this figure are the equivalent single-tooth winding values rather than the actual fully-pitched winding values as these are much easier to interpret. The phase voltage trace was obtained by filtering the 10kHz PWM voltage signal with a 470 μ s continuous time low pass filter. Allowing for the effects of the filter lag, the phase voltage is almost constant over each flux ramp section, and the amplitude is proportional to the flux ramp gradient.

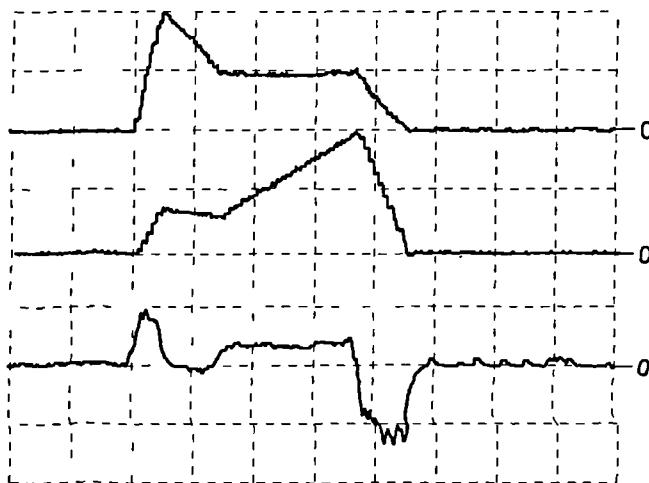


Figure 7.16 Flux-linkage controller response for a constant flux ramp reference. Trace 1 - Equivalent phase current I_1 , 10A/div. Trace 2 - Equivalent phase flux-linkage ψ_1 , 0.5Wb/div. Trace 3 - Equivalent phase voltage V_1 , low pass filtered, approximately 400V/div. Horizontal - 2ms/div..

The flux-linkage reference in Figure 7.17 was designed using flux ramps that could be tracked at all times by the controller, in contrast to those in Figure 7.15 where due to large step changes the controller is unable to follow the flux reference over a large part of the cycle. Using a flux reference that can be tracked throughout the cycle gives much improved performance. Figure 7.17 shows operation at 730rpm which is 30% higher than the limiting speed for the current derived flux reference. Although the speed is higher it has been possible to produce approximately 30% more torque for a similar peak phase current.

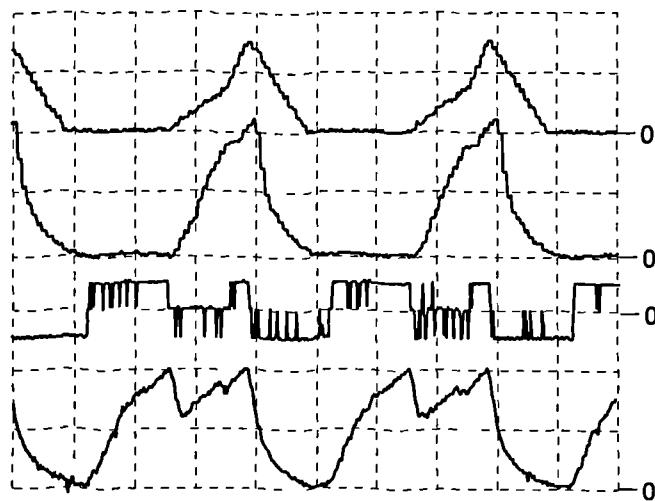


Figure 7.17 Flux-linkage controller with flux ramp reference at 730rpm 32Nm. Trace 1 - Flux-linkage demand ψ_1^* , 0.5Wb/div. Trace 2 - Equivalent phase current I_1 , 10A/div. Trace 3 - Phase voltage V_a , 1kV/div. Trace - 4 - Phase current I_a , 5A/div. Horizontal 2.5ms/div.

7.3.2 Improving High Speed Torque Capability

As speed increases there comes a point where the flux ramps up and down merge together and ultimately a voltage square wave is applied to the phase windings. Under these conditions operation is the same as would occur if using a current rather than flux controller. Phase current and torque are now dependant on the phase angle between rotor position and the applied voltage. Although the phase voltage is the same in each case, using a flux controller does have advantages over a current controller as an offset can easily be added to the flux reference. When a voltage squarewave is applied to a phase the flux-linkage follows a triangular waveshape if the resistive voltage drop is neglected. With the current controlled so that it just reaches zero at its minimum point, the flux-linkage waveform also has a minimum of zero, with the peak to peak amplitude defined by the applied volt-seconds. In steady state the same volt-seconds can be used to produce a flux waveform with the same peak to peak value but offset from zero. By offsetting the flux waveform from zero is possible to produce more torque for the same peak current and peak to peak flux-linkage. Figure 7.18 illustrates for a conventional single-tooth winding switched reluctance machine how this is possible. Shaded areas indicate the theoretical bounds of the flux-linkage/current locus that will produce maximum torque, with the darker area indicating the difference in total enclosed area. In each case peak phase current i and peak to peak flux-linkage $\Delta\psi$ are equal, but in Figure 7.18(b) the minimum flux is offset from zero.

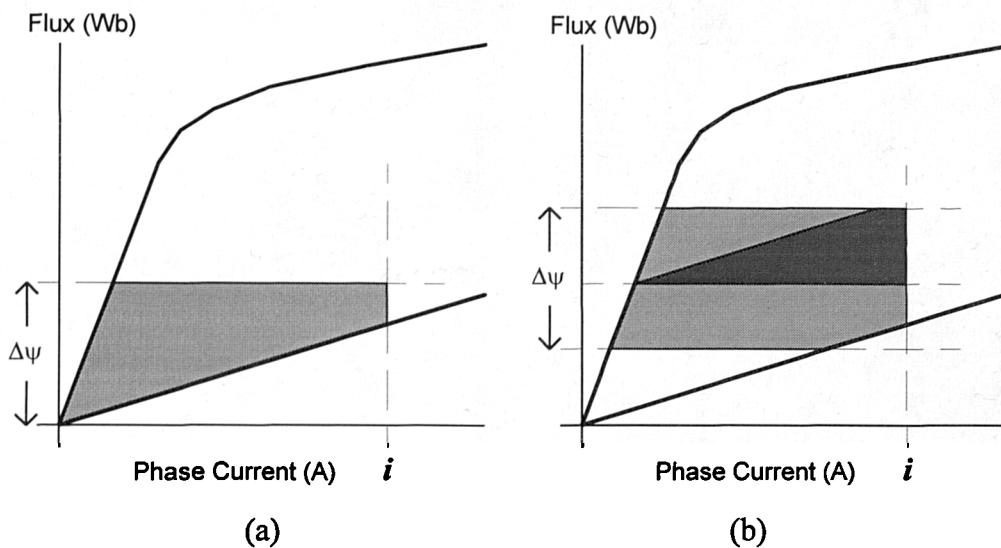


Figure 7.18 Maximum flux-linkage/current loci for a given peak current and peak to peak flux excursion.

It is evident that the shape of the aligned and unaligned magnetisation curves allow a greater area to be enclosed within the locus by offsetting the flux from zero. If this principle is applied to the experimental single-tooth winding machine the two loci shown in Figure 7.19 are obtained. In each case the peak current is limited to less than 25A and the peak to peak flux-linkage is 0.47Wb.

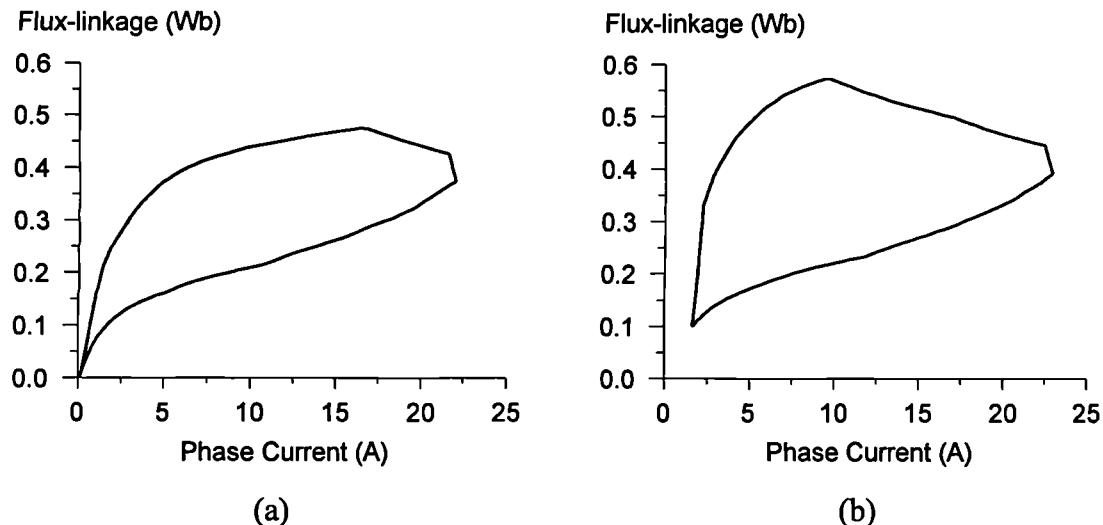


Figure 7.19 Flux-linkage loci for the experimental single-tooth winding machine with flux-linkage offsets of 0.0Wb (a) and 0.1Wb (b). Voltage squarewave applied.

Without a flux offset the average torque is 17.4Nm, but by offsetting the flux by 0.1Wb the average torque is increased to 22.8Nm. As flux is offset from zero phase current is no longer discontinuous and there is an increase in RMS phase current, in this case from 10.0A to 11.9A. Thus by offsetting the flux reference torque per ampere increases by 11% from 1.74Nm/A to 1.92Nm/A.

This principle can also be applied to a fully-pitched winding switched reluctance drive with similar beneficial results. Simulation of the experimental fully-pitched winding switched reluctance drive at a speed of 2000rpm shows the peak to peak flux-linkage is 0.93Wb (the equivalent single-tooth winding flux-linkage is 0.31Wb). With a flux offset of zero the torque at optimum phase advance is 11.2Nm. By offsetting the flux reference by 0.2Wb (the equivalent single-tooth winding flux offset is 0.2Wb) and allowing a moderate increase in phase current, the torque can be more than doubled to 23.0Nm. Peak and RMS phase currents increase from 10.8A to 14.1A, and from 5.6A to 9.3A respectively. The overall effect is to increase the torque per RMS ampere of phase current from 2.01Nm/A to 2.47Nm/A, a gain of 23%.

7.4 Summary and Conclusions

Consideration of a simple inductive circuit and unity gain feedback control system has shown that the discrete time control function required for 'dead-beat' flux response is very simple: requiring only a multiplication by the controller sample frequency. A 'dead-beat' flux-linkage controller has been implemented for switched reluctance machine control and shown to give excellent phase current control. Comparisons with hysteresis and PID current controllers having the same sample rate showed the flux-controller response to be superior in every case.

Operation of the experimental switched reluctance drive using the flux-linkage controller has been investigated at low and high speed. Low speed performance is very good but current and torque waveforms are more sensitive to rotor position errors than current control. At high speed the flux-linkage controller can give poor results with inappropriate reference waveforms. Taking advantage of the simple relationship between flux-linkage and voltage, reference waveforms can be derived which give excellent high speed performance. In experimental results 30% more torque was produced at a 30% higher speed (69% more shaft power) when using a reference waveshape designed for flux control instead of one designed for current control. By addition of an offset to the flux reference it has been shown that the torque per RMS ampere can be increased at high speed. In simulated results for the fully-pitched winding switched reluctance drive the average torque was doubled and the torque per RMS ampere increased 23% by offsetting the flux reference 0.2Wb from zero.

Although very good performance can be achieved using a flux-linkage controller it does require machine flux-linkage characteristics to be accurately known, and high resolution rotor position sensing must be used to ensure that the desired waveshapes are achieved.

Chapter 8

Constant Torque Operation

For many applications, such as tensioners, torque is the controlled parameter. In others, speed controllers often incorporate a torque reference as part of their control structure, so it is desirable to have a linear relationship between a reference value and the mean shaft torque. Servo applications may require smooth movement, and to reduce stress and wear on transmissions, constant torque with respect to rotor position is desirable. However, the doubly-salient and saturating nature of a practical switched reluctance machine results in torque characteristics which are highly non-linear with both phase current and rotor position. A simple switched reluctance drive, with constant phase current excitation, gives a shaft torque which is not linearly proportional to the current demand, nor is it constant with respect to rotor position. It is these shortcomings that have restricted the acceptance of switched reluctance drives for high performance applications and contributed to their perception as being rough and noisy.

By using more advanced control schemes, which use knowledge of machine characteristics to calculate the required excitation, it is possible to produce a high performance switched reluctance drive giving constant torque with respect to rotor position and having a linear torque/reference characteristic.

8.1 Theory of Operation

In an idealised switched reluctance machine reluctance varies linearly with tooth overlap and for a constant MMF constant torque is produced. In a real machine saturation and field fringing mean that the torque produced for a constant phase current is a non-linear function of position. Figure 8.1(a) shows the experimental machine torque characteristics. From these it clear that simple excitation patterns producing a constant MMF will result in substantial torque ripple, especially at higher current and torque levels. In order to produce constant shaft torque the current must be profiled. Phase current contours required for the experimental machine to produce constant torque are shown in Figure 8.1(b). As the aligned or unaligned rotor position are approached the current required increases dramatically as the torque per ampere falls towards zero.

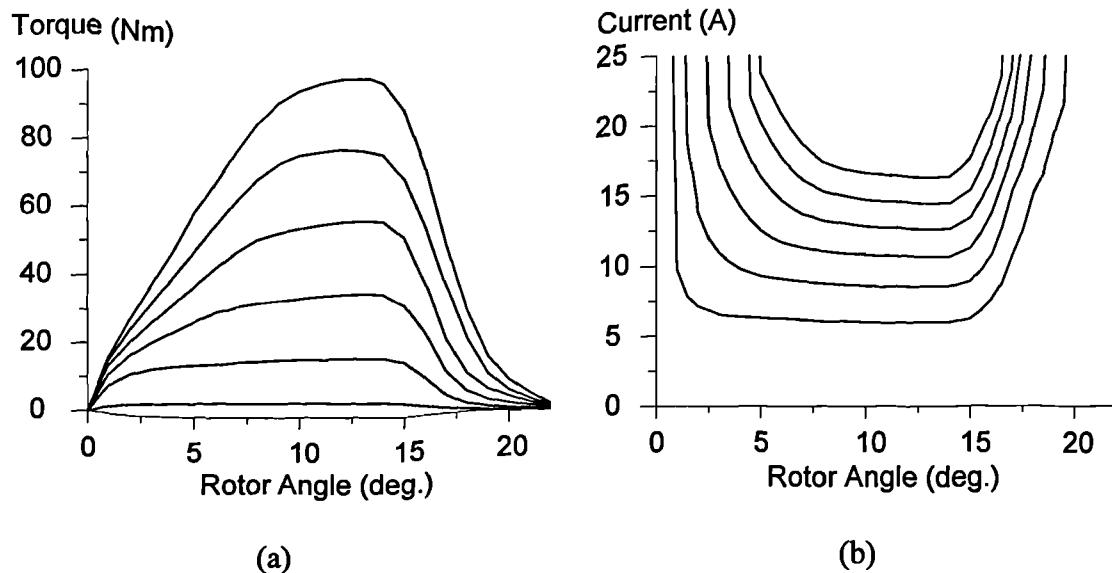


Figure 8.1 Constant current and constant torque machine characteristics. (a) Torque position characteristics for phase current of 4A to 24A in steps of 4A. (b) Current position characteristics for torque of 10Nm to 60Nm in steps of 10Nm.

Considering operation of the experimental three phase single-tooth winding switched reluctance machine:- During the positive torque half cycle the first 60° and last 60° overlap with the preceding and following phases, leaving the middle 60° during which only one phase can produce positive torque. At low speed phase currents can be controlled so as not to produce any negative torque contributions. Hence during the middle 60° the phase current profile must follow that required for constant torque, i.e. the characteristic in Figure 8.1(b). As the torque per ampere near the aligned and unaligned positions falls very rapidly, two phases must contribute to the output torque so as not to exceed the current capability of the converter. It is the different angles over which torque production commutes from one phase to the next, and the current and torque profiles during this period, that distinguish different modes of constant torque operation. Some of the published methods used to calculate the required current and torque profiles are reviewed in the next section.

8.2 Review of Constant Torque Methods

A review of the work published on constant shaft torque operation of switched reluctance machines reveals that the emphasis has been on low speed operation, and using closed loop current controllers. All of the published work relates to operation of conventional single-tooth winding switched reluctance drives. Several of these methods are examined, and have been compared using simulation, and experimentally using the fully-pitched winding switched reluctance drive.

8.2.1 Basis of Methods

There are several basic methods published for the torque control of switched reluctance machines. The methods that will be reviewed here can be summarised as follows:-

- Analytical waveform calculation.
- Control of switching angles.
- Constant rate of change of torque during phase commutation.
- Feedback linearisation.
- Constant rate of change of current during phase commutation.
- Application of Neural Networks.

Analytical Waveform Calculation and Control of Switching Angles.

Work by Matsui *et al* [8.1] models a switched reluctance machine as a synchronous reluctance machine and derives the phase currents required for constant torque by algebraic methods. In order to facilitate this analysis the machine is assumed to be magnetically linear and further assumptions are made about the phase self and mutual inductances. Amor *et al* [8.2] also derive the phase currents required for constant torque operation by analytical derivation based on using the rate of change of co-energy to calculate torque. Co-energy is calculated from flux-linkage characteristics, but as with Matsui *et al* the machine is assumed to be magnetically linear in order to simplify the equations enough for an analytical solution to be found. Switched reluctance machines are generally designed to saturate quite heavily in order to maximise output and so this assumption may not be valid for practical machines. Jiang *et al* [8.3] minimise torque ripple by control of the conduction and advance angles used for phase current control. Although torque ripple can be reduced by correct selection of the control angles, the minimum value is determined by the ripple in machine torque/position characteristics when operating with constant phase current. No attempt is made to control the phase currents to overcome this limitation and so this method cannot truly be considered suitable for constant torque operation.

Constant Rate of Change of Torque.

Methods based on a constant rate of change of torque during commutation between phases has been explored by Kavanagh *et al* [8.4], Rochford *et al* [8.5], and Schramm *et al* [8.6]. In each case the machine static torque characteristics are measured directly or calculated from other indirect measurements, such as flux-linkage. The torque reference for each phase follows a ramp and pedestal waveshape, where the ramp rate and pedestal width are chosen so as to optimise the desired parameter. The desired phase

currents are then found from the torque reference by calculation off line, or from lookup tables. In each case the phase current reference is passed to a closed loop current controller. Rochford *et al* use an additional voltage feed-forward term in the current controller to improve control bandwidth. Operation is only demonstrated for low speed operation so that the phase current controller acts as an almost ideal current source.

Feedback Linearisation.

Taylor [8.7] and Ilic-Spong *et al* [8.8] use feedback linearisation in a switched reluctance motor driven position control system to achieve a linear relationship between the demand and actual shaft torque. This is achieved by current shaping, so that the current references are non-linear functions of torque demand and position. The torque profile for each phase, method of current reference calculation, and phase current control used by Taylor are all implemented as described above for the work by Kavanagh *et al*, Rochford *et al* and Schramm *et al*. Once again only low speed operation is considered so that the real phase currents can track the references accurately.

Constant rate of change of current.

It is possible to have a constant current ramp, instead of a constant torque ramp, during phase commutation and this is investigated by Wallace *et al* [8.9] and Wallace *et al* [8.10]. Current is commutated between phases at the point where each gives the same torque per ampere. Either side of this position the incoming phase current ramps up linearly and the outgoing phase current ramps down linearly. The phase current reference for the active phase is calculated from machine static torque characteristics; closed loop current controllers then impress these currents on each phase. As with the work by Kavanagh *et al*, Rochford *et al*, and Schramm *et al* only low speed operation is considered so that the real phase currents track the reference accurately.

Work by Steiert *et al* [8.11] also makes use of constant current ramps during commutation but only at switch off. Advantage is taken of the lower phase inductance at the unaligned position to give a greater current bandwidth than at the high inductance aligned position. A linear current ramp down is used approaching the aligned position, with the ramp rate selected such that the voltage required does not exceed the nominal dc-link voltage. At other times during a cycle the phase current is calculated from machine torque characteristics, so as to give constant shaft torque. Current reference values are stored in lookup tables and closed loop current controllers are used to regulate the phase currents. Machine torque characteristics are determined from co-energy, which in turn is found from flux-linkage measurements made with sinusoidal

voltage excitation. In contrast to the other methods reported in this section results are presented for a speed range of 100rpm to 1000rpm.

Application of Neural Networks.

Reay *et al* [8.12] and Reay *et al* [8.13] have published several papers on the application of Neural Networks to switched reluctance motor control. Neural Network inputs are torque demand and rotor position, with phase current references the outputs. During training of the network measured torque is also fed back as a network input. Training of the network is accomplished using initial estimates for phase currents, and torque transducer feedback to produce phase current references that result in constant shaft torque. After training the resulting phase torque and current profiles are quite similar to those obtained with a constant torque ramp during phase commutation. The actual phase currents are regulated with closed loop current controllers. Speed is not used as an input to the network so that only low speed operation is possible, although it was stated that further work would consider $di/d\theta$, allowing speed to be taken into account.

Other Methods.

It is also possible to calculate the phase currents required to obtain constant torque with minimum copper losses. Minimum RMS phase currents dictates that the phase which produces the greatest torque per ampere at the desired operating point and rotor position should produce the largest proportion of the required torque. At the point where two phases produce the same torque per ampere the current in each phase will be equal. For the fully-pitched winding switched reluctance machine there may be two or three phases excited at any one time and the currents can be designed for use with unipolar or bipolar converters. The non linear torque characteristics of the experimental machine mean that numerical methods must be used to find the phase currents which give minimum loss.

8.2.2 Simulation

Programs were written in the high level language 'C' running in a DOS environment to calculate torque and current profiles for various constant torque methods. The program assumes ideal current controllers and uses experimental machine characteristics. Single-tooth winding phase currents are determined from torque/position characteristics, and then transformed to their fully-pitched winding equivalents. Results are presented graphically, and data can be saved for further analysis or as a lookup table for real time operation. The constant torque method selected can be any of the following:-

- Constant rate of change of torque during phase commutation.
- Constant rate of change of current during phase commutation.
- Minimum instantaneous RMS currents, unipolar or bipolar.
- Sinusoidal phase currents (for comparison purposes).
- Unipolar two phase on, 240° conduction (for comparison purposes).

In each case the various control parameters appropriate to that control method are optimised for minimum peak or RMS phase currents, for either fully-pitched winding or single-tooth winding switched reluctance machines.

Simulated results are presented here for three modes of operation, all designed to give an average torque of 30Nm, with the current profiles optimised for operation with the fully-pitched winding switched reluctance machine.

Constant torque ramp operation is illustrated in Figure 8.2, with torque commutation from one phase to the next over a period of 25° (electrical). From the equivalent single-tooth winding current, it is evident that current peaks at either end of the conduction interval are approximately equal, and the corresponding fully-pitched winding current has a large peak in the middle of conduction. A shorter commutation period tends to increase the peak current required as a single phase has to produce all the torque over a longer period.

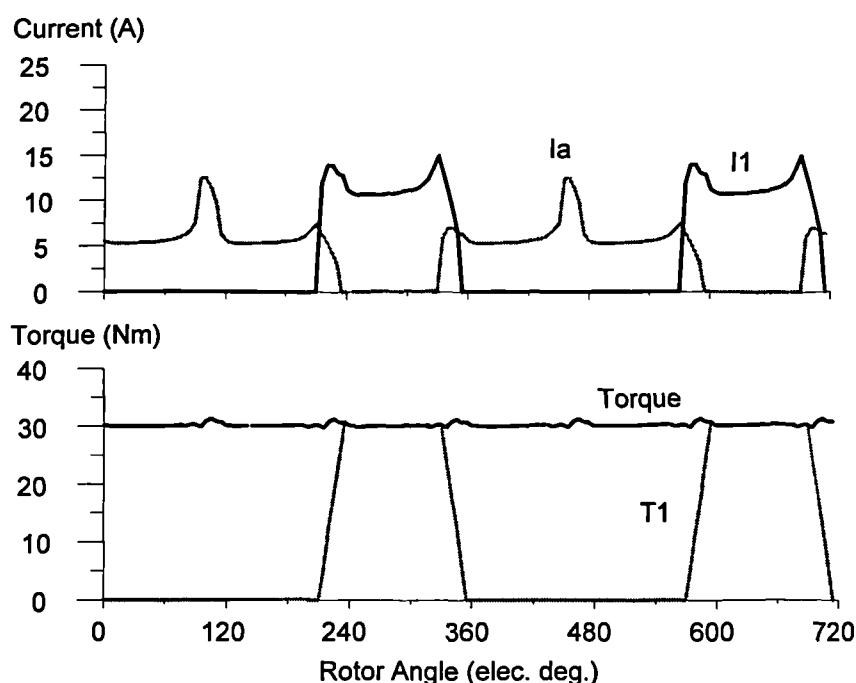


Figure 8.2 Simulated constant torque operation, constant torque ramp method. Real and equivalent single-tooth winding current, single phase and total torque.

The constant current ramp method used by Steiert *et al* is used to calculate the single-tooth winding current profile in Figure 8.3. In this case the current ramps down over 45° towards the aligned position resulting in a pronounced peak in equivalent single-tooth winding current at the start of conduction. This is because the long current ramp down restricts torque production near to the aligned position (360° and 720°) and so there is a large torque contribution required from near the unaligned position. Reducing the ramp period increases the peak current at the end of the conduction period and reduces the peak at the start of the conduction period.

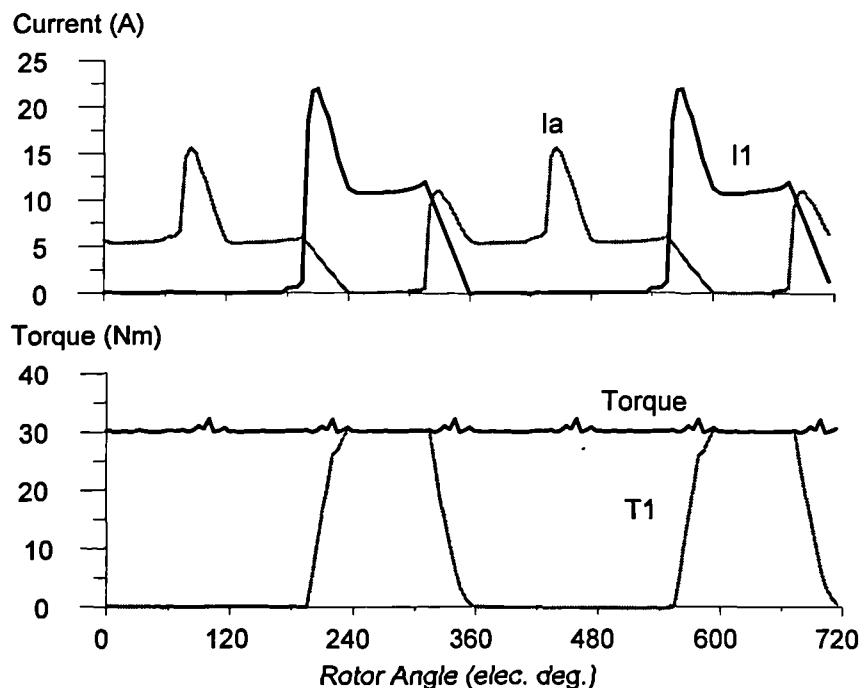


Figure 8.3 Simulated constant torque operation, constant current ramp method. Real and equivalent single-tooth winding current, single phase and total torque.

Figure 8.4 shows simulated operation with unipolar minimum RMS phase currents. Looking at the equivalent single-tooth winding current and single phase torque in Figure 8.4, it is evident that torque production commutes from one phase to another over a shorter period than with the previous two methods. With this method of operation the current peak at the end of the equivalent single-tooth winding conduction period is greater than at the beginning. If the currents are optimised for minimum RMS currents in the equivalent single-tooth windings, then commutation of torque production occurs over a longer period. Note that the torque ripple is due to the current quantisation level used in the simulation program.

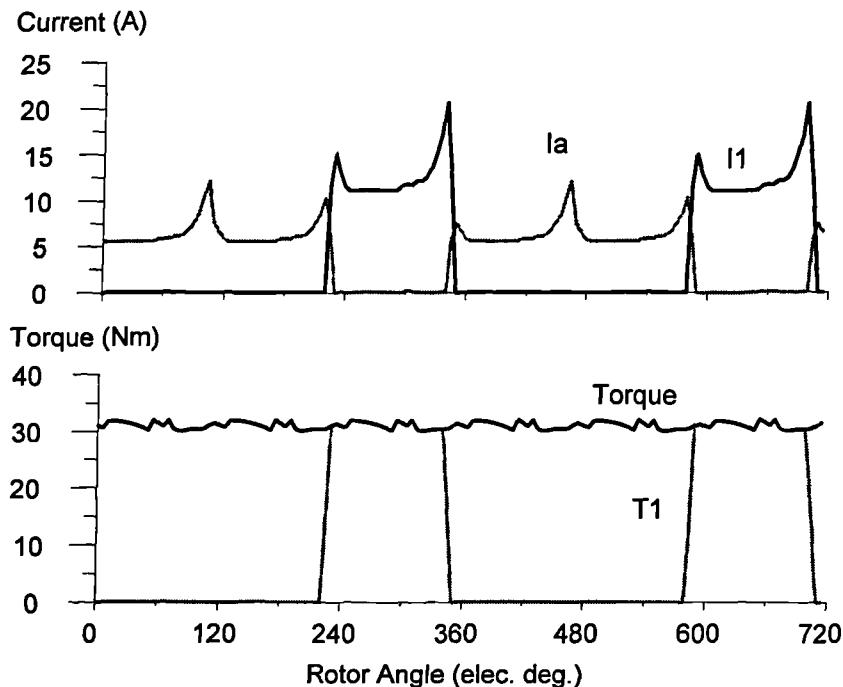


Figure 8.4 Simulated constant torque operation, minimum RMS current method. Real and equivalent single-tooth winding current, single phase and total torque.

Table 8.1 compares the fully-pitched winding and equivalent single-tooth winding currents for various methods, each giving an average shaft torque of 30Nm. The upper section of the table gives figures for constant torque operation and the bottom two lines figures for standard waveforms, for comparison. Each of the constant torque methods gives low torque ripple whilst sinusoidal operation has 32% peak to peak ripple, and unipolar 240° operation has 75% peak to peak torque ripple. From figures for phase current it is clear that constant torque operation with unipolar phase currents extracts a high penalty in peak current rating. Peak fully-pitched winding currents are at least twice the value for unipolar 240° operation, and the RMS currents are also slightly higher. Bipolar minimum RMS operation has a much reduced peak rating compared to unipolar operation but twice as many switching devices are required. Comparing figures for the two bipolar modes of operation, minimum RMS, and sinusoidal. The minimum RMS method has a peak current only 8% higher, and the lower RMS current would give 19% lower copper losses. Thus for constant torque operation, unipolar operation requires a much greater converter rating than for operation with simple current profiles, but bipolar operation has a very small penalty in converter rating and can substantially reduce machine copper losses.

Type	Optimisation	Commu- tation period (°)	Fully-pitched winding current (A)		Single-tooth winding current (A)	
			Peak	RMS	Peak	RMS
Torque ramp	Min. RMS	5	12.9	5.24	21.1	7.3
	Min Pk	25	12.5	5.42	15.0	7.2
Current ramp	Min. RMS	5	12.5	5.29	21.14	7.3
	Min. Pk	25	12.4	5.43	16.9	7.2
	-	45	15.7	6.44	22.1	8.3
Min. RMS Unipolar	-	-	13.5	5.34	22.6	7.45
Min. RMS Bipolar	-	-	±7	4.15	±14	7.86
Sinusoidal	-	-	±6.5	4.6	±12.4	9.24
Unipolar 240°	-	-	6	4.9	12	6.9

Table 8.1 Comparison of phase currents for an average shaft torque of 30Nm.

8.2.3 Experimental Comparison

Several of the constant torque methods described in the last two sections have been studied experimentally to compare various performance aspects. Simulation software was used to generate lookup tables for the following modes of operation:-

- Constant torque ramp, commutation over 25°.
- Constant torque ramp, commutation over 5°.
- Constant current ramp, commutation over 45°.
- Constant current ramp, commutation over 25°.
- Unipolar minimum instantaneous RMS currents.
- Bipolar minimum instantaneous RMS currents.
- Sinusoidal phase currents (for comparison purposes).
- Unipolar two phase on, 240° conduction (for comparison purposes).

Early experiments used a closed loop current controller but results were unsatisfactory because the current controller did not have a sufficiently fast response to track the complex current profiles adequately. The flux-linkage controller described in Chapter 7 section 7.1.4 was used for all the results in this chapter as it has a fast well damped response, well suited to constant torque operation. Lookup tables holding the equivalent single-tooth winding flux-linkage provide the phase references, and the power converter is configured for bipolar operation with an 'H' bridge per phase. In each case the phase reference is calculated to give an average shaft torque of 30Nm. Equivalent single-tooth winding parameters and instantaneous torque are calculated in real time by the controller from measured phase currents and position.

Low Speed Operation

Figure 8.5 shows operation at 57rpm for constant torque ramp operation (based on the method used by Schramm *et al*), and constant current ramp operation (based on the method used by Steiert *et al*). Comparing the fully-pitched winding and single-tooth winding current traces with those in Figure 8.2 and Figure 8.3, it is clear that the actual currents are able to track the reference values used in the simulation. Figure 8.5 also shows the equivalent single-tooth winding flux-linkage.

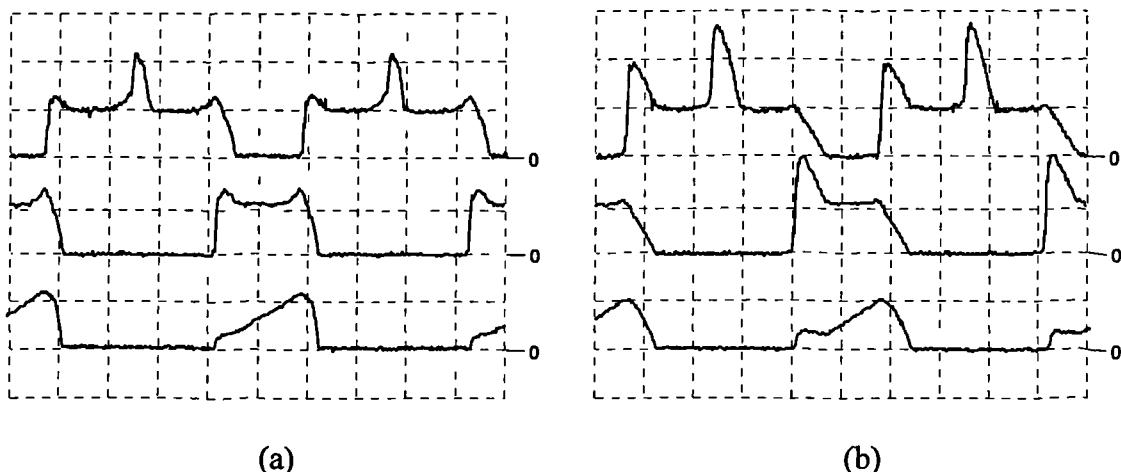


Figure 8.5 Constant torque operation at 30Nm, 57rpm. (a) - Constant torque ramp, 25° commutation. (b) - Constant current ramp, 45° commutation. Trace 1 - Phase current I_a , 5A/div. Trace 2 - Equivalent single-tooth winding current 10A/div. Trace 3 - Equivalent single-tooth winding flux-linkage 1Wb/div. Horizontal 25ms/div.

Plotting the equivalent single-tooth winding flux-linkage/current loci for various modes of operation reveals several aspects of operation. Figure 8.6 shows the flux-linkage/current loci for four modes of constant torque operation; constant torque ramp with a slow and fast ramp, constant current ramp, and bipolar minimum RMS.

Comparing Figure 8.6(a) and (b) it can be seen that the constant current ramp method has a large torque contribution near the unaligned position and little from near the aligned position, whilst the constant torque ramp method has equal contributions from each position. Figure 8.6(c) shows constant torque ramp operation but with the commutation period reduced from 25° to 5° . From the locus it can be seen that the peak MMF near the aligned position is substantially greater than in (a), although the area in this region is very small. The large MMF is required to maintain constant torque because the torque per ampere drops rapidly as the aligned position is approached. However, in terms of mechanical energy delivered to the load the large MMF peak provides very little. Although the bipolar minimum RMS loci encloses twice the area, one cycle corresponds to a rotation of two rotor tooth pitches, and hence is swept at half the rate of unipolar modes, therefore the resulting average torque is the same.

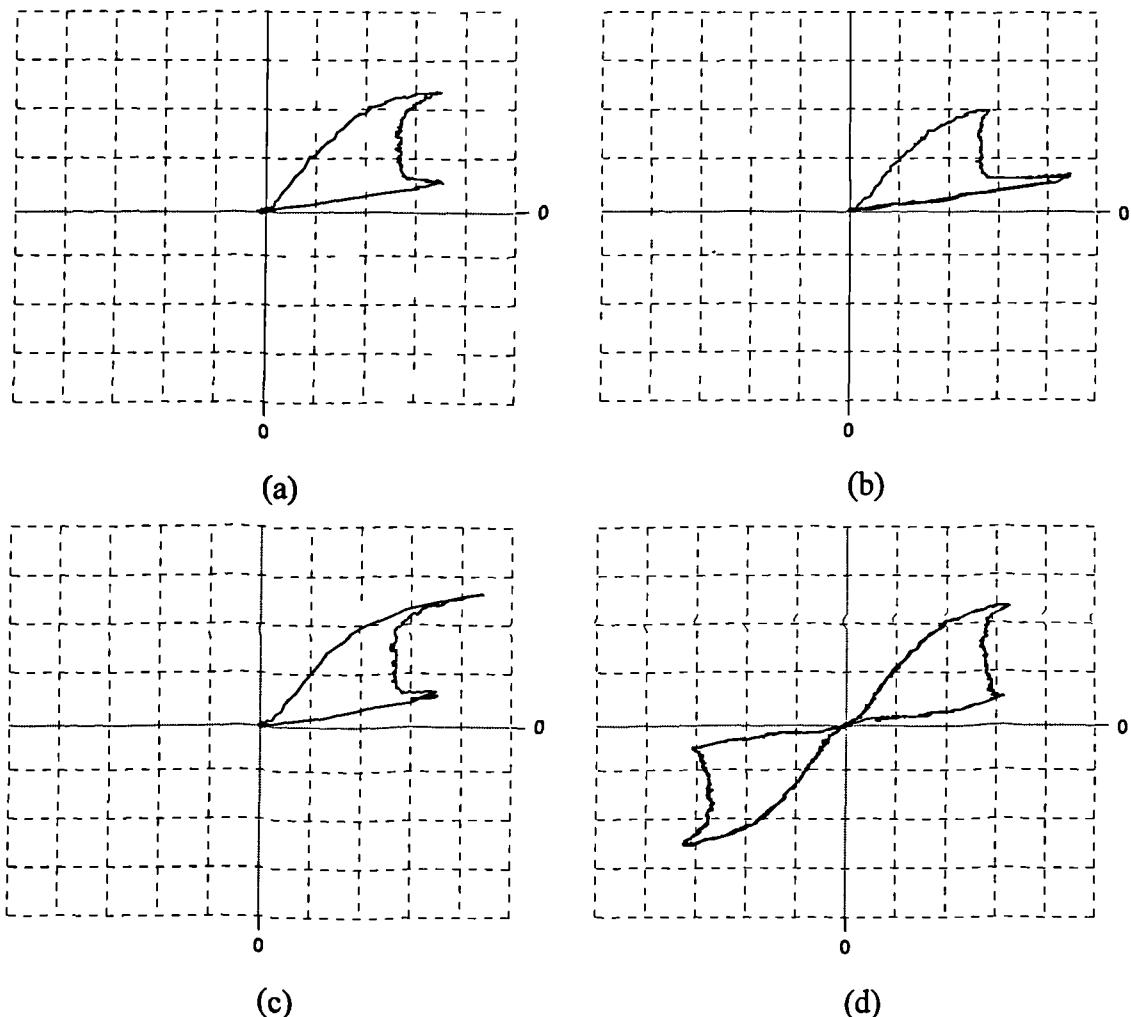


Figure 8.6 Constant torque operation equivalent single-tooth winding flux-linkage/current loci. All 30Nm average torque at 57rpm. (a) - Constant torque ramp, 25° commutation. (b) - Constant current ramp, 45° commutation. (c) - Constant torque ramp, 5° commutation. (d) - Bipolar minimum RMS currents. Horizontal - Current 4A/div. Vertical - Flux-linkage 0.5Wb/div.

Each method of operation has a quite distinctive audible sound which leads to conclusions about the mechanisms of noise generation in a switched reluctance drive. Comparisons are made here are on a subjective basis only. The first and most notable conclusion is that low torque ripple in itself does not mean low acoustic noise, as shown by lack of correlation between the torque ripple and acoustic noise for waveforms tested. This concurs with the conclusion drawn by Cameron *et al* [8.14]. The second conclusion is that the major noise source is due to ovalising forces acting on the stator. This conclusion also agrees with the findings of Cameron *et al* and Wu *et al* [8.15]. The amplitude of acoustic noise which results from radial stator vibrations depends on the strength of the ovalising forces and their rate of change. Ovalising forces are greatest at the aligned position where the tangential force reduces to zero and the radial force reaches a peak, thus a rapidly reducing MMF near the aligned position will result in maximum oscillation. This is shown experimentally by the very much lower acoustic noise produced by the constant current ramp method with 45° commutation and sinusoidal operation, compared to the constant torque ramp method with 5° commutation. Sinusoidal and current ramp modes have a low and slowly reducing MMF near the aligned position, whilst the torque ramp method has a high and rapidly reducing MMF near the aligned position.

Measured equivalent single-tooth winding current, flux-linkage and electromagnetic torque for three modes of operation are shown in Figure 8.7. The methods illustrated are constant torque ramp, constant current ramp, and unipolar minimum RMS currents. Torque is calculated in real time by the DSP controller from measured phase currents and rotor position using stored machine torque characteristics. The torque spikes apparent in Figure 8.7(c) are due to the rapid change in phase current required as torque production commutes from one phase to another. Due the finite dc-link voltage, even at the low operating speed the flux is unable to rise and fall at the desired rate, and the resulting current error produces a torque perturbation. In Figure 8.7(a) and (b) no torque perturbations are evident because the actual flux is able to track the demanded value.

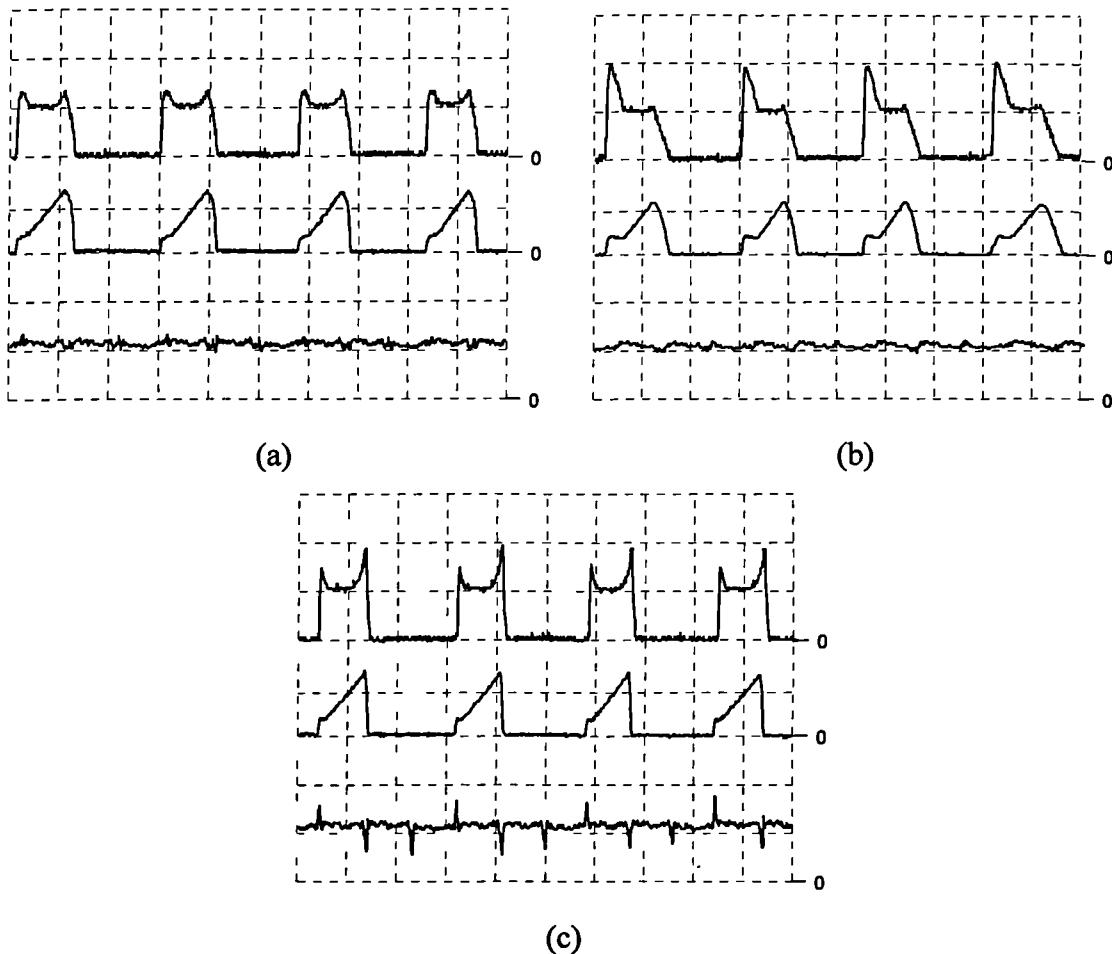


Figure 8.7 Current flux-linkage and torque for constant torque operation at 30Nm, 57rpm. (a) - Constant torque ramp, 25° commutation. (b) - Constant current ramp, 45° commutation. (c) - Unipolar minimum RMS currents. Trace 1 - Equivalent single-tooth winding current 10A/div. Trace 2 - Equivalent single-tooth winding flux-linkage 1Wb/div. Trace 3 - Torque 25Nm/div. Horizontal 50ms/div.

Operation at Higher Speeds

As speed increases the required $d\psi/dt$ and voltage increases for a given reference waveform. In general the largest $d\psi/dt$ occurs as torque production commutes between phases, and it is at this point that the actual flux-linkage is unable to track the reference. This leads to torque perturbations (as in Figure 8.7(c)) which get larger as speed increases. Waveforms such as the constant current ramp method with 45° commutation, which have a lower $d\psi/dt$ than other methods, can sustain constant torque operation to a higher speed than other modes. However all modes reach a point where constant torque operation is no longer possible because the required flux-linkage is changing faster than can be achieved in practice.

As speed increases Steiert *et al* reduce the current ramp rate approaching the aligned position so that the voltage required does not exceed what is available. The reduced current near the aligned position is compensated for by increasing the current near the low inductance unaligned position, where lower volts are required for a given di/dt . Although this method works well for conventional single-tooth winding switched reluctance drives, to the point where the unaligned voltage also reaches the available voltage, it does not work so well for a fully-pitched winding switched reluctance drive. This is because the fully-pitched winding phase voltage required depends on all three equivalent single-tooth winding voltages, which with overlapping periods can require a fully-pitched winding voltage greater than that available.

8.3 Simulation

In an effort to find a control method that would give constant torque operation over a wide speed range several simulation programs were written in the 'C' language to try out various ideas. Two basic approaches were taken: current control, and voltage control on a PWM cycle by cycle basis.

8.3.1 Current Control for Constant Torque

The basis of current control for constant torque is to select phase currents which give the required level of torque with minimum RMS currents and hence copper losses. The low speed operation of this method has been shown theoretically and experimentally in sections 8.2.2 and 8.2.3 of this chapter. To extend operation to higher speed the voltage required to produce the desired current waveform must be taken into account. The method used can be summarised in pseudo-code as follows:-

```

FOR all rotor positions
    FOR a range of possible single-tooth winding currents
        calculate the fully-pitched winding currents
        calculate the change in flux-linkage
        calculate the fully-pitched winding voltages
        IF current < limit AND voltage < dc-link voltage
            calculate shaft torque
            calculate losses
            IF torque > demand AND loss < lowest found so far
                save currents
        NEXT current
    NEXT position

```

The code carries out a systematic search at each rotor position for the set of phase currents that will give the lowest loss for a given torque. For each potential solution the change in flux-linkage from the last position is used to calculate the voltage required. If this voltage exceeds the dc-link voltage, or the fully-pitched winding phase current would exceed the current limit level, then the solution is rejected.

This method works well at low speed as the change in flux-linkage requires a voltage less than the dc-link voltage. As speed increases the voltage limit comes into play for short periods. When this happens the current slew rate is effectively limited and slightly sub-optimum currents (in terms of loss) must be selected. There may be short periods of negative torque contributions, but these are compensated for by increasing the current in phases which can give positive torque. Increasing the speed further extends the periods over which negative torque is produced. This happens because until the point where two equivalent single-tooth winding phases have the same torque per ampere, a single phase is producing most of the torque. As this point is only about 30° (electrical) from the aligned position there is little time for the current to ramp down to zero from the peak value. As current takes longer to reach zero the negative torque producing period extends and grows in magnitude. This then requires a larger current in the following phase to compensate for the negative torque contribution, and the system quickly goes unstable as successive phases try in vain to compensate for the previous phase.

Not only does this method not work at anything other than low speed, it also requires extensive computation at each position, and each possible set of currents at each position. Thus this method would not be suitable for real time implementation.

8.3.2 Voltage Control for Constant Torque

With the current control method several possible sets of currents were evaluated and checked for voltage limiting. Checking the voltage can be eliminated and the number of possible solutions reduced by using voltage rather than current control. Voltage is controlled on a PWM cycle by cycle basis, with the duty cycle set to 0% or 100%. Over a PWM cycle (of $100\mu\text{s}$) the average voltage applied to each phase can be $+V_{dc}$, 0, or $-V_{dc}$, or in per-unit values +1, 0 or -1. As each phase can be any one of three voltages, and there are three phases, there are a total of 27 possible fully-pitched winding voltage vectors. The corresponding single-tooth winding voltages have per unit values of +2, +1, 0, -1, or -2. There are thus 125 (5^3) single-tooth winding voltage vectors, but only 27 of them can be produced by application of voltages to the fully-pitched windings.

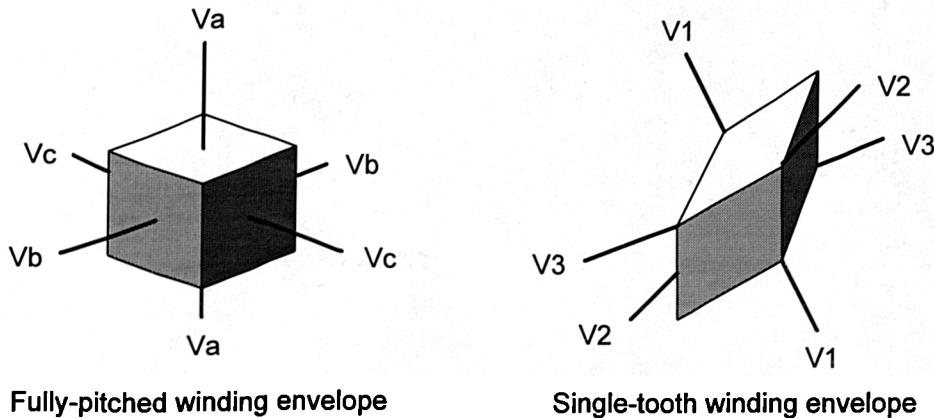


Figure 8.8 Three dimensional representation of voltage vectors.

Figure 8.8 shows the relationship between the fully-pitched winding and single-tooth winding voltage envelopes. In three dimensions the fully-pitched winding voltage envelope is a unit cube. When transformed to equivalent single-tooth winding voltages the cube becomes a three dimensional rhombus, i.e. it is a cube which has been skewed in all three dimensions, and has a peak magnitude of 2 per-unit on each axis. Although the required current and flux-linkage, and hence voltage vector, are much easier to calculate in single-tooth winding parameters, the complex relationship between the fully-pitched winding and single-tooth winding voltage envelopes makes it difficult to predict which of the desired single-tooth winding voltage vectors will fall outside the envelope of possible voltages. For this reason all 27 fully-pitched winding voltage vectors are tested, even though many will not give the required single-tooth winding vector. The algorithm used can be represented by the following pseudo-code.

```

FOR each sample interval
    FOR 27 fully-pitched winding voltage vectors
        calculate the single-tooth winding voltage vector
        calculate the single-tooth winding flux-linkage
        calculate the single-tooth winding current
        calculate the fully-pitched winding current
        IF fully-pitched winding current < limit
            calculate the torque
            calculate the losses
            IF torque > demand AND loss < lowest found so far
                save voltage vector
    NEXT voltage vector
NEXT sample interval

```

For each sample interval all 27 possible voltage vectors are tried, this involves calculating the equivalent single-tooth winding flux-linkage, from which the current and torque can be found. If any vector would result in a current that exceeds the limit it is rejected, otherwise the vector that gives the required torque with the lowest loss is selected.

As with the current control method, voltage control works well at low speed but suffers from the same problems at higher speed, namely successive phases being unable to compensate for negative torque produced by the last phase. Much less computation is required than the current control method, and if fully optimised the algorithm could be used for real time implementation.

These simulation studies have shown that it is not possible to have constant torque operation over a wide speed range by considering the present conditions only. A strategy must be developed which takes into account what will happen much later in the cycle, in fact up to 360° later.

8.4 Flux Ramps for Constant Torque

All of the published methods reviewed in section 8.2.1 of this chapter used closed loop current controllers to regulate the machine phase currents. It was assumed that the current controller would be able to track the current references, and except for the work by Steiert *et al*, have the same reference waveshape for all speeds of operation, indeed most work is only demonstrated at a single low speed. In the experimental evaluation of some of the published methods (section 8.2.3) a flux-linkage controller was used. However, this was only to achieve the required current control bandwidth and response, the reference waveshapes were still based on those used for current control and no account was taken of the voltage required to achieve the desired waveshape. Other simulation work (section 8.3) has shown that in order to extend constant torque operation to higher speed operation the waveform needs to be planned over the cycle as a whole and cannot be done on a point by point basis.

When designing a current reference waveform for constant torque operation one of the biggest problems is maintaining constant torque operation over a wide speed range and with variations in dc-link voltage. This is because the non-linear machine characteristics make it difficult to take into account the rate of change of flux-linkage that will be required for a given current reference. Although constant torque operation is often designed using machine torque/current characteristics to derive a current reference waveform, it would be just as easy to derive a flux-linkage reference waveform given the machine flux-linkage/current or torque/flux-linkage characteristics. By using a flux-

linkage reference waveform is much easier to take into account practical limitations because if resistive voltage drop is neglected, the rate of change of flux-linkage is simply equal to the applied voltage. Furthermore when a large change in flux-linkage is required the converter output will tend to saturate and the flux will ramp at a linear rate (resistance neglected). These considerations lead to the idea of flux ramps for constant torque operation. With this scheme the flux-linkage reference waveform is composed of a series of linear ramps, each with a gradient no more than the greatest actually achievable. Having linear ramps makes it possible to predict operation later on in the cycle, as the time taken to ramp to a new flux-linkage value is easily calculated.

8.4.1 Theory of Operation

Equations derived in Chapter 7 section 7.3.1 showed that for a single phase carrying current i , with resistance R , flux-linkage ψ and with an applied voltage V , the rate of change of current with time is a non-linear function of phase current and position. This makes it difficult to calculate phase current references that the current controller will be able to follow. It was also shown in Chapter 7 that if resistance is neglected the rate of change of flux-linkage is simply equal to the applied voltage, and is independent of phase current and rotor position. Expressing the rate of change of flux with respect to position in terms of rotor speed, ω , and applied voltage gives:-

$$\frac{\partial \psi}{\partial \theta} = \frac{V}{\omega} \quad (8.1)$$

Therefore the maximum rate of change of flux-linkage with respect to rotor position can be found from the rotor angular velocity and the maximum available phase voltage.

8.4.2 Flux-linkage Profiles for Constant Torque

The flux-linkage profile required to produce 30Nm in a single phase of the experimental single-tooth winding switched reluctance machine is shown in Figure 8.9. This figure also shows flux profiles for the constant current ramp, constant torque ramp, and minimum RMS phase current methods, all for a shaft torque of 30Nm. It is clear from this figure that the flux-linkage rises in a linear fashion over a large proportion of the torque producing period. In all the flux-linkage profile required for constant torque operation could be considered to comprise of four distinct sections; an initial rapid rise in flux-linkage, a modest fall as the torque per ampere increases rapidly, a constant rise in flux linkage as the rotor moves towards the aligned position, and finally a rapid fall in flux-linkage at the end of the excitation period.

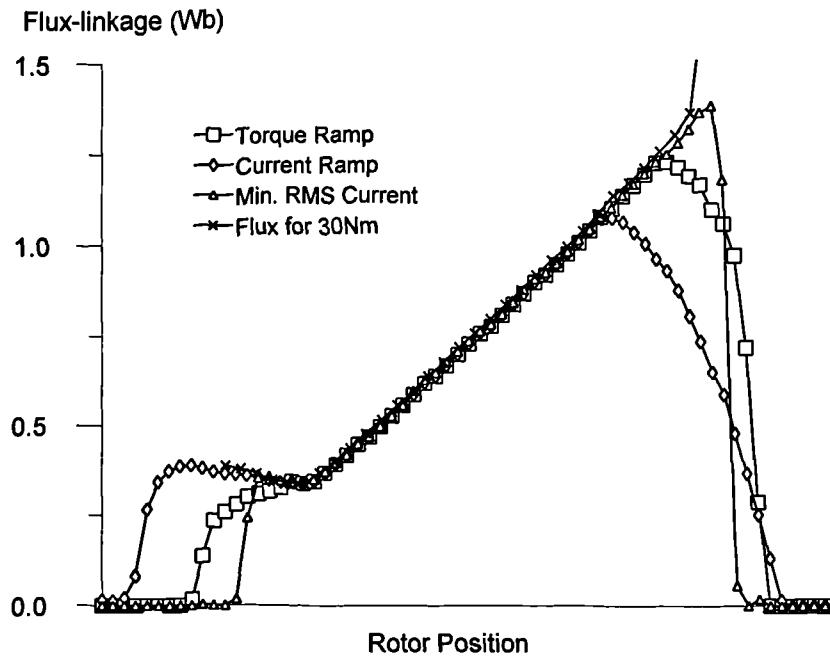


Figure 8.9 Machine 30Nm flux-linkage profile and flux-linkage characteristics for various 30Nm constant torque control methods.

There are two main differences between the flux-linkage waveforms of the three control methods shown in Figure 8.9; the rotor positions at which the flux is ramped up and down at either end of the conduction period, and the shape of the flux-linkage curve during these ramps. It can be seen that the minimum RMS current method requires the flux to fall much more quickly at the end of the conduction period than the constant current ramp method.

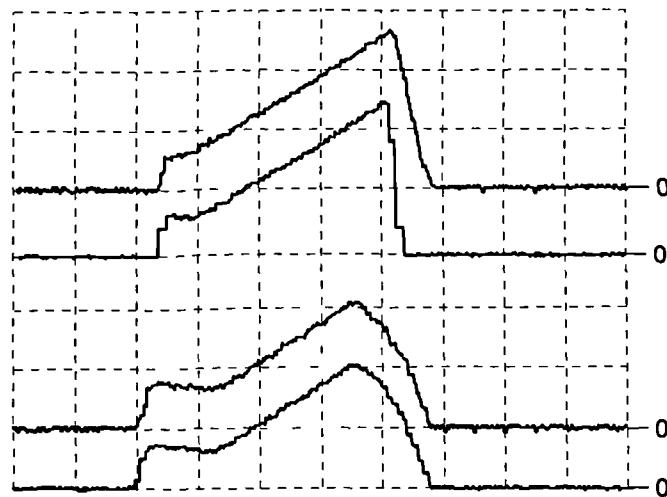


Figure 8.10 Constant torque flux profiles showing the effects of controller saturation. Top two traces - Minimum RMS current method. Bottom two traces - Constant current ramp method. Traces 1 and 3 - Actual flux-linkage. Traces 2 and 4 - Flux-linkage demand. Torque 30Nm, speed 150rpm.

This is further illustrated in Figure 8.10 which compares the measured controller flux responses for these two methods. Even at the relatively low speed of 150rpm the actual flux is unable to track the rapid fall required by the minimum RMS current method, but is able to track the slower fall of the constant current ramp method. As speed or torque demand is increased there is a point when the controller can no longer produce the required rate of change of flux. The flux-linkage profile at either end of the constant torque period then become linear ramps, as the controller output is saturated.

Under these conditions the phase flux-linkage profile can be approximated by the profile shown in Figure 8.11: with a linear ramp at either end of the conduction period, and linear ramps following the constant torque characteristic. Thus the flux-linkage reference for constant torque operation is composed of a series of linear 'flux ramps'. Neglecting the resistive voltage drop each linear flux ramp is achieved by application of a constant voltage to the phase. The gradient of each flux ramp determines the maximum speed at which it is possible to track the constant torque flux characteristic for a given dc-link voltage.

The single-tooth winding flux-linkage profile for constant torque is defined by five points (marked with circles in Figure 8.11) and the linear flux-linkage ramps which connect them. The period from θ_{adv} to θ_d could be considered the phase conduction period as ψ_{off} is usually zero. The main torque producing period is from θ_a to θ_c , during which time the flux demand follows the machine constant torque flux-linkage characteristic. Commutation of torque production between phases occurs mainly over the periods θ_{adv} to θ_a , and θ_c to θ_d , although as speed increases commutation extends beyond these intervals.

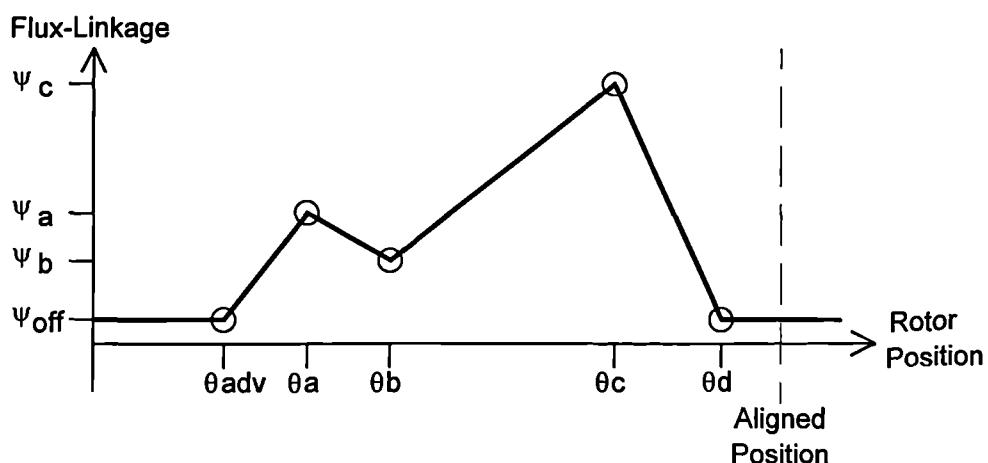


Figure 8.11 Flux-linkage profile composed of constant flux ramps.

The flux-linkage profile shown in Figure 8.11 is of a typical low speed motoring waveform, and has a similar shape to that of the constant current ramp profile shown in Figure 8.9. When producing negative torque for braking or generating, the flux-linkage profile is reflected about the aligned position. In this case the beginning of the conduction interval (in time) is from θ_d to θ_c , rather than from θ_{adv} to θ_a . For constant shaft torque operation the flux-linkage points must be chosen to meet several criteria:-

- The flux-linkage at any point must not correspond to a phase current greater than the maximum rated value.
- The rate of change of flux-linkage between successive points must not exceed the maximum value possible for a given speed and dc-link voltage (equation 8.1).
- The sum of torque produced by all machine phases should minimise the RMS torque ripple whilst maintaining the required average torque.

This is the basis of constant torque operation using 'flux ramps'. This method can be used for fully-pitched winding switched reluctance drives and single-tooth winding switched reluctance drives. Flux-linkage reference design is always done in single-tooth winding parameters as this enables the phases to be de-coupled. For fully-pitched winding switched reluctance drives parameters must be transformed from the single-tooth winding parameters to calculate the actual phase currents and voltages. An additional complication with the application of flux-ramps to fully-pitched winding switched reluctance drives is that the phase voltages are no-longer de-coupled, this makes checking of the flux-linkage ramp rate more complex.

8.4.3 Simulation

Several programs have been written to simulate constant flux ramp operation. The 'C' language was used and programs run in a DOS environment. Programs were designed to simulate the operation of fully-pitched winding or single-tooth winding drives, and use flux and torque data for the experimental machines. The various control angles and flux-linkage values are entered (as in Figure 8.11) and from this the programs work out the phase current, voltage and torque. Resistive voltage drop is neglected when calculating the phase voltage as this only accounts for about 2% of dc-link voltage. Results are displayed graphically with numerical output of key parameters such as peak current, average torque and torque ripple. Results can also be saved to a data file for post-processing.

Figure 8.12 shows simulated operation with the following flux ramp parameters:-

$$\begin{bmatrix} \Psi_{off} \\ \Psi_a \\ \Psi_b \\ \Psi_c \end{bmatrix} = \begin{bmatrix} 0.0 \\ 0.38 \\ 0.34 \\ 1.06 \end{bmatrix} \text{ (Wb)} \quad \text{and} \quad \begin{bmatrix} \theta_{adv} \\ \theta_a \\ \theta_b \\ \theta_c \\ \theta_d \end{bmatrix} = \begin{bmatrix} 185 \\ 20 \\ 35 \\ 80 \\ 50 \end{bmatrix} \text{ (Electrical Degrees)}$$

Initial estimates for the flux-linkage points were based on the constant current ramp flux-linkage profile for 30Nm, as shown in Figure 8.9. The points were then adjusted to minimise torque ripple. Average torque is 29Nm with a torque ripple of $\pm 3\text{Nm}$ peak, 1.9Nm RMS.

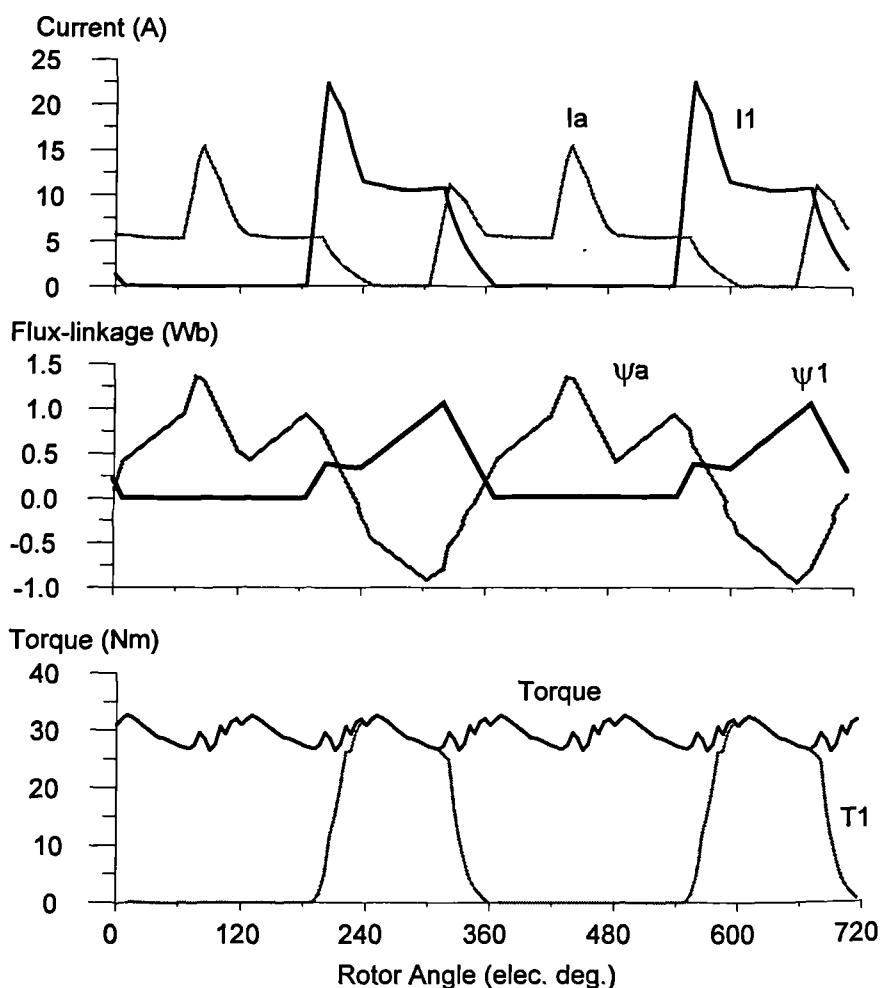


Figure 8.12 Simulated constant flux ramp operation.

8.4.4 Systematic Search

Using the flux ramp simulation program it is possible to manually optimise the flux ramp parameters by trial and error to get less than 10% RMS torque ripple. However, improvements beyond this are difficult to achieve by trial and error due to the large number of parameters. With nine parameters to optimise and many local optima it is not easy to predict which parameter to adjust and by how much. For this reason a systematic search is used to fine tune parameters. With the large number of parameters even trying just two points either side of the present position for each variable would require the current and torque during a cycle to be evaluated 1,953,125 (5^9) times, this would take a considerable amount of processing time. To reduce the number of calculations required the simulation programs can optimise just a few parameters at a time, for example optimising just the angles at two points either side of the present position requires 3,125 sets of calculations. However, just optimising some of the parameters over a small range is unlikely to find the global optimum solution, and even the reduced search takes a significant amount of time.

Although a simple systematic search can help improve on a good starting estimate it is too slow and searches too small an area to find the optimum solution. What is required is an algorithm that can search a large space in an efficient manner.

8.4.5 Genetic Algorithm Search

An optimisation method that is able to make a directed search in an arbitrarily large search space is a Genetic Algorithm (GA). Genetic Algorithms originated from work by Holland [8.16] and are based on the principles of natural selection. For those unfamiliar with Genetic Algorithms Appendix E gives an introduction to their origins, terminology and theory of operation. Further reading can be found in Rawlins [8.17], Whitley [8.18] and Buckles *et al* [8.19].

Many of the terms used in Genetic Algorithms come from natural genetics such as population, generation, chromosome, breeding and mutation. A GA attempts to find the optimum solution to a problem by evaluating a set of possible solutions (a population of organisms) using some fitness function. The best solutions are then combined together in a breeding process in an attempt to find a better solution (the next generation). Mutation adds a random element to the possible solutions which adds genetic diversity enabling the search to move to new areas within the search space. The breeding and mutation process is repeated until an acceptable solution is found or some other terminating condition met. Figure 8.13 illustrates this process.

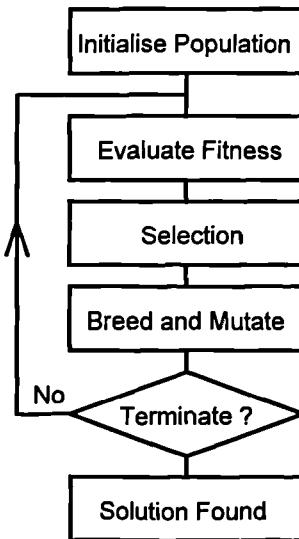


Figure 8.13 Genetic Algorithm flow chart.

A Genetic Algorithm is essentially an empirical method working on a trial and error basis. For the problem of optimising flux ramps for constant torque operation GAs have several attractive features. The most important aspect is that they are able to search a large area effectively. This is achieved because each organism covers a different area of hyperspace (multi-dimensional search space) and each successive generation can move to new areas. In evaluating the fitness of each organism information is gathered from many areas of hyperspace which enables the search to be directed. By directing the search using information gathered from many areas of hyperspace optima can be located quickly, thus requiring less computation time than other methods. Random mutation of organisms moves the search area to regions not explored by breeding alone. The second most important aspect of GAs is that local maxima can be located accurately whilst covering a large search space. Mutation plays an important role in this as small mutations can move a solution already near an optimum a little closer, whilst larger mutations allow the search to explore large areas of space.

Just as a Genetic Algorithm is empirical in its operation, selection of GA parameters such as population size, degree of mutation, etc. are also made on an empirical basis. The details of GA operation outlined below were arrived at through extensive testing during which the development of the population through successive generations was monitored.

The Population

The population is made up of a set of possible solutions or organisms. Each organism is represented by a chromosome (vector) having nine alleles (elements) which correspond

to the points on the flux-linkage graph Figure 8.11. Expressed as a vector it has the following format:-

$$U(n) = \{ \psi_{off}, \psi_a, \psi_b, \psi_c, \theta_{adv}, \theta_a, \theta_b, \theta_c, \theta_d \}$$

The four flux-linkage alleles have an alphabet (set of possible values) bounded by $0.0Wb \leq \psi \leq 2.0Wb$ in steps of approximately $3mWb$, and the five angle alleles have an alphabet of all integer values such that $0^\circ \leq \theta \leq 359^\circ$. The resulting search area is a nine dimensional hyperspace with 9.5×10^{23} points within it (the experimental hardware is in fact capable of measuring 3.48×10^{41} points).

Selection of population size is a trade off between optimisation performance and the time required to find an acceptable solution. If a population is too small it tends to contain little genetic diversity (variety of possible solutions) which leads to premature convergence possibly at a sub-optimum solution (Grefenstette [8.20]). A large population contains greater genetic diversity and thus searches more areas of hyperspace, it is therefore more likely to find a global optimum and less likely to converge prematurely. A large population does however require more processing time for each generation and hence the solution is found more slowly. In work by Grefenstette [8.20] and Goldberg [8.21] optimum population size is found to be between 30 and 100 or 200. For flux ramp optimisation a population size of 40 organisms in each generation was found to give sufficient genetic diversity without requiring excessive processing time.

Initialisation of the First Generation

It is possible to start optimisation with a completely random first generation. However with the large search space involved good results are obtained much quicker if the first generation gives approximately the desired torque. As machine torque characteristics are known it is possible to use these to generate the first generation. Values from machine data tables are taken near the aligned position and two near the unaligned position to give estimates for the flux values ψ_a , ψ_b and ψ_c at angles θ_a , θ_b and θ_c . The maximum flux ramp rate for a given speed and voltage is then used to calculate angles θ_{adv} and θ_d from the flux values ψ_a and ψ_c respectively. These flux and angle values are then randomised to produce the first generation by adding a random element to every allele. A quarter of the population has a small random element added to the flux and angle estimates, with the remainder of the population having a much larger random element. This scheme preserves some of the population with initial values near to the machine characteristics which may prove to be near optimum, but also gives a sufficiently wide coverage to find other solutions quickly.

Evaluation of Fitness

The main criteria used to evaluate the fitness of an organism is RMS torque ripple expressed as a fraction of the average torque. For each organism the torque is evaluated at m points in a cycle, with the average and RMS torque calculated using the following equations:-

$$\bar{T} = \frac{\sum_{n=0}^m T_n}{m} \quad (8.2)$$

$$T_{RMS} = \sqrt{\frac{\sum_{n=0}^m T_n^2}{m}} \quad (8.3)$$

The RMS torque ripple can be found using the following equation:-

$$T_{Ripple} = \sqrt{\frac{\sum_{n=0}^m (T_n - \bar{T})^2}{m}} \quad (8.4)$$

By expansion of the sum and simplification of the resulting terms the equation can also be written as:-

$$T_{Ripple} = \sqrt{T_{RMS}^2 - \bar{T}^2} \quad (8.5)$$

The fitness of an organism is found from the following equation:-

$$Fitness = \frac{T_{Ripple}}{\bar{T}} \quad (8.6)$$

The ideal solution is one with zero torque ripple, the fittest organism is therefore the one with the lowest proportion of torque ripple.

There are three further criteria used to rate the fitness of an organism, but these are really used to reject organisms rather than grade them. If the required flux-linkage would result in a phase current greater than the maximum allowed, or if the rate of change of flux-linkage is greater than the maximum possible, or if the average torque is too high or low, then the fitness of that organism is set to a value which will force that organism to the bottom of the pile and it will be dropped from the population in the next generation.

Rather than just using torque ripple to assess the fitness of an organism, better overall solutions may be found if the RMS phase current is considered as a factor. This allows a waveform with much lower losses but only very slightly worse torque ripple to be

selected for the next generation. Also having a penalty proportional to the phase current overload or average torque error would prevent the rejection of organisms only just outside the limits, but which with a little modification may provide the best solution. Although taking these other factors into account would undoubtedly improve the overall solution, the fundamental problem is knowing what weight to give each factor so that the primary goal of low torque ripple is still met. Further work in this area is outside the scope of this thesis but would merit future investigation. It would be possible to use another GA to help find the optimum weights for the main GA fitness criteria. This principle has been used by Grefenstette [8.20] to find the optimum GA population size, mutation rate and the like.

Breeding and Selection

The selection strategy used is an elitist one in combination with a modified liner rank scheme. An elitist strategy means that the fittest organism in the current generation survives intact into the next generation (Grefenstette [8.20]) which ensures that the best solution to date is not lost through mutation. A linear rank selection strategy is employed because it requires less computation than proportional selection, it also has the advantage that it is less prone to premature convergence (Grefenstette *et al* [8.22]) caused by organisms which are much fitter than other members of the population. Linear rank selection requires each organism to be selected with a probability proportional to its position in an ordered list. This is modified slightly to reduce the amount of computation required and simplify implementation. A linear rank is still employed, i.e. organisms are sorted by fitness, but a simplified selection strategy is employed. Half of the next generation is produced by breeding of the eight fittest organisms, each selected with equal probability. The other half of the next generation is bred from the top 32 organisms, again each with equal probability. In this way the 8 least fit organisms from each generation have no input into the next generation.

Mutation

With little mutation the search remains in the nearby area and this helps in homing in on local maxima. A large amount of mutation makes the search more random and can take the search into new areas of hyperspace, this helps in finding global maxima by preventing all organisms converging to a local maxima. The degree of mutation needs to balance these two requirements (Buckles *et al* [8.19]). In practice several levels of mutation are employed in an attempt to optimise the search. The top four organisms have no mutation applied which ensures that the best solution found so far is not lost. Three levels of mutation are then applied to the remaining organisms, a small mutation to pinpoint local maxima, a medium amount of mutation to find nearby maxima, and a

large mutation to move the search to completely new areas of hyperspace. Each level of mutation is applied to one quarter of the population, with the smallest mutation being applied to the fittest individuals. Only one allele is mutated in each organism.

Terminating Condition

Although it is possible to run a Genetic Algorithm as a continuous process, for example to track parameter variations over time, it is usual to stop the breeding process when some terminating condition is met. Two popular methods are to run for a set number of generations, or to stop when all organisms in a generation are the same (Buckles *et al* [8.23]). The method employed is to stop after a fixed number of generations as this is simple to implement and the breeding and mutation strategy ensures that not all organisms are the same. A figure of 50 generations was found to be a good compromise. If there are too few generations there is insufficient genetic diversity to get near to the optimum solution, too many generations increases the processing time required and gives little improvement in the solution. This is because there is a law of diminishing returns with successive generations as shown in Figure 7 of Grefenstette [8.20].

8.4.6 An Array of Optimised Flux Ramps

The Genetic Algorithm is able to optimise the flux ramp parameters for a given torque, dc-link voltage and speed, i.e. at a single operating point. When implemented on the experimental drive hardware, using the TMS320C31 DSP to perform the Genetic Algorithm optimisation, it takes about 8 seconds to process 50 generations with exclusive processor use, and in the order of 100 seconds when the processor is also controlling the drive in real time. For steady state optimisation this may be acceptable but for more realistic operation with speed and torque varying many times per second it is clear that the GA cannot optimise operation at each operating point in real time.

To enable real time operation it is necessary to optimise the flux ramps at many operating points off line and store the results so they can be used in real time. Each operating point is defined by three parameters; torque, speed and the dc-link voltage. In terms of the flux ramp parameters there are only really two variables; torque, and the maximum rate of change of flux, because the speed and dc-link voltage combine in equation 8.1. Thus a two dimensional array of solutions can cover all operating points. The two array indexes are *torque*, which covers 0Nm to 50Nm (106% of rated) in steps of 2Nm, and *ramprate* which is defined as follows:-

$$\text{ramprate} = K_{rr} \times \frac{\text{Speed(rpm)}}{\text{DCLinkVoltage(V)}} \quad (8.7)$$

The constant K_{rr} is chosen so that *ramrate* has a range of 1 to 20 (integer values only) which at nominal dc-link voltage covers the speed range 50rpm to 1000rpm. Note that this equation is the inverse of equation 8.1 which ensures a good spread of index values at high speed, where the limit on flux ramp rate is most relevant.

Each element in the array is a data structure which holds the four flux values and five control angles that define a flux ramp waveform. For negative torque the same flux ramps are used but the controller mirrors the control angles about the aligned position.

To generate the two dimensional array of optimised flux ramps every other element is initialised with an estimate based on machine flux-linkage/torque characteristics. Each of these points is then optimised. Interpolation is used to generate initial estimates for the intermediate operating points before every point is optimised using the GA. The resulting 26×21 array of structures, each having 9 elements, is stored in a data file.

8.4.7 Experimental Results

In this section experimental results for constant torque operation using Genetic Algorithm optimised flux ramps is presented. The torque reference is set manually using front panel controls and provides the '*torque*' index value. From measurement of the dc-link voltage and rotor speed the controller calculates the '*ramrate*' index value. The appropriate set of GA optimised flux ramp values are then read from the stored data array. This process is repeated approximately every $500\mu\text{s}$ so the flux ramp values can track speed and torque demand changes rapidly. The flux-linkage controller used is based on the design described in Chapter 7 but the flux-linkage reference calculation is altered. Flux reference values are calculated in real time from the points on the flux waveform shown in Figure 8.11. As each section of the waveform is a linear ramp the calculation involved is relatively straight forward. As with previous results instantaneous flux-linkage and torque are calculated in real time from measured phase currents using stored machine data.

Typical Flux-linkage and Torque Waveforms

Figure 8.14 and Figure 8.15 illustrate the flux ramps produced by the GA optimiser for a variety of operating points. In each case the single-tooth winding flux-linkage and electromagnetic torque are shown, along with the *torque* and *ramrate* index values. The array index values in subsequent figures have the following format:- $\text{Index}(\text{torque_index}, \text{ramrate_index})$. From these index values the actual torque demand, and speed for a dc-link voltage of 480V, can be found using the following equations:-

$$Torque(Nm) = 2 \times \text{torque_index} \quad (8.8)$$

$$Speed(rpm) = 48 \times \text{ramprate_index} \quad (8.9)$$

From the shape of the flux-linkage waveforms it is evident that there is no particular set pattern in the shape of the flux-linkage profiles, or their development with speed and torque. This is due to the random nature of the GA optimisation process which means that two adjacent operating points can have quite different flux profiles although the torque ripple in each case may be very similar.

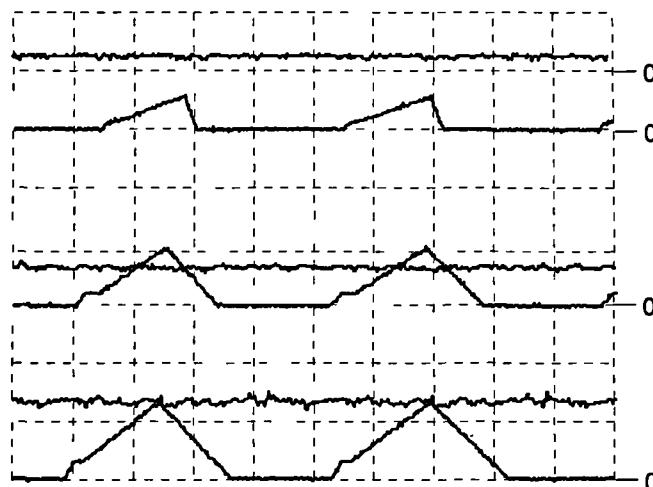


Figure 8.14 Constant torque flux ramps for low speed. Traces show single-tooth winding flux-linkage, 1Wb/div, and electromagnetic torque, 25Nm/div. Traces 1 and 2 - Index(2,2), 20ms/div. Traces 3 and 4 - Index(7,2), 20ms/div. Traces 5 and 6 - Index(17,2), 20ms/div.

Note that the bottom flux-linkage trace in Figure 8.14 which is for Index(17,2), 34Nm and up to 96rpm at nominal voltage, has a long flux ramp down even though operation is at a low speed so the ramp could be much shorter. This is because of all the possible solutions the GA tested this waveform gave the lowest RMS torque ripple. However, it is likely that the GA tested another possible solution with a much faster flux ramp down and only slightly higher torque ripple. A shorter ramp down would reduce the RMS phase current as the effective conduction interval would be shorter. This highlights the problem of having a single fitness criteria, and shows that although the solution is optimum in terms of torque ripple a price is paid in higher RMS current.

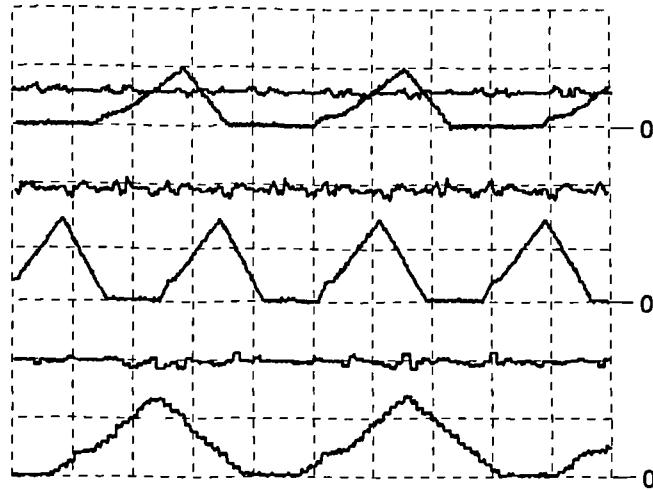


Figure 8.15 Constant torque flux ramps for medium and high speed. Traces show single-tooth winding flux-linkage, 1 Wb/div, and electromagnetic torque, 25 Nm/div. Traces 1 and 2 - Index(7,8), 5ms/div. Traces 3 and 4 - Index(23,6), 10ms/div. Traces 5 and 6 - Index(24,18), 2ms/div.

Constant Torque Motoring

Low speed operation at maximum torque is shown in Figure 8.16. Average torque is 50Nm (index value 25) at 110rpm with a dc-link voltage of 480V (index value 3). The long tail in the phase current is clearly evident and is due to the long ramp down in flux-linkage. Note that this long ramp down of flux and current approaching the aligned position makes machine operation relatively quiet.

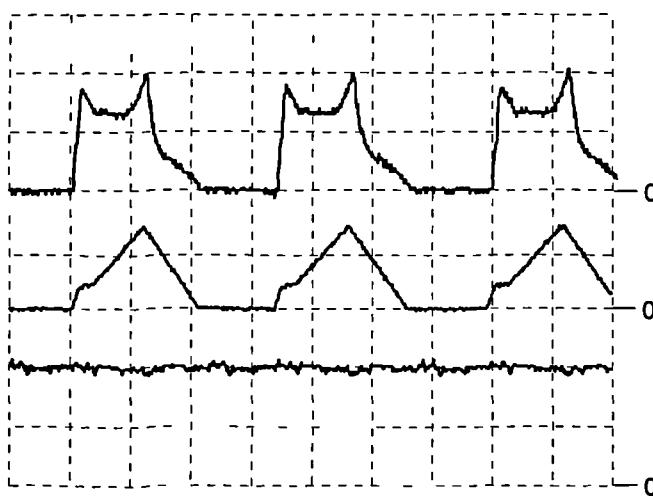


Figure 8.16 Operation at 110rpm and 50Nm with GA optimised flux ramps. Trace 1 - Phase current 20A/div. Trace 2 - Flux-linkage 1Wb/div. Trace 3 - Torque 25Nm/div. Horizontal 20ms/div.

Results in Figure 8.17 are for operation at 140rpm with an average shaft torque of 15Nm. If the phase current and flux-linkage profiles are compared with those in Figure 8.16, it is evident that although the speed in each case is similar the flux-linkage profiles are very different. When producing 15Nm the flux has a short ramp down at the end of the torque producing period, and a large current peak at the start of the period. For 50Nm it is not possible to simply 'scale up' this waveform as the peak current at the start of the period would exceed the power converter rating. The inability to produce most of the required torque early in the cycle means that the short fall must be made up at the end of the period, hence the long flux ramp down evident in Figure 8.16.



Figure 8.17 Motoring at 140rpm and 15Nm with GA optimised flux ramps. Trace 1 - Phase current 20A/div. Trace 2 - Flux-linkage 1Wb/div. Trace 3 - Torque 25Nm/div. Horizontal 20ms/div.

Constant Torque Generating

Figure 8.18 shows operation as a generator at the higher speed of 530rpm. As the required torque opposes the direction of rotation teeth are excited after they have passed through the aligned position. Peak flux now occurs at the start of the waveform with the current profile the 'mirror' of motoring operation. When producing negative torque the flux-linkage controller uses the positive torque flux reference for that magnitude of torque but subtracts the measured rotor position from a constant so that the rotor angle appears to run backwards and hence the reference waveform is reversed. Note that the slightly higher dc-link voltage whilst generating is automatically taken into account by the flux controller in selection of the *ramprate* index and controller gain.

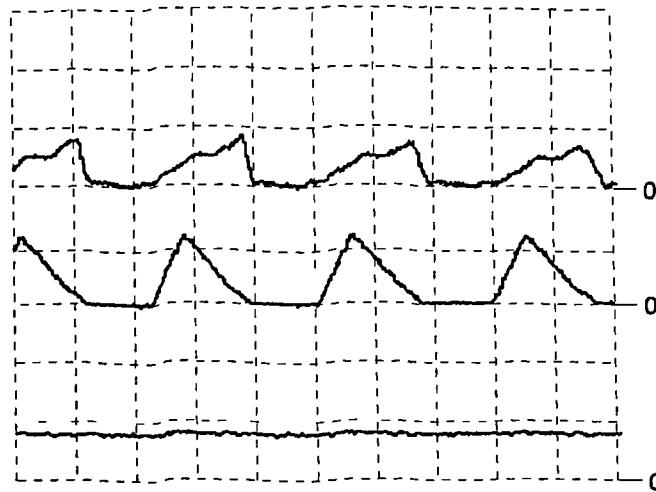


Figure 8.18 Generating at 530rpm and 20Nm with GA optimised flux ramps. Trace 1 - Phase current 20A/div. Trace 2 - Flux-linkage 1Wb/div. Trace 3 - Torque 25Nm/div. Horizontal 5ms/div.

Looking at the flux profile in Figure 8.17 it is evident that the flux ramps at either end of the conduction period in the upper trace are quite short, whilst in Figure 8.18 they are long, so that the final ramp down almost merges into the rest of the ramp down from peak flux. This is because as speed increases the rate of change of flux over the main torque producing interval is almost equal to the maximum possible rate, it is thus not possible to ramp the flux down any faster at the end of the conduction interval. As with motoring, the generating torque response shows that the shaft torque is almost constant with rotor position

8.5 A Speed Controller

The array of constant torque flux ramps produced by the GA optimiser enable a known torque to be produced over a wide speed range. This opens up the possibility of implementing a speed controller which produces a torque demand output, with the phase currents determined by the constant torque flux ramps. A linear relationship between demand and actual torque enables standard linear design and analysis methods to be employed in formulating the speed controller.

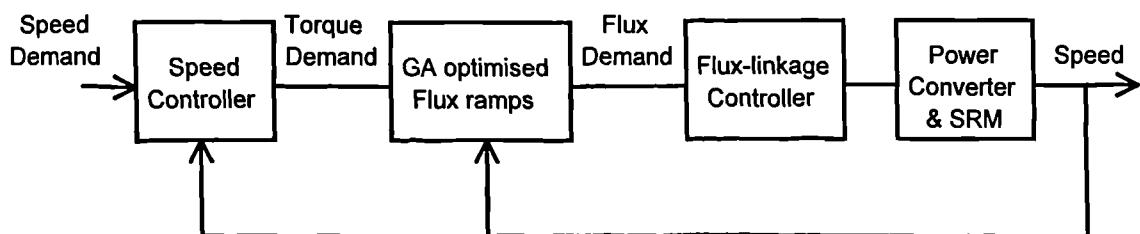


Figure 8.19 Speed control system.

Figure 8.19 is a simplified diagram of the speed control system. A torque demand is calculated by the speed controller from the actual and demand speeds. Torque demand and speed are then used to select a particular flux ramp profile that will give the required torque at the operating speed. This provides the flux reference input to the flux-linkage controller which controls the power converter so as to achieve the desired machine flux. The design of the speed control system is described in the following sections with software implementation details given in sections D.7 and D.8 of Appendix D.

8.5.1 System Model and Controller

The speed control system has three main elements; the load, the speed measurement system and the controller. Each of these will be explored in the following sections.

Load Model

For speed controller tests the DC load machine armature was connected to a load resistor rather than the Ward-Leonard set. This allows speed to vary with load torque over a wide range. The Laplace transfer function for the first order mechanical system model is derived in Appendix B section B.4, and is:- .

$$\frac{\omega}{T_{srm}}(s) = \frac{1}{K_l + sJ} \quad (8.10)$$

Experimental measurements of the speed response to a torque step input showed that the load time constant is approximately 375ms. At nominal field current and load resistance the steady state gain is approximately 22rpm/Nm.

Speed Measurement

Speed is measured by timing the period of pulses from the rotor position encoder as described in Appendix D.7. The mechanical system has several resonant frequencies at which there is considerable torsional twisting in the motor coupling shafts. This results in the apparent speed (as measured by the encoder) having an oscillation superimposed on the true average speed. At the main shaft resonance peak to peak speed oscillation is approximately 50% of the true speed. Clearly this size of oscillation is not acceptable for a speed controller feedback signal. It was therefore decided to low pass filter the measured speed. The filter is implemented digitally and has a 'z-domain' transfer function of:-

$$H(z) = \frac{1 - e^{-\omega T}}{1 - e^{-\omega T} z^{-1}} \quad (8.11)$$

The cut off frequency is ω and the sample interval T . Implemented as a difference equation the formula is:-

$$y_n = e^{-\omega T} \cdot y_{n-1} + (1 - e^{-\omega T}) \cdot x_n \quad (8.12)$$

It was found that a filter time constant of 10ms, giving a cut off frequency of 100rad/s, gave a good compromise between ripple attenuation and speed of response.

Speed Controller

A simple proportional plus integral (PI) controller was selected for the speed controller. An additional pole rolls off the response at high frequency to reduce sensitivity to noise and speed measurement ripple. Figure 8.20 shows the general gain response required from the controller.

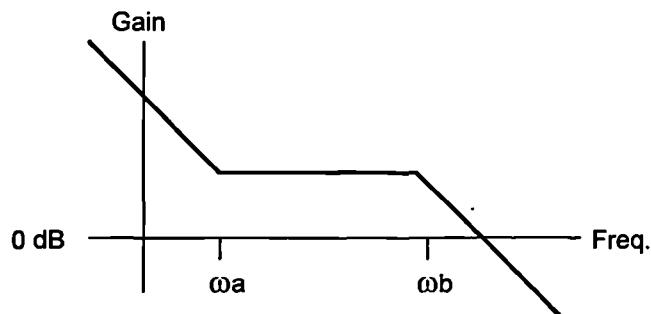


Figure 8.20 Desired speed controller gain characteristics.

The 's-domain' transfer function of the gain characteristic in Figure 8.20 is as follows:-

$$D(s) = K_p \cdot \frac{1}{s} \cdot \frac{\omega_a + s}{\omega_a} \cdot \frac{\omega_b}{\omega_b + s} \quad (8.13)$$

The discrete time equivalent with a sample time of T is:

$$D(z) = K_p \cdot \frac{1 - K_b}{1 - K_a} \cdot \frac{(1 - z^{-1} K_a)}{(1 - z^{-1})(1 - z^{-1} K_b)} \quad (8.14)$$

Where $K_a = e^{-\omega_a T}$ and $K_b = e^{-\omega_b T}$. Expressed as a difference equation the controller discrete time transfer function $D(z)$ is:-

$$y_n = y_{n-1} \cdot (1 + K_b) - y_{n-2} \cdot K_b + (x_n - x_{n-1} \cdot K_a) \cdot \left(\frac{K_p \cdot (1 - K_b)}{1 - K_a} \right) \quad (8.15)$$

The Closed Loop System

Combining the load, measurement and controller sections gives the system block diagram illustrated in Figure 8.21

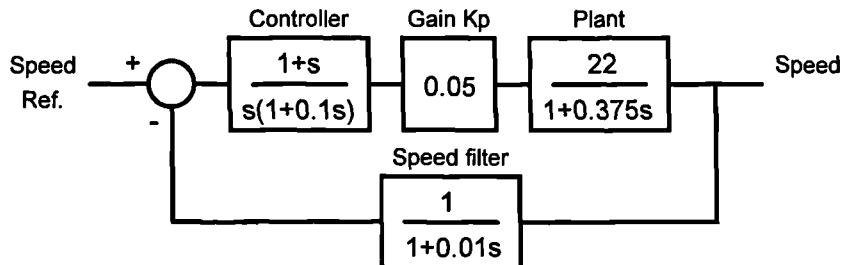


Figure 8.21 Speed control system block diagram.

The load gain and time constants used are the nominal values, with the controller gain and time constants selected with the aid of simulation using Matlab™ and Simulink™. Controller time constants were selected to be 1s and 0.1s, with a proportional gain term of 0.05. The resulting open loop gain and phase responses are plotted in Figure 8.22.

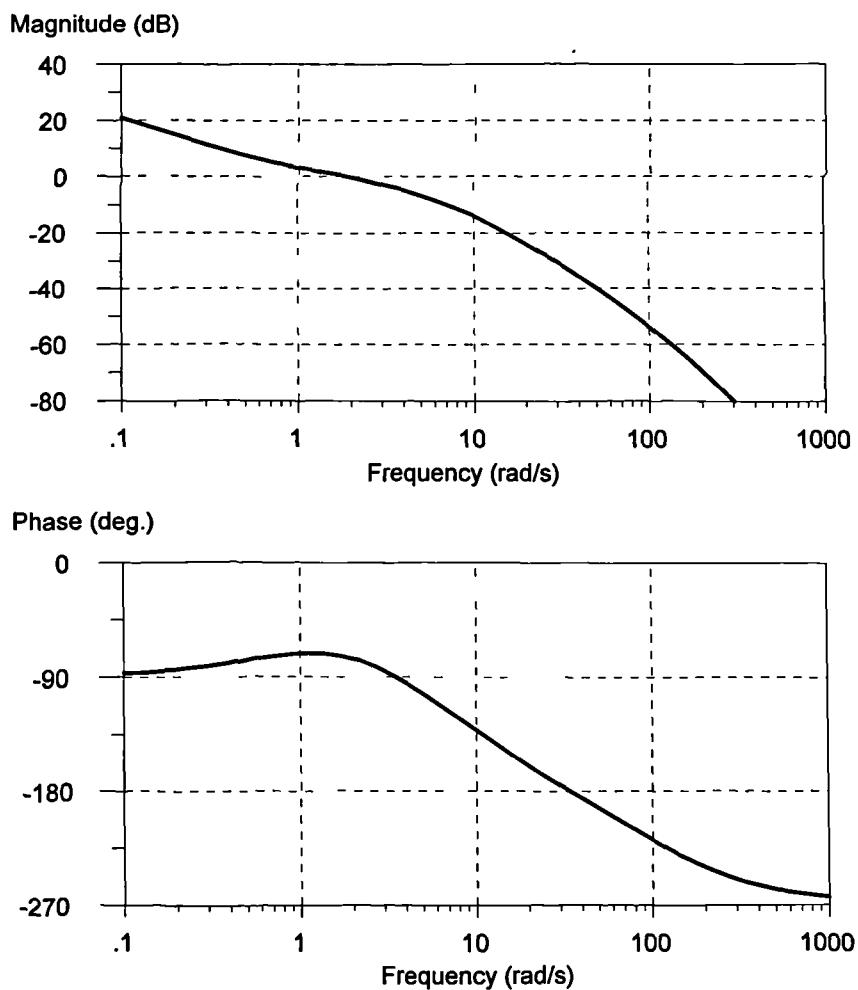


Figure 8.22 Open loop system gain and phase response.

Note that the bandwidth is limited by the load time constant of 375ms, which is due to the large load machine inertia. The relatively low gain was chosen because although the nominal load gain is 22, this increases at light loads to 130 or more reducing stability margins. For stability on minimum load nominal load response is compromised.

8.5.2 Experimental Results

The effects of the relatively low gain are most evident when step changes are made to the speed demand or load torque: the response is slow and little use is made of the available torque. The top two traces in Figure 8.23 show the simple PI controller speed and torque response. With the load set to 50% of maximum torque and the speed at approximately 300rpm the direction is reversed twice before reducing the load torque to minimum and then finally restoring the load torque to 50% of maximum. With the PI controller as the speed demand is reversed there is an initial peak of near maximum torque (50Nm) but this drops off very quickly resulting in a slow change in speed. Response to load step changes is quite good with only small perturbations in speed.

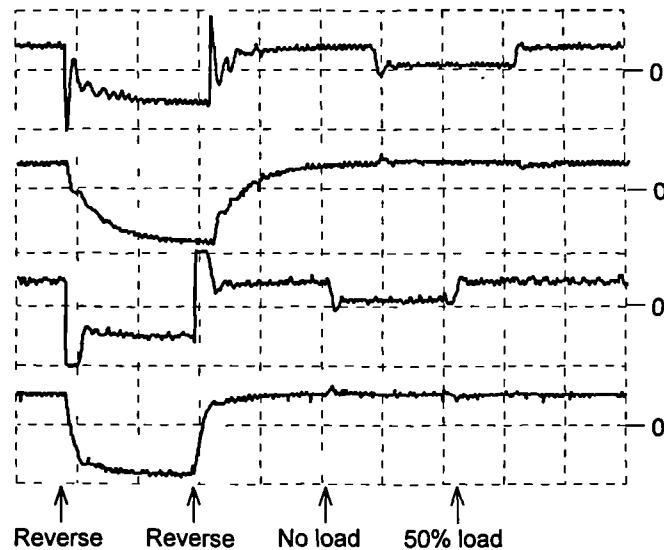


Figure 8.23 Speed controller response to demand and load step changes. Traces 1 and 2 - Non-adaptive controller. Traces 3 and 4 - Adaptive controller. Trace 1 - Torque, 50Nm/div. Trace 2 - Speed, 500rpm/div. Trace 3 - Torque, 50Nm/div. Trace 4 - Speed, 500rpm/div. Horizontal 2s/div.

The lower two traces in Figure 8.23 show the speed and torque response for an adaptive controller. This controller was developed to overcome the short comings of the simple PI controller caused by the low gain necessary to maintain stability. The adaptive controller uses a PI gain block when the speed error is less than 100rpm: this ensures a stable response for normal linear operation. When the speed error exceeds 100rpm a proportional only controller is used with a gain of 10 times that of the PI controller

proportional gain. With this gain an error of 100rpm or more produces a torque demand of 50Nm, the maximum torque. As the controller response is saturated under these conditions the response is stable, note however that this gain would produce an oscillatory response at light loads in the linear mode of operation. There is a clear improvement in torque and speed response using the adaptive controller as shown in Figure 8.23. When the demand speed reverses full torque is applied and maintained resulting in a rapid change of direction. The size of the load step change in Figure 8.23 is not sufficient to cause a transient speed error of greater than 100rpm, and so the PI controller maintains control throughout these transients. As a result the load step change response of the adaptive controller is identical to the simple PI controller response.

8.6 Summary and Conclusions

Many authors have proposed methods to overcome the problem of torque non-linearity with respect to phase current and rotor position in switched reluctance machines. Some of these methods have been reviewed, with simulated and experimental results used to compare low speed performance. Although the methods reviewed work well at low speed, little account has been taken of high speed performance. Further simulation work has shown that in order to produce constant torque over a wide speed range account must be taken of what excitation will be required significantly further on in a cycle.

It is relatively easy to predict performance later on in a cycle when using a flux-linkage controller, as flux-linkage is simply equal to the applied volt-time integral, neglecting resistance. Examination of machine flux-linkage/position characteristics for constant torque shows them to have a linear rise in flux over much of a cycle. These observations lead to the concept of using a flux-linkage waveform composed of a series of linear ramps for constant torque operation. By limiting the flux ramp rate, reference waveforms have been calculated to give constant torque over a wide speed range. The large number of possible flux-linkage waveforms makes conventional search algorithms unsuitable for the task of finding the optimum flux ramp reference waveforms. A Genetic Algorithm was therefore employed as it can search a large search space effectively. It has been used successfully to find optimum flux profiles for constant torque operation over a 20 to 1 speed range, and 25 to 1 torque range. Experimental results have shown the Genetic Algorithm optimised flux-linkage waveforms give the desired torque characteristics. A speed controller has been implemented that uses an array of optimised flux profiles to give a linear torque response. Linear design methods were used in formulating the controller which gives a stable but relatively slow response. By dynamic gain changing in response to large speed errors the transient response has been considerably improved.

Chapter 9

Comparison of Operating Modes

9.1 Torque Speed Characteristics

Torque/speed characteristics for the experimental drive are shown in Figure 9.1. For each of the four current waveforms the current reference is 10A RMS, and the dc-link voltage is 420V. It is evident that at 100rpm, where the actual currents differ very little from the reference values, each waveform gives approximately the same average torque. The characteristic for trapezoidal conduction falls faster than other modes partly because the reference conduction width is fixed at 120° per half cycle, whereas for unipolar operation the advance angle and conduction width are both altered as speed increases in order to optimise torque output. At speeds of 500rpm and above unipolar operation gives more torque than the other modes. The reasons for this are explained in Chapter 6 section 6.3.2, and will not be repeated in full here. Suffice to say that the low torque is essentially a function of the current reference, and given a suitable reference waveshape for the operating speed, the same torque can be produced by unipolar or bipolar operation. Note however that this complicates controller design, and bipolar currents in general require twice as many switching devices as unipolar operation.

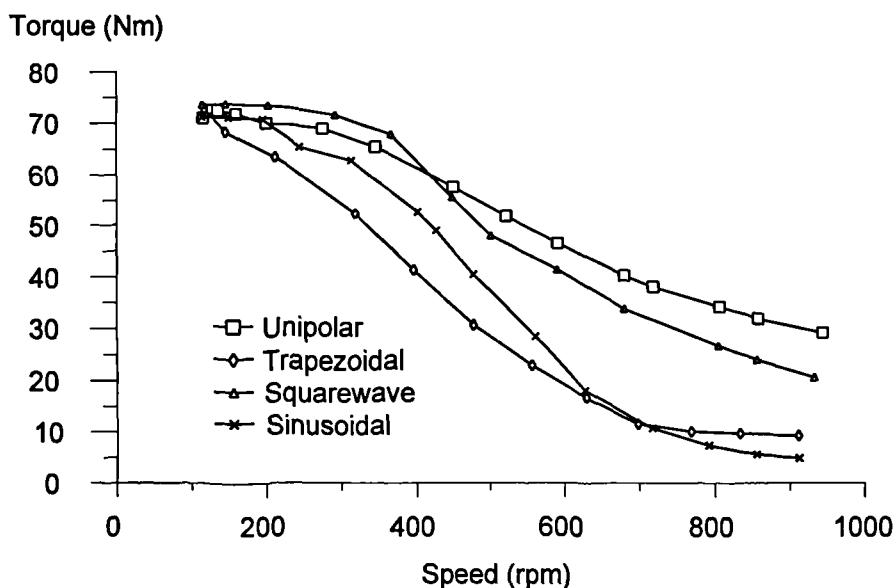


Figure 9.1 Motoring torque speed envelopes for 10A RMS phase currents.

9.2 Sensitivity to Control Angles

Torque production in switched reluctance drives is very dependant on correct synchronisation of current to rotor position as shown in Figure 9.2, which is for unipolar operation. The two curves are simulated values for low speed operation with ideal current control, and high speed operation under voltage control. In each case 0° corresponds to the advance angle which gives maximum average torque. These results show that in both cases torque falls by 10% for errors of about 20° , and by 50% for errors of 40° to 50° .

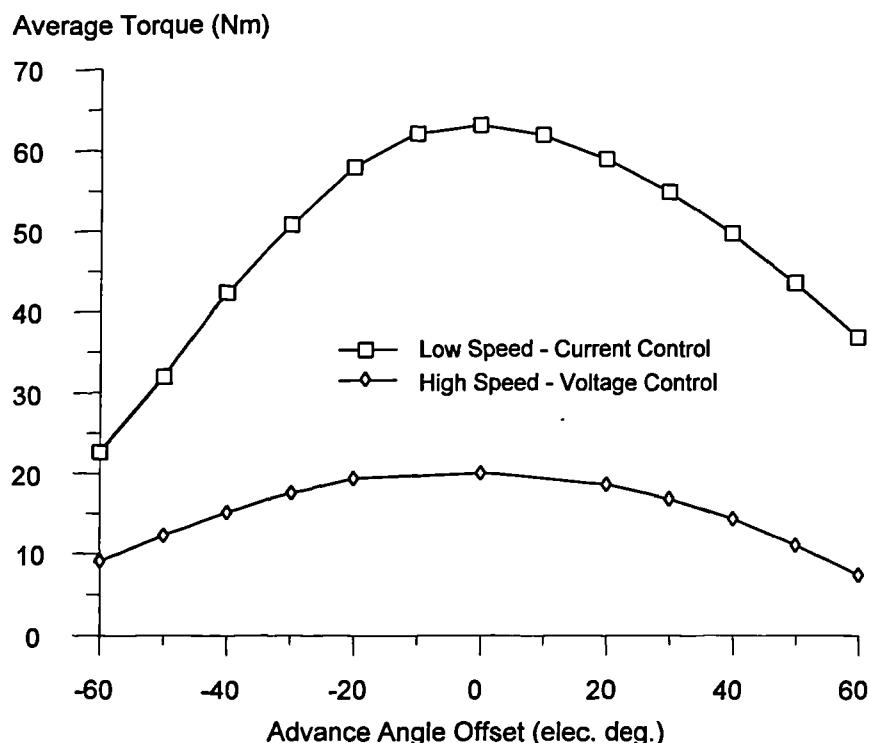


Figure 9.2 Simulated data showing the variation in torque with advance angle. Low speed - Current control, Unipolar, 10A peak, optimum advance angle 20° . High speed - Voltage control, unipolar, 500V, 1500rpm, optimum advance angle 280° .

Figure 9.3 shows measured values for the variation in average torque with advance angle for unipolar and sinusoidal operation. In each case phase reference currents are 10A RMS and the dc-link voltage is 420V. From these results it is evident that the effect of varying advance angle at low speed does not depend on the current waveshape. The overall trend also corresponds well with the simulated values.

For unipolar operation the conduction angle (nominally 240° at low speed) can be varied as well as the advance angle. From the curves in Figure 9.4 it is evident that the effect of varying pulse width is almost the same as varying advance angle.

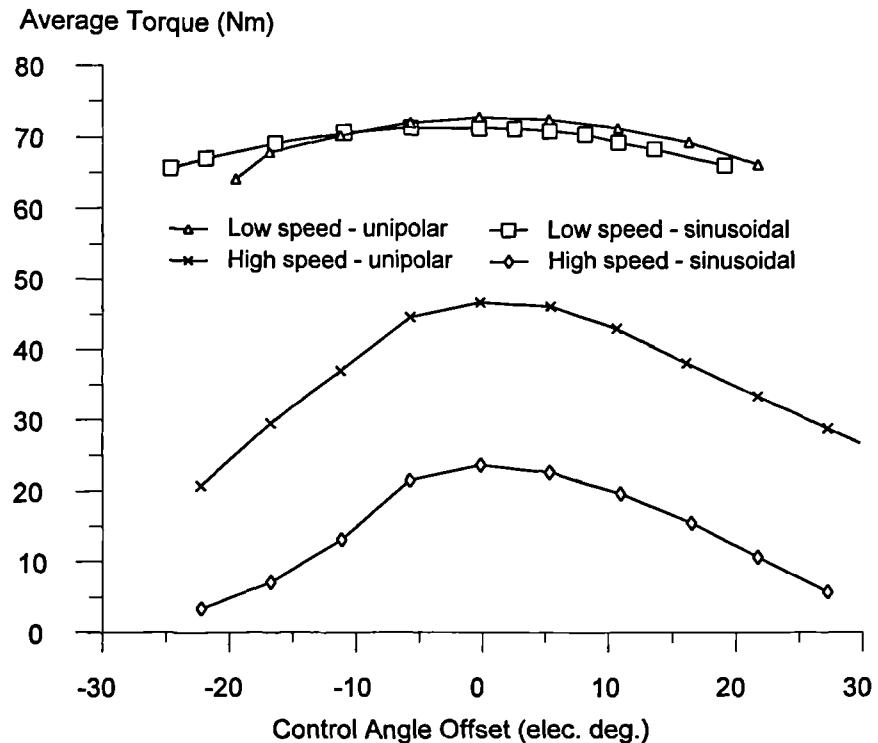


Figure 9.3 Measured variation in average torque with advance angle for unipolar and sinusoidal operation. Both modes are for 10A RMS, 420V. Low speed - 130rpm, high speed - 580rpm.

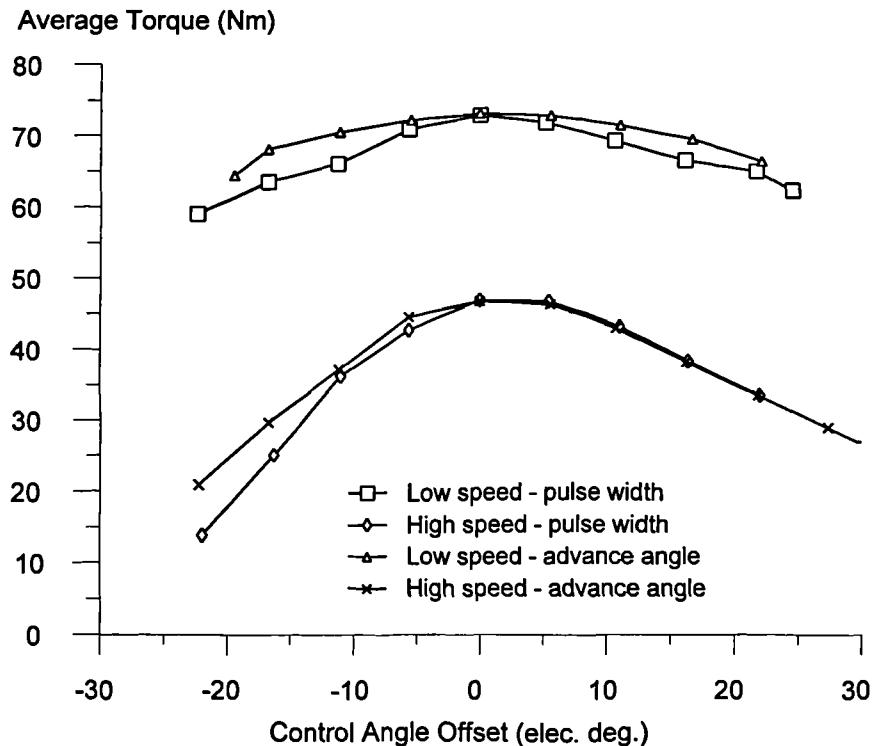


Figure 9.4 Measured variation in average torque with pulse width and advance angle. Unipolar operation, 10A RMS, 420V. Low speed - 130rpm, high speed - 580rpm.

Constant torque operation is very sensitive to rotor position errors in terms of peak to peak ripple. The experimental drive uses a 1024 position encoder which has a resolution of 0.35° mechanical, 2.8° electrical. Simulation shows that an error of just 3° electrical, increases the peak to peak torque ripple by between 130% and 230% for constant current ramp and constant torque ramp modes.

9.3 Torque Ripple and Torque per Unit Loss

In this section figures are compared for four 'standard' modes of operation and two constant torque modes of operation. The four standard modes are; unipolar, trapezoidal, squarewave and sinusoidal conduction. Constant current modes are; constant torque ramp and constant current ramp, as described in Chapter 8 section 8.2.

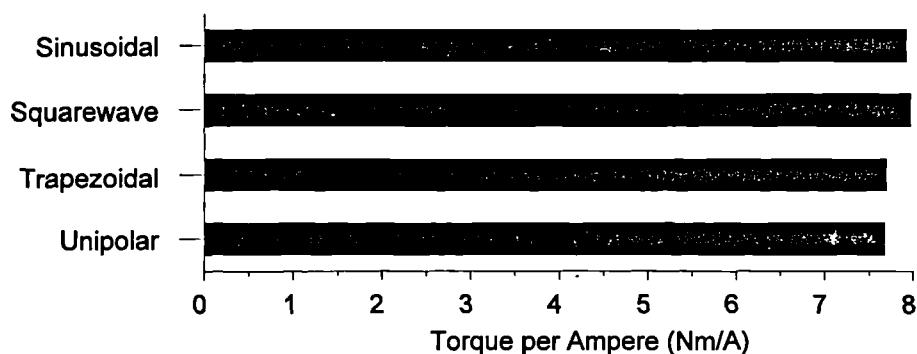


Figure 9.5 Torque per RMS ampere for standard modes.

Figure 9.5 shows figures for low speed torque per RMS ampere of phase current. The phase current for each waveform is 8A RMS, which gives an average torque of approximately 63Nm. It is evident that the torque per ampere for each waveform is almost the same, hence in terms of machine copper loss no waveform offers any significant advantage.

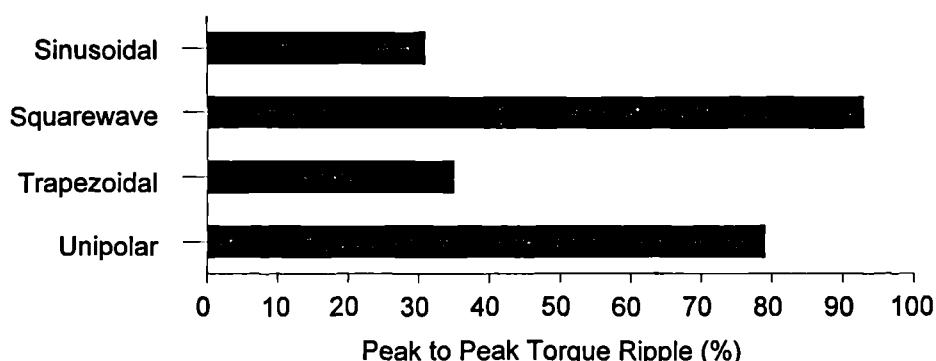


Figure 9.6 Peak to peak torque ripple for standard modes.

Peak to peak torque ripple, expressed as a percentage of the average torque, for the previous conditions is shown in Figure 9.6. This shows that sinusoidal, and trapezoidal operation have considerably less torque ripple.

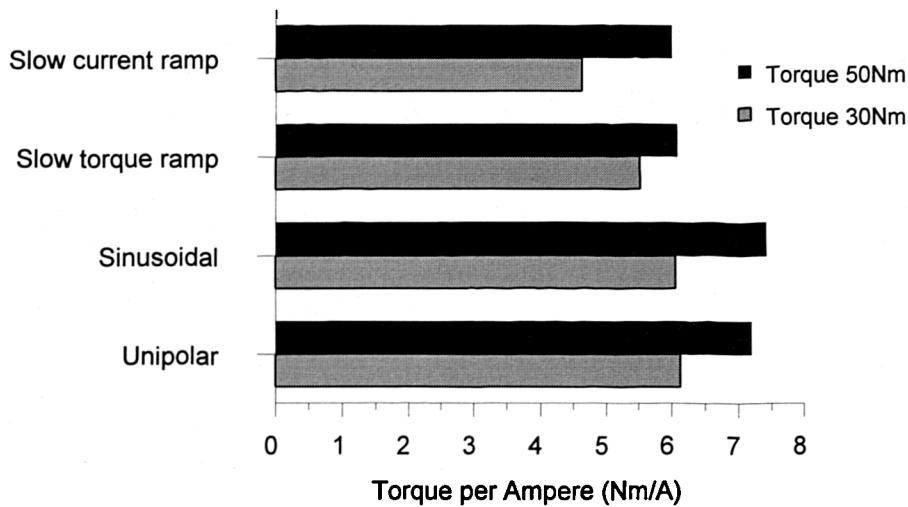


Figure 9.7 Torque per RMS ampere for standard and constant current modes.

If results for the standard waveforms are compared to those for constant torque operation, then it is evident from Figure 9.7 that there is a significant penalty in terms of RMS current, and hence machine copper losses for constant torque operation. Although the increase in RMS current is not very great, the peak phase current for constant torque operation can be very much higher than for standard modes of operation, as shown in Figure 9.8. This could have considerable consequences for the power converter rating, and could lead to a significant cost penalty. Note that not all constant torque modes have a high peak current penalty: the bipolar minimum RMS current mode has a peak current of 7A and an RMS value of 4.15A for 30Nm - similar to the figures for sinusoidal operation.

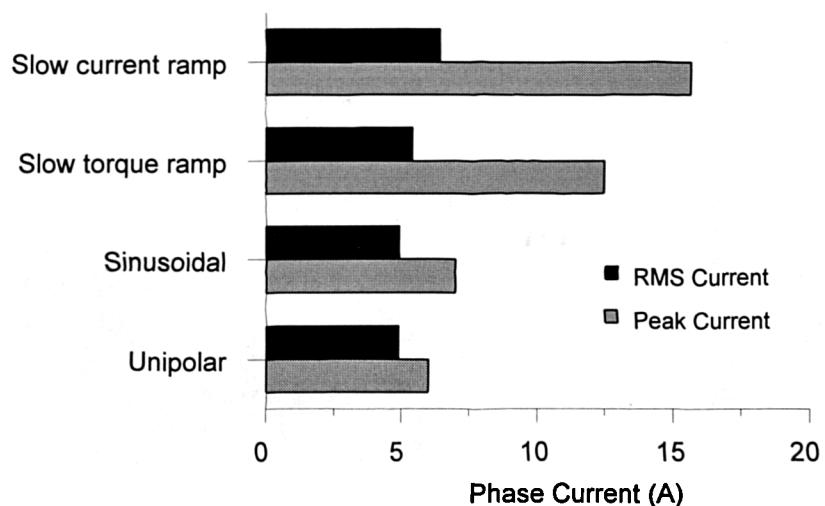


Figure 9.8 Peak and RMS current for standard and constant torque modes giving 30Nm.

9.4 Acoustic Noise

The perceived acoustic noise produced by different modes of operation differs markedly, in both amplitude and signature. In order to quantify the perceived differences in noise, measurements were taken for a variety of waveforms and operating conditions.

9.4.1 Measurement of Acoustic Noise

Figure 9.9 shows the basic experimental arrangement used for the audio noise tests. The audio level meter, a Dawe Instruments model 1400G, was positioned 1m away from the experimental switched reluctance machine, on axis and approximately half way along the stator housing. The test environment was the department machines laboratory, used for all other experimental work. This is a large room, approximately 25m by 15m with a 5m high ceiling. There are no soft furnishings or coverings on the walls, and the laboratory contains many hard objects, though fortunately most are acoustically dead. Thus the laboratory does not make an ideal environment for acoustic measurements due to the many reflective surfaces. The large volume does however prevent too much reverberation from the walls. The laboratory has not been acoustically characterised, and no allowance for its effect on readings has been made. The background noise level was between 46dBA and 47dBA.

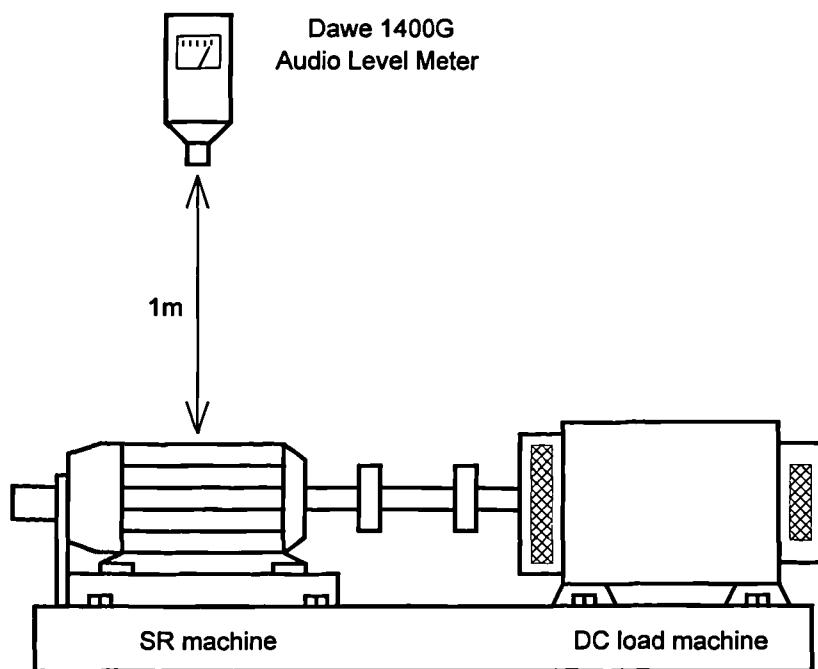


Figure 9.9 Experimental arrangement for acoustic noise measurements.

All noise levels are given in dB and are 'A' weighted (dBA), this gives a reading that corresponds to human perceived 'loudness', rather than absolute sound pressure level.

9.4.2 Variation of Noise with Speed

A 'base-line' for the variation of noise with speed was established by driving the unexcited switched reluctance machine with the DC machine. This is the bottom line in Figure 9.10. Below 100rpm the noise level is only about 1 or 2dB above the background level, rising to 70dBA at 1000rpm. A significant resonance is evident at about 260rpm and is associated with a mechanical resonance in the DC machine.

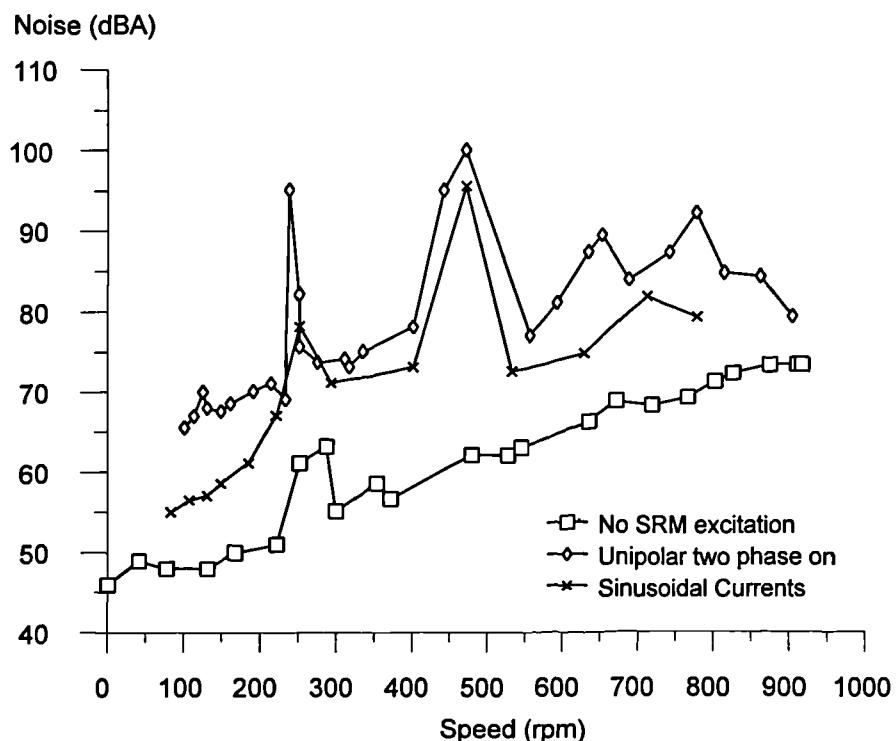


Figure 9.10 Variation in acoustic noise with speed, unloaded, and loaded with unipolar and sinusoidal currents.

Two motoring characteristics are shown in Figure 9.10, unipolar operation with 10A peak phase current (low speed torque $\approx 56\text{Nm}$), and 10A peak sinusoidal currents (low speed torque $\approx 48\text{Nm}$). At around 100rpm sinusoidal operation is about 10dBA quieter than unipolar operation, and remains significantly quieter throughout the speed range. The main reason for this is the low rate of change of flux-linkage and current approaching the aligned position. This reduces excitation of switched reluctance machine stator housing vibrations and hence acoustic noise. Low torque ripple also helps reduce mechanical vibrations in the test rig as a whole.

Several resonant peaks are evident in the noise characteristics, the main ones being around 250rpm and 470rpm, with additional ones grouped around 700rpm. The peak at about 250rpm corresponds to the resonant frequency of the two machine rotor masses at either end of the shaft, and the peak at about 470rpm is due to the switched reluctance

machine stator housing natural frequency. At the lower speed sinusoidal peak noise level is lower than for unipolar operation, due to the lower torque ripple producing less shaft oscillation. Note that the fall in noise level approaching 1000rpm is due to the fall in phase currents and torque.

9.4.3 Variation of Noise with Supply Voltage

Cameron *et al* [9.1] found that the major noise source in switched reluctance drives is stator oscillations. Variation in the radial and tangential stresses, between rotor and stator cause excitation of these oscillations. As the air-gap stresses are proportional to the square of the field strength, and its rate of change is determined by the applied voltage, then it can be expected that noise level will vary with supply voltage.

Figure 9.11 shows measured results for unipolar operation at low speed, and confirms the dependence of noise level on dc-link voltage. Current demand was set at 10A peak, giving approximately 56Nm at 103rpm. Unipolar operation was chosen because it is a 'hard-switched' waveform, i.e. the current transitions are made as quickly as possible. In practice this means that at the switching instants the full dc-link voltage is applied to a phase, the rate of change of flux is thus determined essentially by the dc-link voltage. Low speed operation was chosen to give a low 'base-line' noise level, and so that the actual current waveform, and hence torque, varied very little with supply voltage.

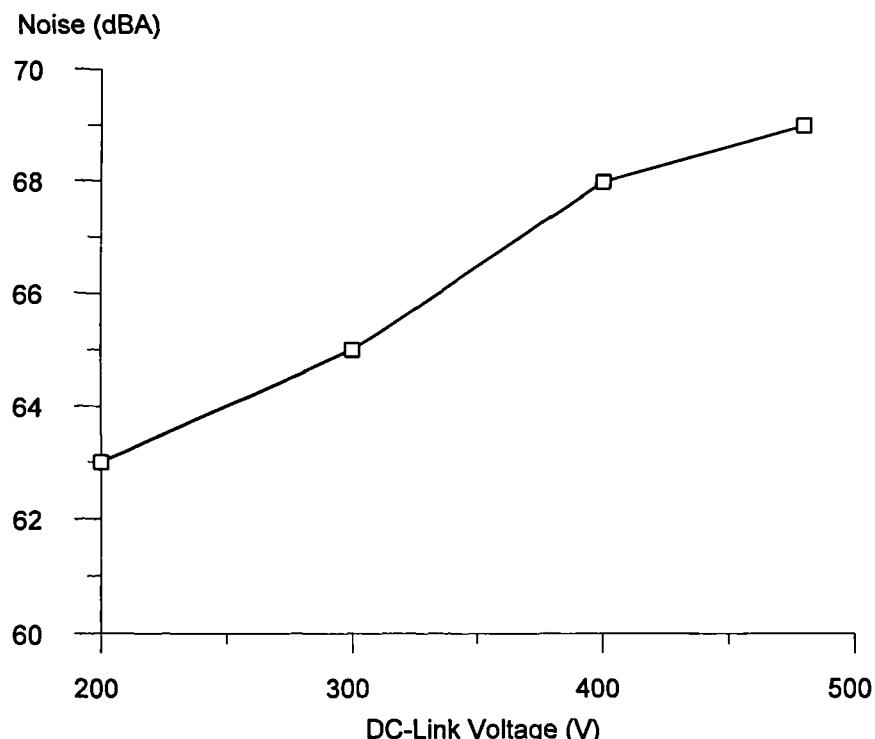


Figure 9.11 Variation in acoustic noise with dc-link voltage for unipolar operation.

9.4.4 Variation of Noise with Current and Torque Waveform

Four 'standard' modes of operation; unipolar, trapezoidal, squarewave and sinusoidal, were characterised for low speed noise. Phase currents were 10A RMS in each case, giving measured average torque values of between 71Nm and 73Nm at 115rpm. Results given in Figure 9.12 show that the noise level varies by 11.5dB between the quietest and the noisiest modes. Sinusoidal operation is much quieter than other modes for reasons previously given. The difference between the other three modes is more difficult to explain. Although the peak MMF is greater, the lower noise figure for squarewave in comparison to unipolar operation can be explained by the MMF dropping significantly before the aligned position, and then remaining constant passing through the aligned position. Radial forces thus change slowly with position reducing stator vibrations. The variation of torque ripple with mode, see Figure 9.6, does not correlate with the noise levels and so could not be considered a major factor in determining low speed noise, as found by Cameron *et al* [9.1].

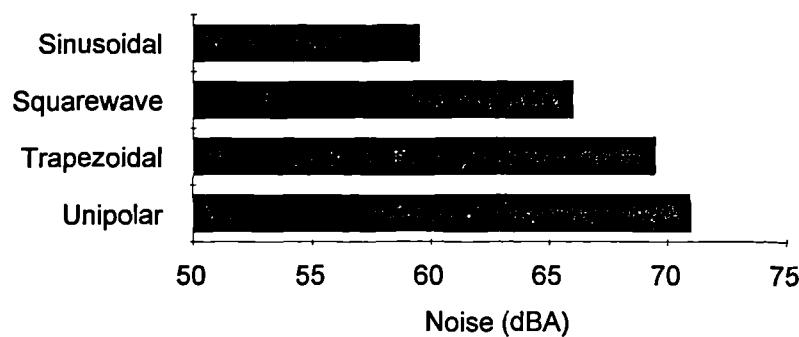


Figure 9.12 Acoustic noise for different 10A RMS current waveforms. All measured at approximately 115rpm, 73Nm, and 420V.

Figure 9.13 shows the equivalent single-tooth winding current near the aligned and unaligned positions, and noise level, for waveforms giving 30Nm average torque. In addition to the standard unipolar and sinusoidal modes, there are six constant torque modes using the constant current ramp, constant torque ramp, and minimum RMS current methods (see Chapter 8 section 8.2). It can be seen that there is a good correspondence (excluding unipolar operation) between peak current in the aligned position and noise level. The higher noise level for unipolar operation can be explained by the faster fall of current at the aligned position and very much higher torque ripple, in comparison to all the other modes. There is no apparent correspondence between peak unaligned current and noise. These results confirm the conclusions of Cameron *et al* [9.1] and Wu *et al* [9.2], that stator ovalising forces are the major noise source. The forces are greatest with the rotor at the aligned position and with a high exciting MMF.

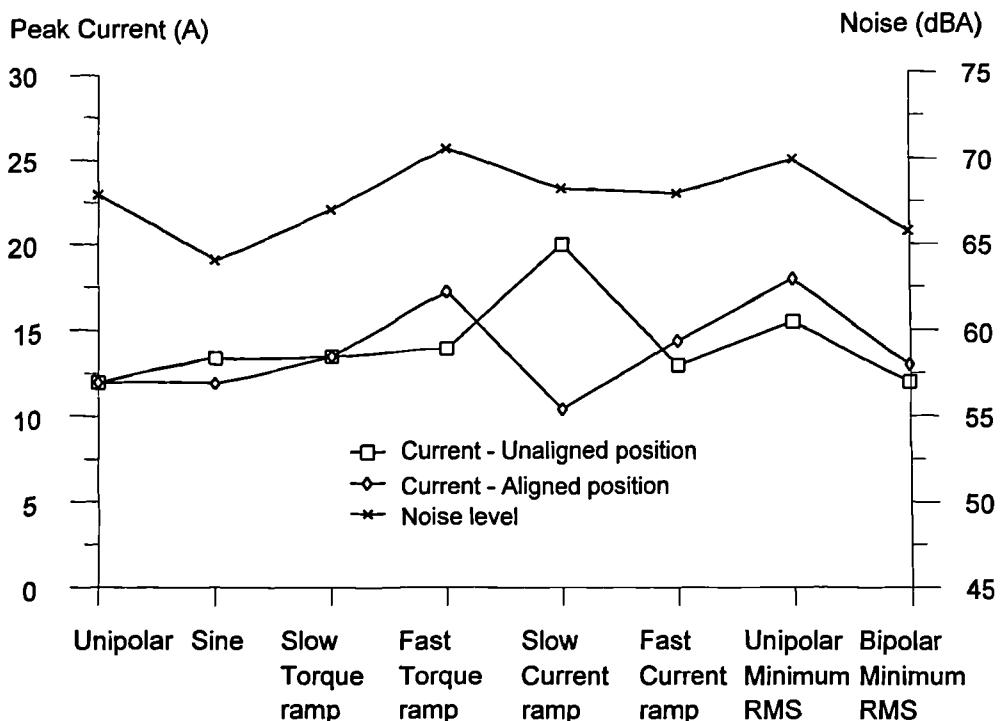


Figure 9.13 Acoustic noise for different current waveforms each producing approximately 30Nm. All measured at approximately 57rpm and 420V.

In addition to the quantitative analysis of noise it is worthwhile making some qualitative statements about the nature of the noise. A feature common to most of the constant torque waveforms is a 'grinding' noise. This can be explained by the high radial forces produced near the aligned position as the tangential component falls and current has to be boosted to maintain torque. Only the slow current ramp waveform does not have this feature and is due to the relatively long and slow reduction of MMF approaching the aligned position. The unipolar waveform gives a 'rattle', with less 'grind' than the constant torque waveforms. High torque ripple and the rapid fall of MMF at the aligned position give rise to oscillations, but the stator ovalising forces are much lower than for some of the constant torque waveforms. Sinusoidal operation is the most pleasant mode on the ear, producing a gentle 'whirring' sound. Low harmonic content, the gentle fall of MMF at the aligned position and low torque ripple combine to give little vibration.

9.5 Summary and Conclusions

At low speed, with equal RMS phase currents and hence equal machine copper loss, there is very little difference between the average torque produced by unipolar, trapezoidal, squarewave and sinusoidal modes. There is however a variation of 35% to 93% in the peak to peak torque ripple produced by these modes, at an average torque of approximately 63Nm. There can be a substantial penalty in increased copper losses and peak converter current with constant torque operation. Some modes investigated result

in peak currents 100% higher, and RMS currents 30% higher than unipolar operation for the same average torque. Not all constant torque modes result in higher losses: bipolar minimum RMS currents give a RMS current 15% lower than sinusoidal operation. Deviation from the optimum advance and conduction angles, for unipolar operation has a detrimental effect on torque for errors greater than about 10° electrical. For constant torque operation a position error of just 3° electrical can double the peak to peak torque ripple.

Acoustic noise is very sensitive to the current and flux-linkage waveform. At low speed sinusoidal operation is 10dBA quieter than unipolar operation for the same shaft torque. Noise has been shown to be dependant on the dc-link voltage and the current profile approaching the aligned position. Excitation patterns having a high and rapidly changing MMF near to the aligned position produce the largest stator vibrations and acoustic noise. Several constant torque waveforms produce significantly more noise because the MMF has to be substantially increased near the aligned position, in order to maintain torque output.

Chapter 10

Conclusions

10.1 Switched Reluctance Machines

Although switched reluctance machines have been around for some time, switched reluctance drives are still a relatively immature technology compared to DC and induction motor drives. With increasing market acceptance resulting in more commercial development and academic research effort, switched reluctance drives are likely to take an increasing large market share in years to come, especially in domestic and automotive applications. Advances in semiconductor technology can be expected to make simple low power drives cheaper and more compact, and improve the performance of servo drives.

The fully-pitched winding arrangement represents one of the most fundamental changes in switched reluctance machines for some time. Simple modelling of phase inductance values shows that for a three phase machine the self inductance is constant, whilst the mutual inductance varies with rotor position. The fully-pitched winding switched reluctance machine can thus be considered the dual of a single-tooth winding switched reluctance machine, as torque is derived from changes in mutual, rather than self inductance. The ability to use periods of both rising and falling mutual inductance to produce positive torque gives much greater flexibility in the choice of excitation. Although the fully-pitched winding arrangement results in longer end-windings, and hence phase resistance, better winding utilisation enables the experimental fully-pitched winding switched reluctance machine to produce more torque than an equivalent single-tooth winding switched reluctance machine for equal copper loss. In terms of converter rating detailed analysis is required to differentiate single-tooth winding and fully-pitched winding drives with unipolar currents, but bipolar currents are likely to require a higher kVA/kW as twice as many devices are required.

Transformation equations allowing electrical parameters in the fully-pitched winding switched reluctance machine to be related to those in an equivalent single-tooth winding switched reluctance machine have been described. By de-coupling the windings of a fully-pitched winding machine, these transformations allow the machine to be modelled

as a conventional single-tooth winding machine, and are a very useful tool for simulation and analysis.

10.2 Operation with Current and Voltage Control

The experimental switched reluctance drive hardware and loading arrangement was designed for maximum flexibility and ease of configuration. Using a power circuit which by control of device signals can emulate several common configurations, and by implementing control functions in software, it has been possible to explore several circuit configurations and control strategies without hardware modifications.

Current control methods used for switched reluctance drives have been reviewed. For a simple analogue implementation, hysteresis or constant off time control gives a simple unconditionally stable implementation. A PI or PID controller and PWM output has the advantage of lower current ripple and is better suited to digital implementation. The current controller used for the experimental drive has several features to improve transient performance under large signal conditions. A simplified plant model was used to make initial estimates for controller gain terms, and the effectiveness of anti-windup features has been demonstrated.

Experimental operation of the fully-pitched winding switched reluctance drive has been explored with unipolar and bipolar phase currents. With unipolar phase currents positive torque can only be produced from periods of rising mutual inductance, and one electrical cycle corresponds to a rotation of one rotor tooth pitch. Bipolar phase currents allow much greater choice in phase current waveshape and several possibilities have been explored. Trapezoidal conduction sequences produce positive torque from periods of falling mutual inductance only, whilst squarewave excitation produces positive torque from periods of rising and falling mutual inductance. For bipolar excitation one electrical cycle corresponds to a rotation of two rotor tooth pitches.

Simulation has been used to predict electrical waveforms and torque for unipolar and bipolar operation, at low speed with an ideal current source, and at high speed with an ideal voltage source. At low speed simulated and measured waveforms correspond very closely, and at high speed the small differences are due to current limiting by the real controller - an effect which is not simulated. Extensive use of transformations has been made to calculate equivalent single-tooth winding parameters from fully-pitched winding parameters, in simulation and to aid the understanding of operating waveforms.

Comparisons made on the basis of slot MMF and stator tooth flux between unipolar operation of a fully-pitched winding SR drive and a conventional single-tooth winding

SR drive show that at low speed, when operating under current control, the tooth fluxes and slot MMFs are very similar, and so it could be said that the underlying mode of operation is the same. However at high speed when operating under voltage control, the tooth fluxes and slot MMFs are quite different.

The torque/speed envelope for unipolar operation of the experimental fully-pitched winding switched reluctance drive is much like that of many other electrical machine drives, with shaft power rising up to a base speed and then remaining constant thereafter. Comparisons made with the experimental single-tooth winding switched reluctance drive show that the torque/speed envelopes are quite similar with the control strategy employed. At low speed the torque per RMS ampere of phase current is almost the same for trapezoidal, squarewave, sinusoidal and unipolar phase currents. There are significant differences in torque ripple, with sinusoidal and trapezoidal excitation having approximately half the torque ripple of squarewave operation. At high speed all three bipolar modes produce substantially less torque than unipolar operation. However, it has been shown that this is due to the current reference and controller. With suitable reference waveforms bipolar operation can equal unipolar power output, but this requires a more complex controller as the fundamental reference waveshape must change as speed increases.

In terms of power converter rating figures suggest that a unipolar fully-pitched winding SR drive would require a converter rating between 9% and 16% higher than a conventional SR drive for the same shaft power. Although in theory bipolar squarewave operation gives substantially more torque for the same peak phase current, in practice this advantage does not materialise. In terms of power converter kVA/kW rating bipolar operation is at a significant disadvantage, because twice as many power switches are required unless the machine is star connected. Delta connection gives very unsatisfactory results due to a large third harmonic circulating current.

10.3 Advanced Drive Control

It has been shown that the discrete time control function required for 'dead-beat' flux response is very simple, and a flux-linkage controller has been implemented for switched reluctance machine control. In comparisons with digital hysteresis and PID current controllers the flux-controller phase current response was superior in every case.

Low speed drive performance using the flux-linkage controller is very good, but at high speed performance can be poor with inappropriate reference waveforms. By taking advantage of the simple relationship between flux-linkage and voltage, and designing the reference waveform for flux control rather than current control, in experimental tests

30% more torque was produced at a 30% higher speed, i.e. 69% more shaft power. Using flux-linkage controller and adding an offset to the flux reference can improve high speed torque capability. In simulated results for the fully-pitched winding switched reluctance drive the average torque at 2000rpm was doubled, and the torque per RMS ampere increased 23% by offsetting the flux reference by 0.2Wb from zero.

Although very good performance can be achieved using a flux-linkage controller it does require machine flux-linkage characteristics to be accurately known, and high resolution rotor position sensing must be used as current and torque waveforms are more sensitive to rotor position errors than with current control.

Many authors have proposed methods to overcome the problem of torque non-linearity with respect to phase current and rotor position in switched reluctance machines. Although the methods reviewed work well at low speed, little account has been taken of high speed performance. Simulation work has shown that in order to produce constant torque over a wide speed range account must be taken of what excitation will be required significantly further on in a cycle.

When using a flux-linkage controller it is relatively easy to predict performance later on in a cycle as flux-linkage is simply equal to the applied volt-time integral, neglecting resistance. By using a flux-linkage waveform composed of a series of linear ramps which follow machine flux-linkage/position characteristics and by limiting the flux ramp rate, reference waveforms have been calculated to give constant torque over a wide speed range. A Genetic Algorithm was employed to find the optimum flux ramp reference waveforms as it can search a large search space effectively. Experimental results have shown the Genetic Algorithm optimised flux-linkage waveforms give the desired torque characteristics over a 20 to 1 speed range, and 25 to 1 torque range. A speed controller using an array of optimised flux profiles has been implemented. Dynamic gain changing in response to large speed errors improves the transient response considerably.

There can be a substantial penalty in increased copper losses and peak converter current with constant torque operation. Some modes investigated result in peak currents 100% higher, and RMS currents 30% higher than unipolar operation, for the same average torque. However not all constant torque modes result in higher losses: bipolar minimum RMS currents give a RMS current 15% lower than sinusoidal operation. Deviation from the optimum advance and conduction angles, for unipolar operation has a detrimental effect on torque for errors greater than about 10° electrical, but for constant torque operation a position error of just 3° electrical can double the peak to peak torque ripple.

Acoustic noise is very sensitive to the current and flux-linkage waveform. At low speed sinusoidal operation is 10dBA quieter than unipolar operation for the same shaft torque. Noise has been shown to be dependant on the dc-link voltage and the current profile approaching the aligned position. Excitation patterns having a high and rapidly changing MMF near to the aligned position, e.g. some constant torque waveforms, produce the largest stator vibrations and acoustic noise.

10.4 Further Work

Although the work outlined in this thesis has answered many questions and presented several control strategies, there are still some questions which remain unanswered and avenues of research which deserve further investigation.

On the question of converter rating of fully-pitched winding switched reluctance drives in comparison to single-tooth winding switched reluctance drives, much more detailed simulation work is required to explore the limits of drive performance within the machine and converter thermal limits. Only then will it be possible to give a definitive answer.

In chapter 6 section 6.3.3 a novel bipolar current waveform was presented. This had the interesting feature that the equivalent single-tooth winding current could have a conduction angle greater than 120°, without doubling the peak fully-pitched winding current. This type of waveform has possibilities for constant torque operation as it would be possible to profile the single-tooth winding currents without a large increase in fully-pitched winding current and hence converter rating.

The use of a Genetic Algorithm for optimising flux parameters for constant torque has been shown to work well but there is scope for improvement. Better results would be obtained if RMS phase current, peak phase current and average torque error were used in addition to torque ripple when calculating the fitness of an organism. It may be possible to use another Genetic Algorithm to optimise the weights of these additional parameters but extensive testing would undoubtedly be required to determine which weights gave the best results over a wide range of operating conditions. A further enhancement would be to add an additional point to the flux-linkage reference near to the aligned position. This would give much greater scope in controlling peak current and torque in this region and yield better results.

Appendix A

Switched Reluctance Machine Characteristics

A.1 The Experimental Machines

There are two switched reluctance machines used for all the experimental and simulated results in this thesis; a fully-pitched winding switched reluctance machine, and its equivalent single-tooth winding switched reluctance machine. Both machines started life as commercial switched reluctance machines produced by Allen-West in the UK, to a design licensed from Switched Reluctance Drives Ltd. The fully-pitched winding switched reluctance machine was produced by taking a conventional, three phase, 12:8, switched reluctance machine with single-tooth windings, an Allen-West model SR75, and replacing the windings with a set of fully-pitched coils. Wire of the same diameter was used, and the total number of turns per phase was also kept the same. No modifications to the original end castings, stator stack, or housing, were required despite the longer end-windings. Another Allen-West model SR75 is used for the experimental single-tooth winding switched reluctance machine, this is a standard unmodified machine. The experimental fully-pitched winding, and equivalent single-tooth winding machines, are thus identical, apart from the winding configuration and small differences due to manufacturing tolerances and lamination magnetic properties.

A.2 Physical Dimensions and Windings

Major mechanical and electrical parameters for the experimental switched reluctance machines are shown in Table A.1. The two machines are essentially identical, apart from the winding configuration, which affects the length per turn, and phase resistance.

Parameter	Value for fully-pitched winding machine	Value for single-tooth winding machine	Units
Frame Size	D132	D132	-
Number of rotor teeth	8	8	-
Number of stator teeth	12	12	-
Stack length	194	194	mm
Rotor outside diameter	120.0	120.0	mm
Stator inside diameter	120.8	120.8	mm
Nominal air gap	0.4	0.4	mm
Rotor tooth face width	16	16	mm
Stator tooth face width	16	16	mm
Number of electrical phases	3	3	-
Number of turns per phase	244	244	-
Phase resistance at 20°C	1.147	0.797	Ω
Copper mass	12.0	8.34	kg

Table A.1 Physical parameters of the experimental fully-pitched winding, and single-tooth winding, switched reluctance machines.

A.3 Flux-Linkage Characteristics

A.3.1 Measurement of Flux-Linkage Characteristics

The flux-linkage characteristics of a machine determine its torque capability, and are required for performance prediction and simulation. To characterise a machine it is necessary to determine the relationship between flux-linkage and current, at a number of fixed rotor positions. The rotors of the experimental machines were held stationary by

fixing a clamping plate to the machine bed-plate, and bolting the shaft coupling to the clamping plate. To reduce the amount of additional equipment required, and allow some degree of automation, the drive controller and power converter were used for the flux-linkage testing. Three parameters need to be measured for the determination of flux-linkage characteristics; phase current, flux-linkage and rotor position. Phase current was measured using a Hall Effect current transducer. To determine flux-linkage, rather than measuring the phase voltage, the dc-link voltage was measured, and the voltage drop across the power switches taken into account, along with the phase resistance, when calculating the flux-linkage. Rotor position was measured using the optical position encoder, giving a resolution of approximately 0.3° mechanical.

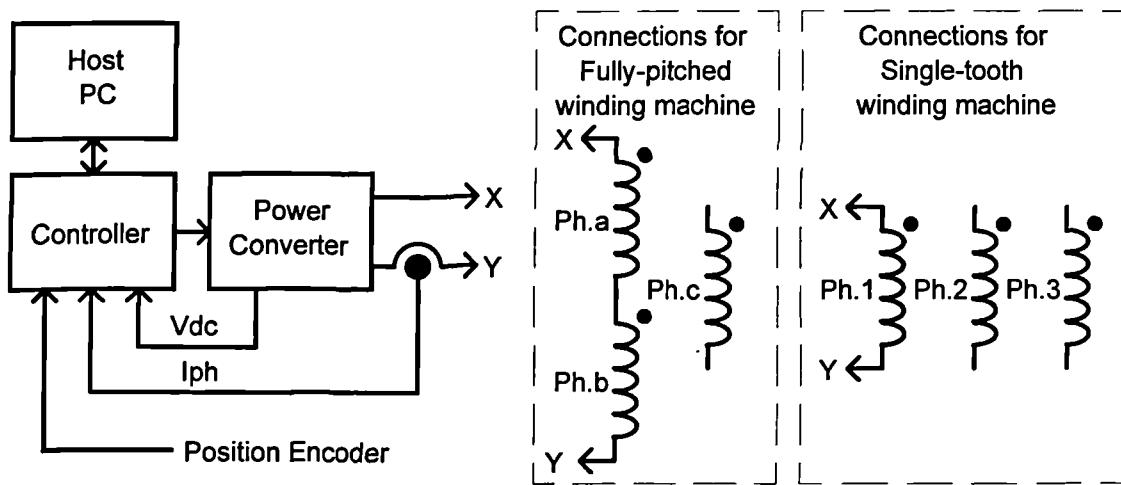


Figure A.1 Flux-linkage measurement circuit.

The system configuration used for flux-linkage tests is shown in Figure A.1. For the fully-pitched winding machine, tests are done with two phases connected in series, with the third open circuit; and for the single-tooth winding machine, only one phase is connected to the power converter. The drive controller displays the rotor angle on a monitor whilst positioning the rotor, before being locked in the desired position. The flux-linkage is characterised automatically by the drive controller which applies maximum voltage, starting from zero phase current, until the phase current reaches 20A. At this point the devices are switched off, allowing the current to decay by flowing in the free-wheeling diodes back to the supply. Phase current and dc-link voltage are sampled every 50μs during the rise in current, and values stored in an array. The drive controller then calculates the phase flux-linkage using the following equation:-

$$\Psi_n = \Psi_{n-1} + (V_{dc} - V_{sw} - i_n R_{sw} - i_n R_{ph}).T \quad (A.1)$$

Where $\Psi_n, V_{dc}, V_{sw}, R_{sw}, R_{ph}, i_n$ and T , are the flux-linkage, dc-link voltage, power switch offset voltage, power switch incremental resistance, phase resistance, phase

current, and sample interval respectively. Having done this the phase current and flux-linkage data is passed to the host PC, via a serial link for storage on disk. The raw data is then processed to obtain a data array which has fixed current intervals, over a defined range. This whole sequence is repeated for each rotor position. For the fully-pitched winding switched reluctance machine it was decided to measure the flux-linkage characteristics at 21 positions per tooth pitch (approximately every 1° of rotation).

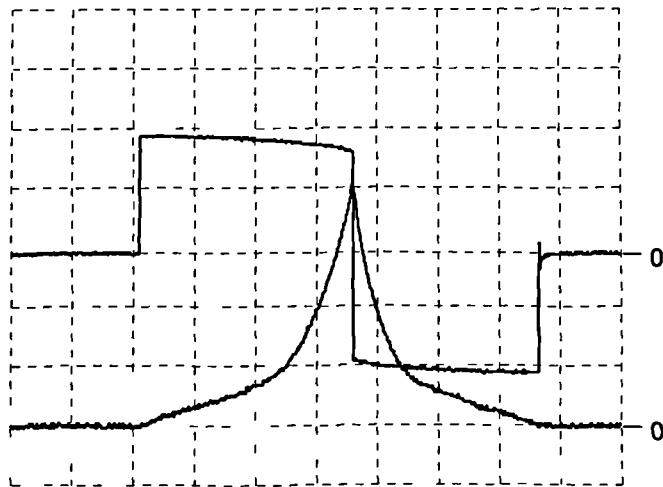


Figure A.2 Flux-linkage testing at the aligned position. Trace 1 - phase voltage, 100V/div. Trace 2 - phase current, 5A/div. Horizontal - 5ms/div.

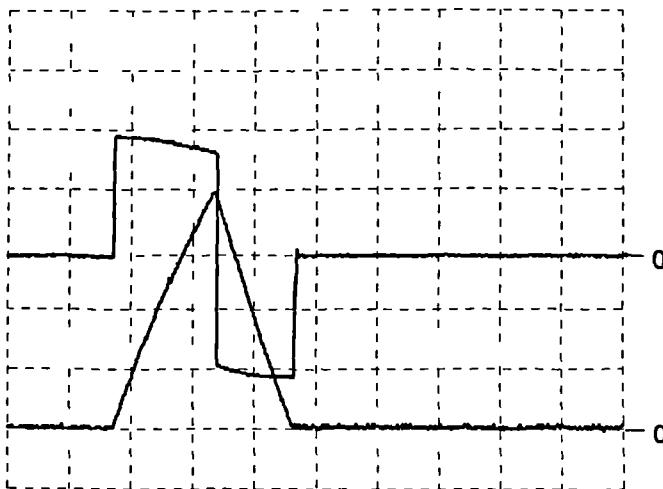


Figure A.3 Flux-linkage testing at the unaligned position. Trace 1 - phase voltage, 100V/div. Trace 2 - phase current, 5A/div. Horizontal - 5ms/div.

Phase voltage and current waveforms during flux-linkage tests are shown in Figure A.2 and Figure A.3, which are for the aligned and unaligned rotor positions respectively. Note that the phase voltage drops slightly as the current increases and the dc-link capacitors discharge.

A.3.2 Fully-Pitched Winding Switched Reluctance Machine Characteristics

The fully-pitched winding switched reluctance machine flux-linkage characteristics, determined using the method described, are illustrated graphically in the following figures.

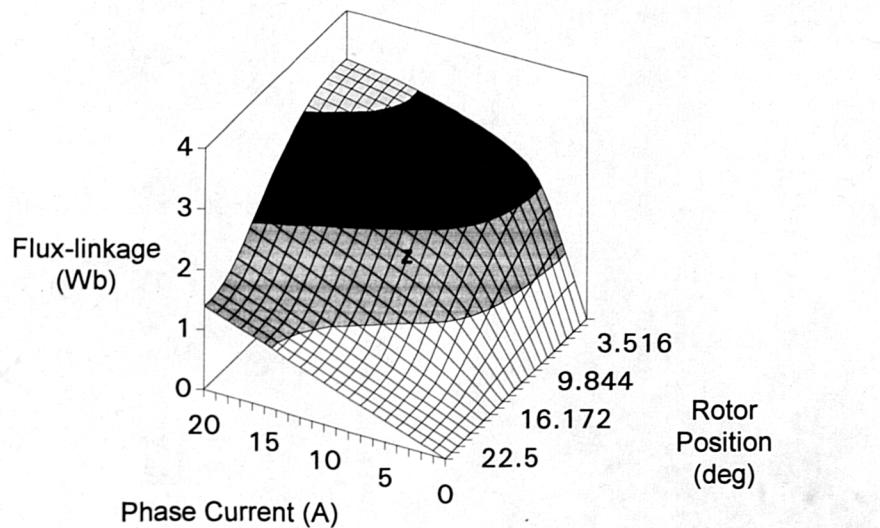


Figure A.4 Three dimensional plot of fully-pitched winding switched reluctance machine flux-linkage characteristics.

Detailed analysis of the measured flux-linkage data revealed that sets of data for some rotor positions did not fit the expected characteristics exactly, although the errors were small, being in the order of 1%. Further study of the possible error sources lead to the conclusion that the errors were due to rotor position inaccuracy. It was originally thought that a position resolution of 0.3° would be sufficient, but the quantised nature of the encoder output, means that the position quantisation error can be a significant factor, in the difference between two actual rotor positions. To correct these errors, a position correction factor was applied to certain sets of data, to bring them into line with the estimated best fit.

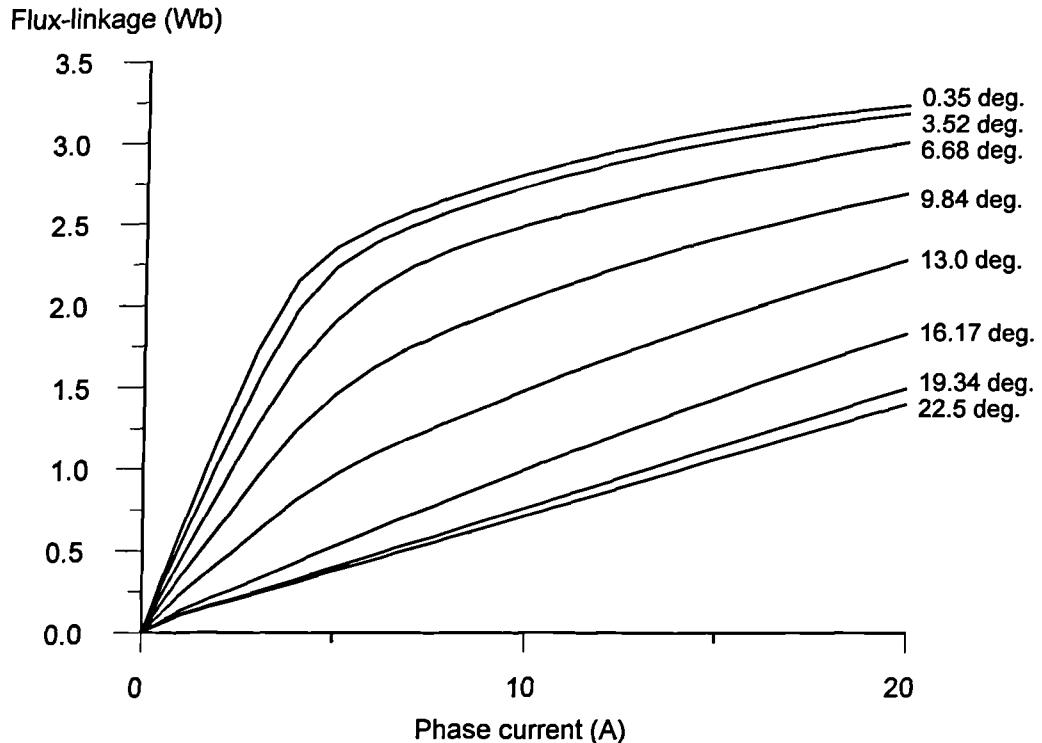


Figure A.5 Fully-pitched winding switched reluctance machine flux-linkage characteristics for constant rotor position.

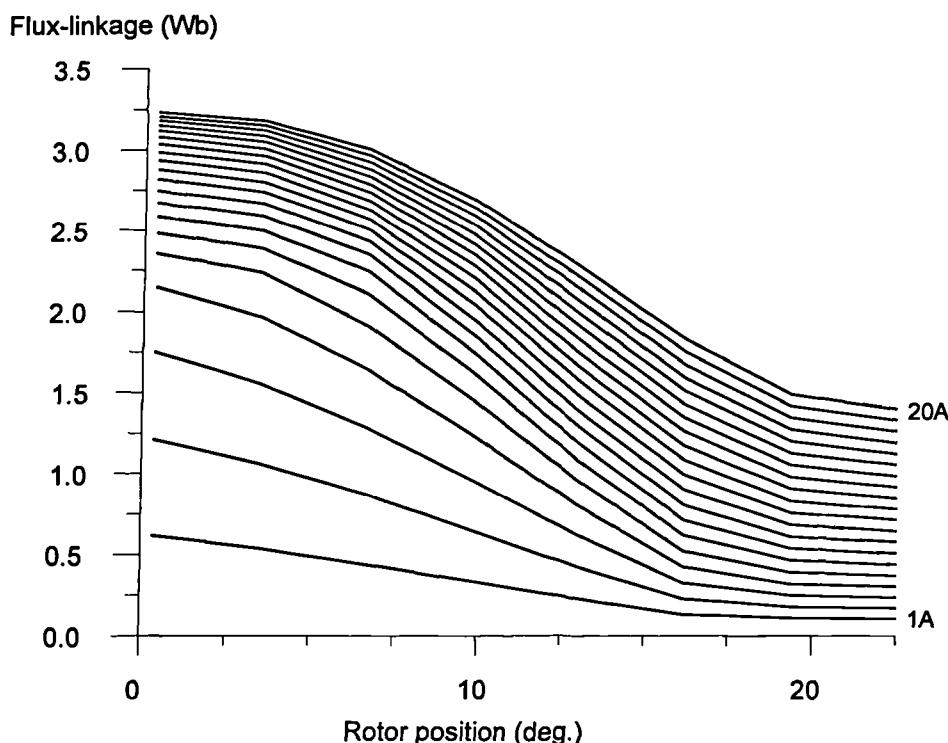


Figure A.6 Fully-pitched winding switched reluctance machine flux-linkage characteristics for constant phase current.

A.3.3 Single-Tooth Winding Switched Reluctance Machine Characteristics

Before being re-wound to make the fully-pitched winding switched reluctance machine, the original single-tooth winding switched reluctance machine was characterised by B.C.Mecrow, using a DC voltage source and data acquisition system. This data and the characteristics of the second single-tooth winding switched reluctance machine, measured by P.G.Barrass using the method described previously, are shown in Figure A.7. The second single-tooth winding switched reluctance machine was only characterised at the aligned, and unaligned positions (shown as heavy lines) for comparison with the first single-tooth winding switched reluctance machine. It is evident that the flux-linkage in both positions is slightly greater for the second machine than the first. This could be due to experimental differences between the methods used by B.C.Mecrow and P.G.Barrass, and differences in the two machines due to manufacturing tolerances and possible differences in lamination characteristics. No tests have been done to determine the exact source of the discrepancies.

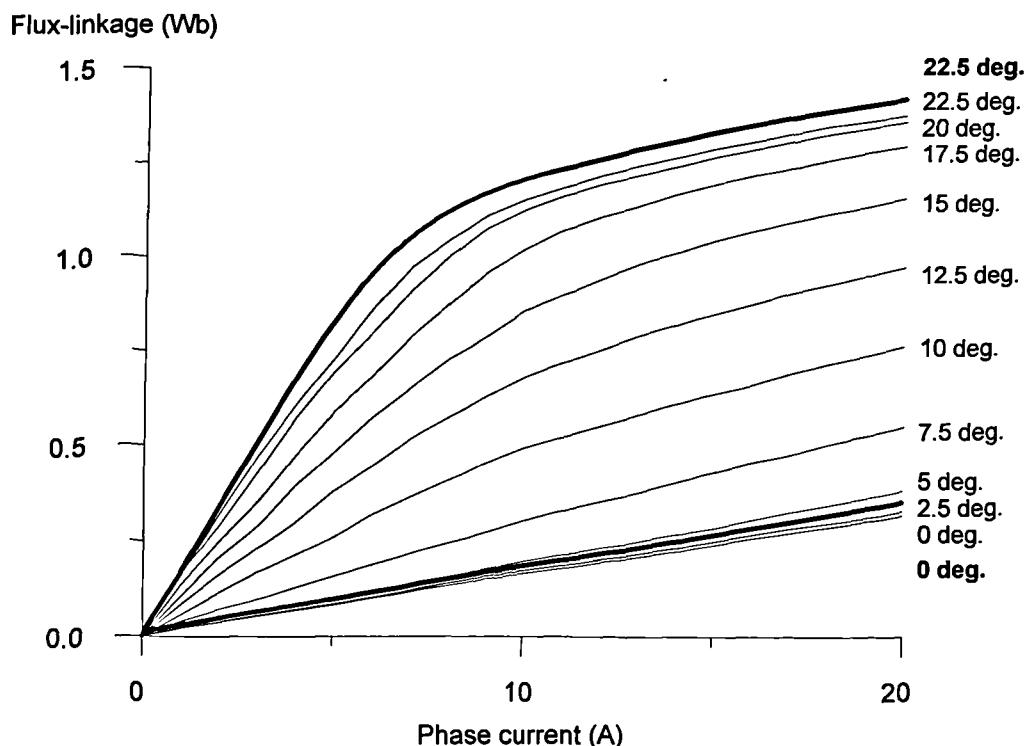


Figure A.7 Single-tooth winding switched reluctance machine flux-linkage characteristics for constant rotor position. Light lines - data measured by B.C.Mecrow. Heavy lines - data measured by P.G.Barrass.

A.4 Static Torque Characteristics

Torque characteristics for the experimental fully-pitched winding switched reluctance machine are shown in the following figures.

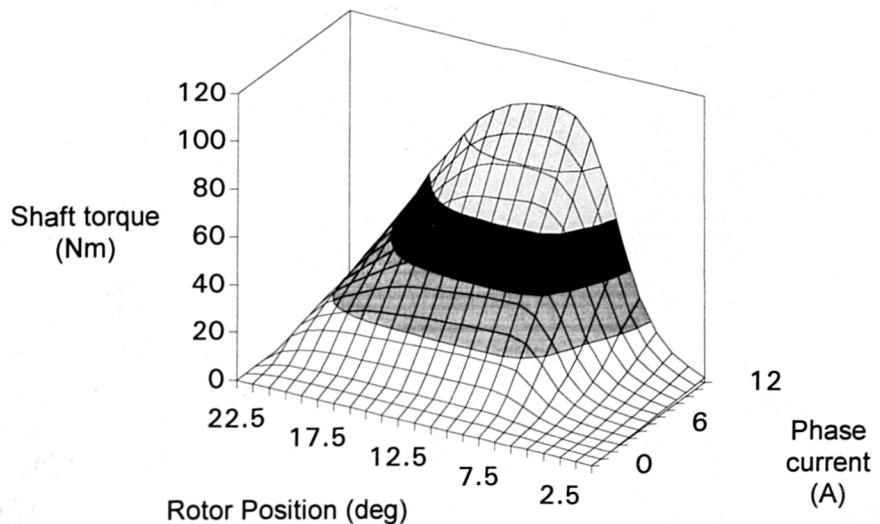


Figure A.8 Three dimensional plot of fully-pitched winding switched reluctance machine torque characteristics for two phase on unipolar excitation.

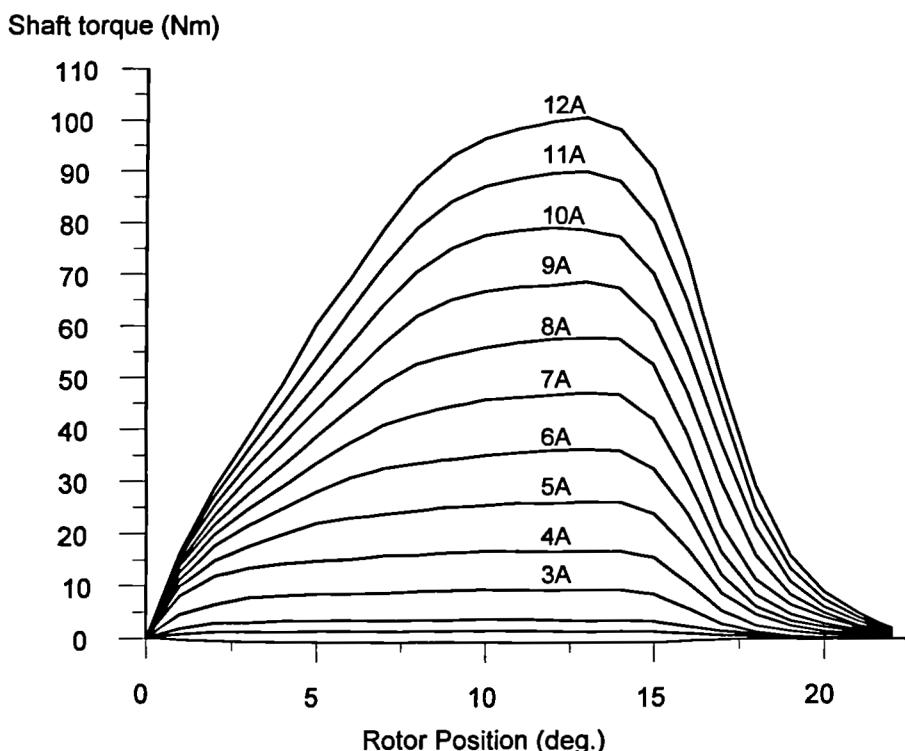


Figure A.9 Fully-pitched winding switched reluctance machine torque characteristics for constant phase current. Two phase on unipolar excitation.

Results are for two phases connected in series with the third phase open circuit. The static torque characteristics were measured by B.C.Mecrow using a Sensor Technology torque transducer and display system. Data used is in each figure is for currents of 0A to 12A per phase in steps of 1A, and for rotor positions of 0° to 22.5° in steps of 2.5°.

A.5 Calculation of Flux-Linkage from Torque

For a loss-less energy conversion system the change in electrical and mechanical energies must be equal. Applying this principle to a switched reluctance machine, the magnetic circuit co-energy can be related to shaft torque. Considering a machine with phase current i , moving from a rotor position θ to $\theta+d\theta$, then equating the electrical input and mechanical output:-

$$W'(\theta + d\theta, i) - W'(\theta, i) = T(i) \cdot d\theta \quad (\text{A.2})$$

This is shown graphically in Figure A.10, where the shaded area in each case represents the change in energy.

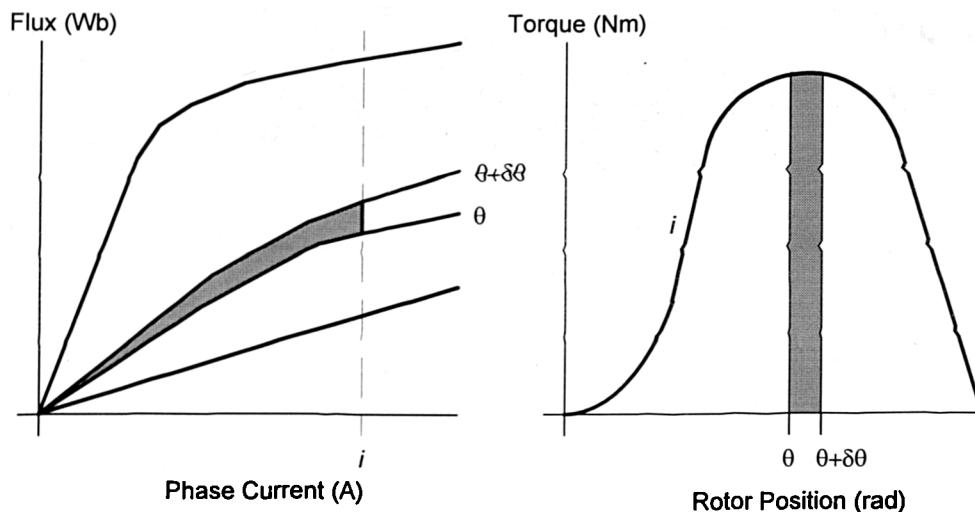


Figure A.10 Comparison of electrical and mechanical energy.

Thus by relating energies, is possible to calculate flux-linkage from torque, just as torque is often calculated from flux-linkage. There is however a problem with the calculation of flux-linkage because only the *change* in flux-linkage can be calculated from the torque. No absolute value can be calculated as the flux integral initial conditions can not be derived from the torque characteristics. Thus to obtain the flux-linkage characteristics is necessary to measure the machine phase flux-linkage at a known position. This is most easily done by measuring the phase inductance at the unaligned position, where the absence of saturation allows a linear relationship between current and flux-linkage to be used.

Appendix B

Machine Loading Rig Characteristics

B.1 Physical Arrangement

The mechanical arrangement of the experimental machine test rig is shown in Figure B.1. Rated speeds of the experimental SR machine and the DC load machine are similar enough to allow direct coupling. Rigid couplings are used so that when a torque transducer is fitted the torque ripple can be measured without the damping effects of a flexible coupling. The use of rigid couplings does however make very accurate machine alignment necessary, to avoid large variations in torque due to shaft stressing. Another disadvantage of rigid couplings is that the mechanical system, comprising of the two machine rotor inertias and the coupling shaft, has a high Q factor. When excited by the switched reluctance machine torque ripple, at or near resonance, there are considerable shaft oscillations, which make torque transducer reading very inaccurate, there is also a considerable increase in acoustic noise.

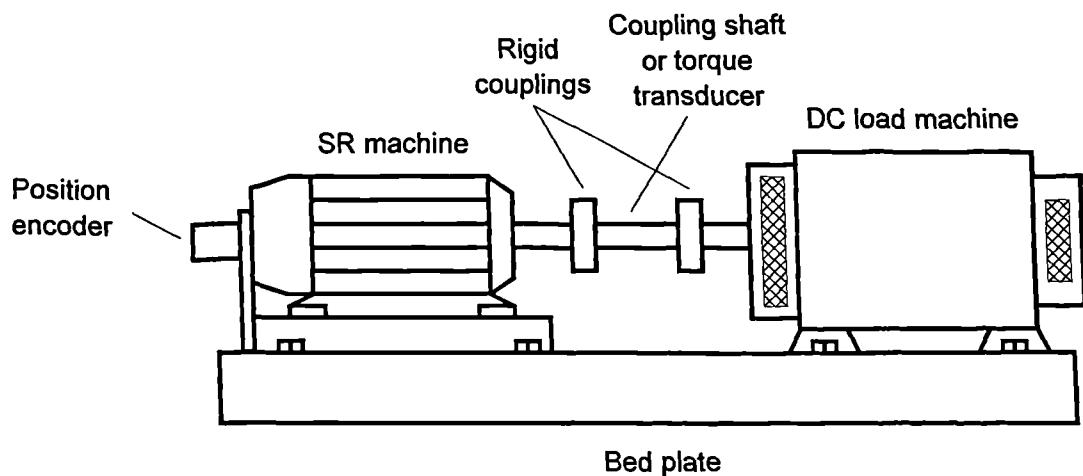


Figure B.1 Mounting arrangement of experimental SR, and DC load machines.

The two machines are mounted on a common steel bed-plate, to which can also be attached a fixing plate, used for locked rotor tests. The optical position encoder is held in

place by a bracket fixed to the bed-plate and is coupled to the switched reluctance machine rotor using a shaft extension mounted at the fan end of the machine, and a flexible coupling.

B.2 DC Machine Characteristics

The load machine is a four pole brushed DC motor, with inter-poles and separate field connections, its name-plate ratings are shown in Table B.1. The field is designed for shunt operation, and when cold (approximately 20°C) the rated field voltage gives a field current of about 1.6A. This drops to about 1.3A at normal operating temperature (outer casing temperature approximately 40°C).

Parameter	Value	Units
Rated power	10	kW
Rated speed	1080	rpm
Armature voltage	240	V
Armature current	50	A
Field voltage	240	V

Table B.1 Name-plate rating of DC the load machine.

The field magnetic circuit is well into saturation at nominal field voltage and current, and so the torque constant/field current characteristic is very non-linear. To get an accurate equation for the machine torque constant two operating points were characterised. The first point is at nominal field voltage, which gives a maximum speed of approximately 1100rpm (for nominal armature volts), with the second point at a reduced field voltage, which gives a top speed of around 1600rpm. The torque constant was measured by driving the DC load machine with the experimental switched reluctance machine, and measuring the open circuit armature voltage at several speeds, for field currents around 1.3A and 0.45A. The results of these tests are shown in Figure B.2 and Figure B.3. In each of these figures the points are measured values, and the line represents the best fit.

Machine Loading Rig Characteristics

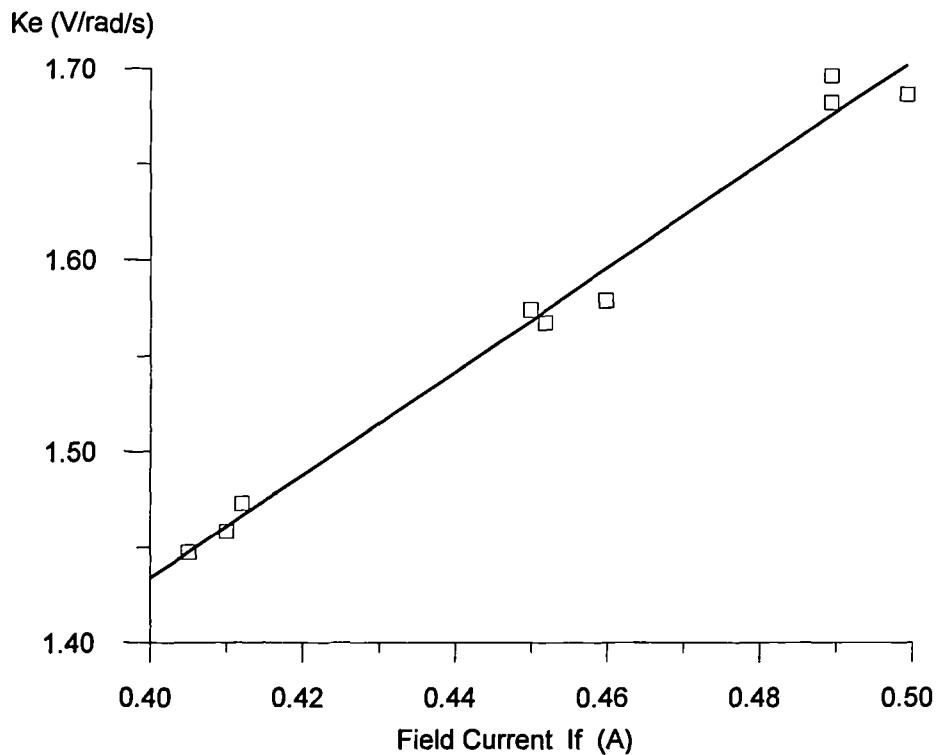


Figure B.2 DC load machine torque constant at low field currents.

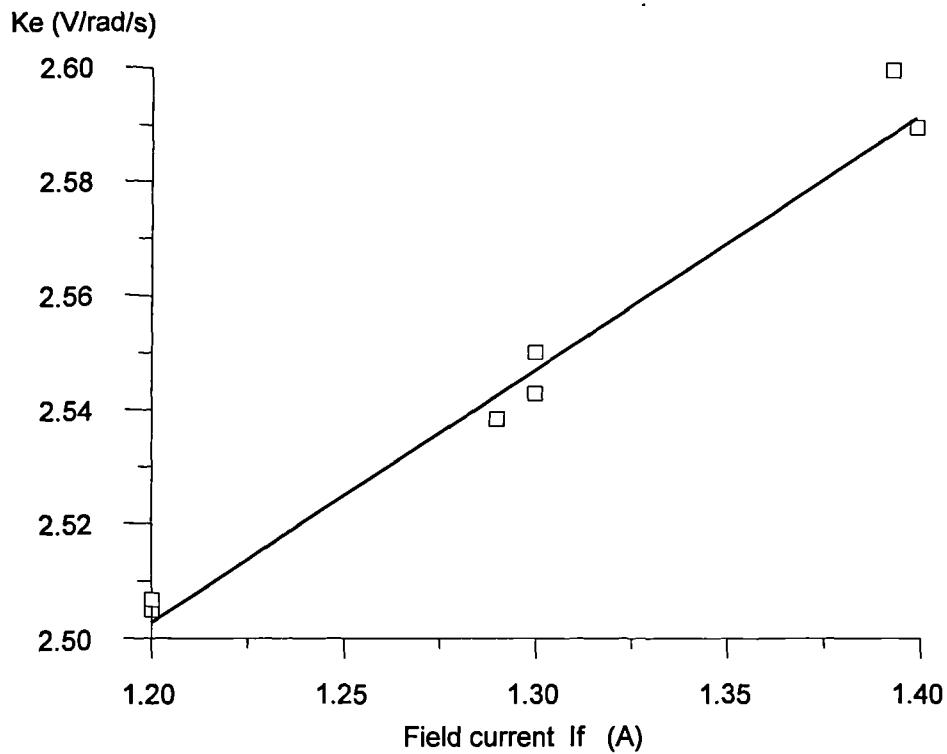


Figure B.3 DC load machine torque constant at high field currents.

The straight line equations for the DC load machine torque constant, at high and low field currents are:-

$$K_e = 1.97 + 0.440I_f \quad (1.2 \leq I_f \leq 1.4) \quad (\text{B.1})$$

and

$$K_e = 0.354 + 2.70I_f \quad (0.4 \leq I_f \leq 0.5) \quad (\text{B.2})$$

Armature circuit resistance was measured with the field current set to zero, and armature held stationary by the switched reluctance machine, so that the back emf would be zero. Current was then passed through the armature circuit, with the current, voltage and power measured. The resistance was then calculated using two methods and the results averaged to give a figure of $R_{arm} = 0.398\Omega$.

To determine the effects of armature reaction in the DC machine, the machine was run at various speeds and loads, and the measured armature voltage compared with that determined from the no-load machine parameters. To within the measurement error bounds, no change in the effective torque constant was found for armature currents up to approximately 30A, so it can be concluded, that armature reaction effects can be neglected at currents up to this level.

Rotating losses in the test rig come from two main sources: the switched reluctance machine, and the DC load machine. Each contributes losses from bearing and seal friction, windage losses, and eddy current losses. When unexcited the iron losses in the switched reluctance machine are assumed to be negligible, as the remnant field is small. Eddy current losses in the DC load machine, when excited, are quite substantial and necessitate the characterisation of rotating losses at different values of field current. Rotating losses were measured by using the DC machine as a motor, and measuring the electrical input power. Armature resistance losses were then subtracted, to yield the rotating losses. Measured losses, and the best polynomial fit are shown in Figure B.4, for field currents of 0.45A and 1.3A.

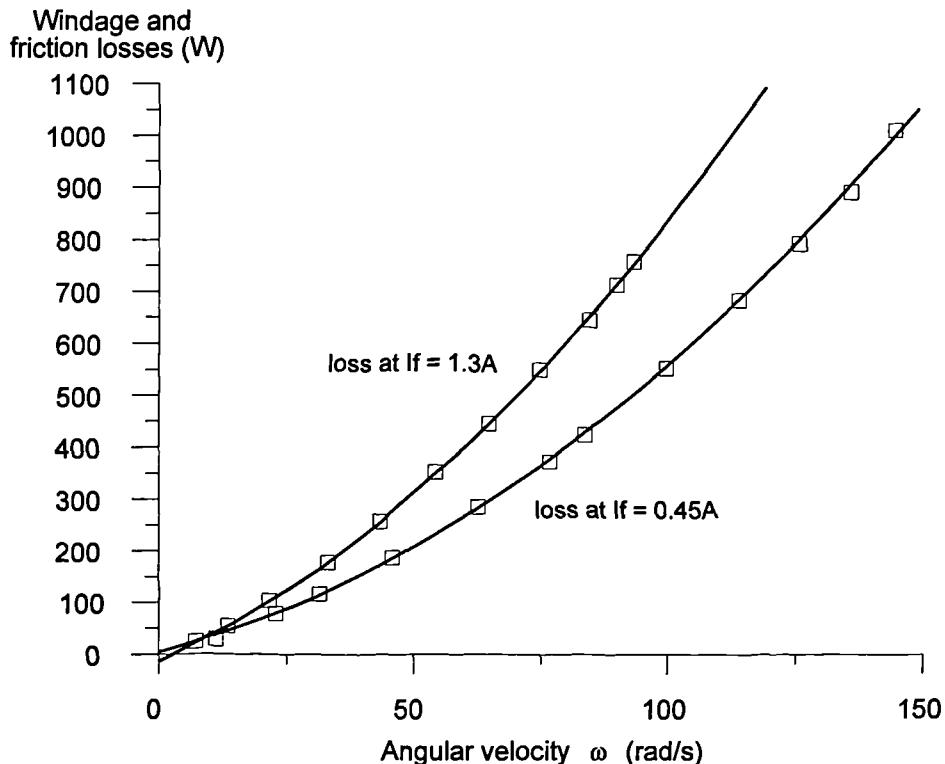


Figure B.4 Machine test rig rotating losses.

The equation for windage and friction loss power at If=1.3A is:-

$$P_{loss} = -11.696 + 4.488\omega + 0.040\omega^2 \quad (B.3)$$

and at If=0.45A is:-

$$P_{loss} = 2.362 + 2.633\omega + 0.029\omega^2 \quad (B.4)$$

B.3 Calculation of Shaft Torque

Shaft torque can be found by one of three methods:-

- Direct measurement using a torque transducer
- Indirectly from measurements of electrical quantities
- Indirectly from measured phase currents, and static machine characteristics

The first method is to some extent the most desirable, as the torque is measured directly using a single instrument. In practice however, several factors have made the use of a torque transducer inappropriate. The torque transducer used is made by Sensor

Technology Ltd., and uses an optical grating to determine shaft torque. Torque variation caused by shaft misalignment in conjunction with the rigid shaft couplings, and large torque pulsations at shaft resonance speeds make the torque transducer readings unreliable. Furthermore, the torque transducer was also found to be generally inconsistent and unreliable, with readings over time and temperature. For these reasons an alternative method of shaft torque measurement was sought.

The third method; using stored static torque characteristics and measured phase currents, requires sufficient processing power to calculate torque in real time. This was not available at the start of the project. This method does not take into account differences in static and dynamic torque, current measurement errors, or errors in the static torque measurements. Thus, although this method could be considered suitable for qualitative torque waveform analysis, and approximate average torque calculation at low speed, it is not suitable for accurate measurement of shaft torque.

This leaves the second method; calculation of torque from electrical quantities. Two methods of torque calculation based on electrical parameters are used: load machine air-gap power, and load machine torque characteristics. For each of these methods the system losses are added to the air-gap power, in order to calculate the switched reluctance machine electromagnetic torque T_{elec} . The two methods are described in more detail below, with reference to Figure B.5, which shows the measured quantities.

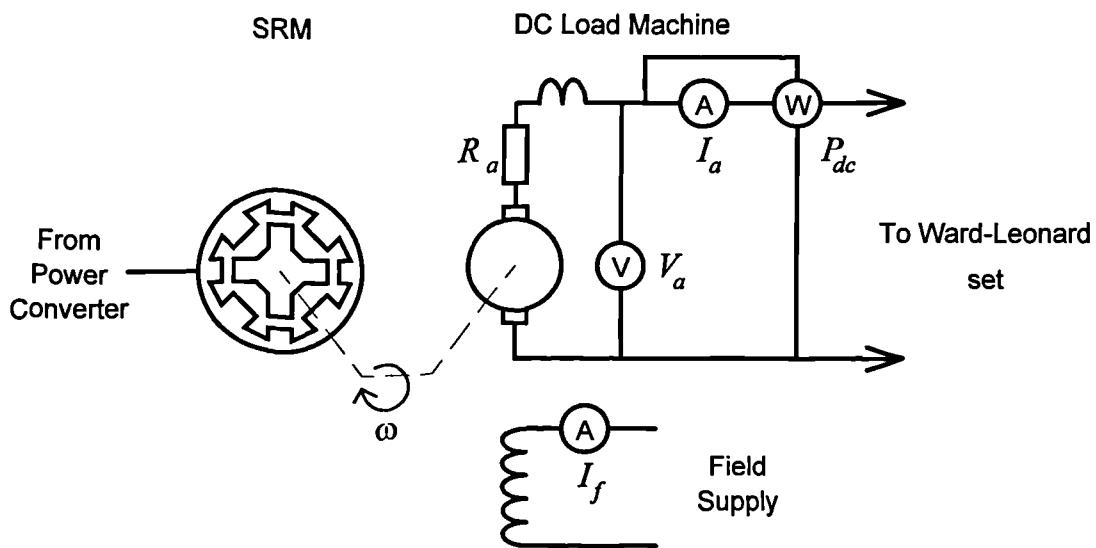


Figure B.5 Torque calculation circuit parameters.

Calculation of Torque from Air-Gap Power.

This method is based on the equation that air-gap power is the product of shaft torque and angular velocity. DC load machine the air-gap power is taken to be the armature terminal power plus the armature circuit losses. Armature terminal power is measured with a Sangamo Weston moving coil analogue wattmeter. Armature circuit losses are determined from the measured armature current and armature resistance. Switched reluctance machine electromagnetic torque is the sum of load machine air-gap torque and the torque due to windage, friction and iron loss. Torque is given by equation B.5, with loss torque calculated using equation B.6. Loss power is given by equation B.3 or equation B.4.

$$T_{elec} = \frac{P_{dc} + I_a^2 \cdot R_a}{\omega} + T_{loss} \quad (B.5)$$

$$T_{loss} = \frac{P_{loss}(\omega)}{\omega} \quad (B.6)$$

Calculation of Torque from Load Machine Torque Characteristics.

This method is based on calculating DC load machine air-gap torque from the torque constant and armature current. For a DC machine the torque constant is equal to the back emf constant, which is given by Equations B.1 and B.2. Armature current is measured using a low resistance shunt and voltage divider, calibrated to give the correct reading on a digital volt meter. Switched reluctance machine electromagnetic torque is the sum of air-gap torque and torque due to windage, friction and iron loss. Torque is given by equation B.7, with the loss torque calculated from equation B.6. The main assumption in this method is that the torque constant is independent of armature current. With a field current of around 1.3A, measurements have shown that the armature reaction is negligible for armature currents of up to 30A.

$$T_{elec} = I_a \cdot K_e (I_f) + T_{loss} \quad (B.7)$$

By measuring rotor angular velocity using a digital output from the rotor position encoder, and a counter timer with timing errors of less than 100ppm, the rotor velocity can be determined with sufficient accuracy. Comparison of the torque calculated by the two methods shows that in general the difference between them is less than 0.5%, results also show good correlation with what could be considered a good set of torque transducer readings.

B.4 Speed Controller Load Model

Connection of the DC load machine to the Ward-Leonard set, as used for constant speed tests, is not appropriate for testing the speed controller. A new loading arrangement is used to allow speed to vary with load torque over a wide range, this is achieved by connecting the DC load machine armature circuit to a resistive load bank. With a constant field current the emf is proportional to speed, and the power dissipated in the load bank proportional to speed squared. For a given torque constant, neglecting the armature resistance (which is negligible compared to the load resistor) the armature voltage is given by:-

$$V_{dc} = K_e \cdot \omega \quad (B.8)$$

With a load resistance of R the power dissipated P is:-

$$P = \frac{V_{dc}^2}{R} = \frac{(K_e \cdot \omega)^2}{R} \quad (B.9)$$

If frictional and windage torque are neglected then the shaft torque is given by:-

$$T_{dc} = \frac{P}{\omega} = \frac{K_e^2}{R} \cdot \omega \quad (B.10)$$

Using a first order system model for the switched reluctance machine and DC load machine with a combined inertia of J , the system equation is:-

$$T_{srm} = T_{dc} + J \cdot \frac{d\omega}{dt} = \frac{K_e^2}{R} \cdot \omega + J \cdot \frac{d\omega}{dt} = K_l \cdot \omega + J \cdot \frac{d\omega}{dt} \quad (B.11)$$

The Laplace transfer function of the load system derived from this equation is:-

$$\frac{\omega}{T_{srm}}(s) = \frac{1}{K_l + sJ} \quad (B.12)$$

Experimental measurements of the speed response to a torque step input showed that the load time constant is approximately 375ms. At nominal field current and load resistance the steady state gain is 22rpm/Nm.

Appendix C

Electrical and Electronic Circuits

C.1 Load and Measurement Circuit

The DC and line frequency power circuit arrangement of the experimental drive and load is shown in Figure C.1.

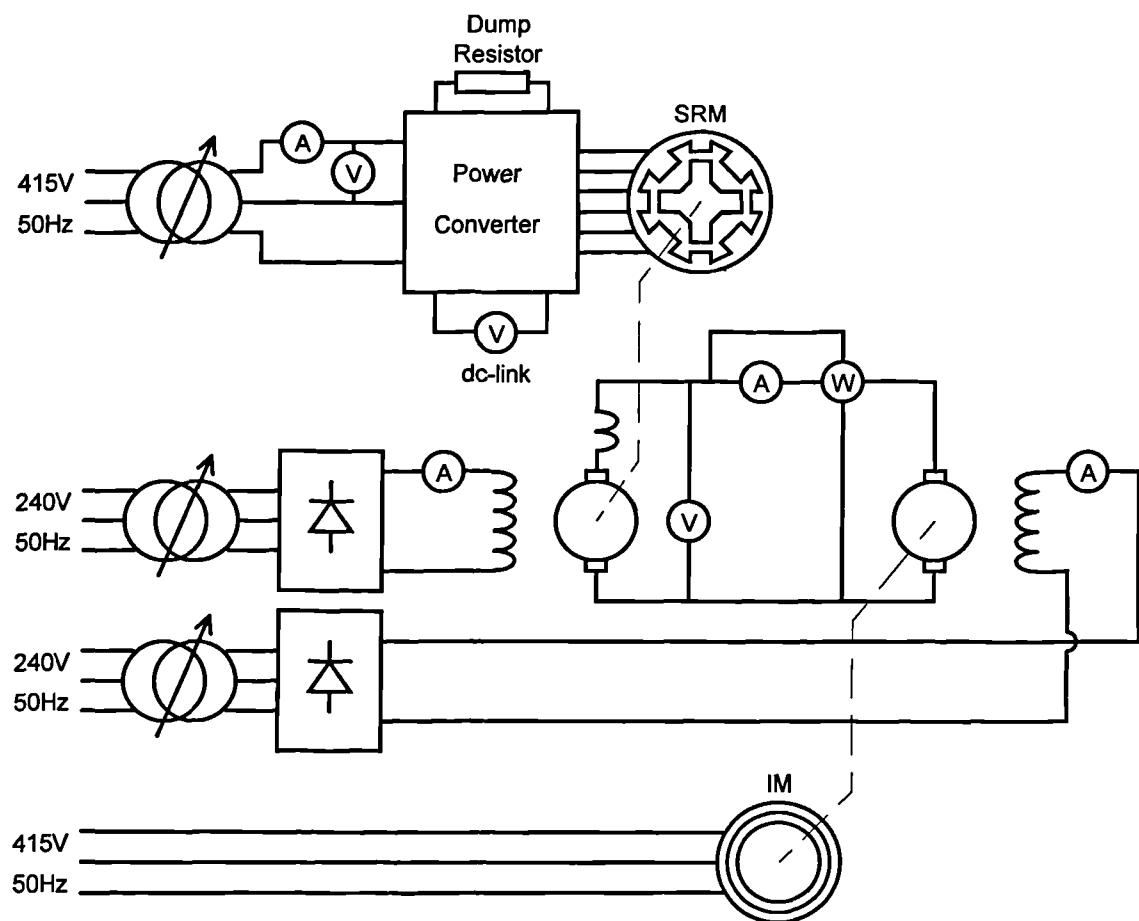


Figure C.1 Load and measurement circuit.

Whilst motoring the experimental switched reluctance drive is supplied from a three phase 50Hz supply. A variac allows the dc-link voltage to be adjusted to the desired level. To provide a flexible and easily controlled loading arrangement, a brushed DC machine is mechanically coupled to the experimental switched reluctance machine. This is a separately excited machine with the field supplied from a three phase supply, via a

variac and rectifier. Field weakening allows operation above base speed. The load machine armature circuit is connected to another DC machine which forms part of a Ward-Leonard set, allowing power to be returned to the AC supply. The Ward-Leonard set uses an induction motor connected to the three phase 50Hz supply, hence it runs at an approximately constant speed. DC machine armature voltage is controlled by adjusting the field circuit voltage, and hence current. As with the DC load machine, a variac and rectifier supplies the field. Both DC machines have sufficient field inductance to give low current ripple with a three phase supply. All ammeters and voltmeters are moving coil or moving iron analogue meters, except the ammeter in the DC machine armature circuit. This meter comprises of a $1.5\text{m}\Omega$ shunt with calibrated voltage divider, feeding a digital voltmeter. Not shown in Figure C.1 is the isolating transformer between the utility supply, and the power converter and DC machine supplies. This allows the power converter dc-link negative voltage rail to be earthed, simplifying measurements using earthed oscilloscopes, and adding a degree of safety.

Load and measurement circuit power ratings are given in Table C.1.

Parameter	Value	Units
Converter dc-link voltage	500	V(DC)
Converter input voltage	355	V(RMS)
Converter input current	16	A(RMS)
DC machine armature voltage	240	V(DC)
DC machine armature current	50	A(DC)
DC machine field voltage	240	V(DC)
DC load machine field current	1.6	A(DC)
DC supply machine field current	4	A(DC)
Induction machine voltage	415	V(RMS)
Induction machine frequency	50	Hz

Table C.1 Load circuit power ratings.

C.2 Power Converter

A power converter is required to perform several functions in a switched reluctance drive; rectification of the utility AC supply, energy storage, and control of the applied phase voltages to achieve the desired current. Initial static torque measurements showed that the fully-pitched winding switched reluctance machine could produce twice the shaft torque of its single-tooth winding equivalent, for the same phase current. For this reason it was decided that the power converter should be rated at twice the single-tooth winding machine power. Initial design parameters for the power converter were; peak phase current 20A, motoring power 15kW, generating power 10kW. The experimental power converter comprises of several blocks, as shown in Figure C.2; the input rectifier and filter, the dc-link dump for clamping the dc-link voltage during generation or braking, six bridge legs - two for each of the three phases, and the auxiliary power supplies. Each of these blocks is described in more detail in the following sections.

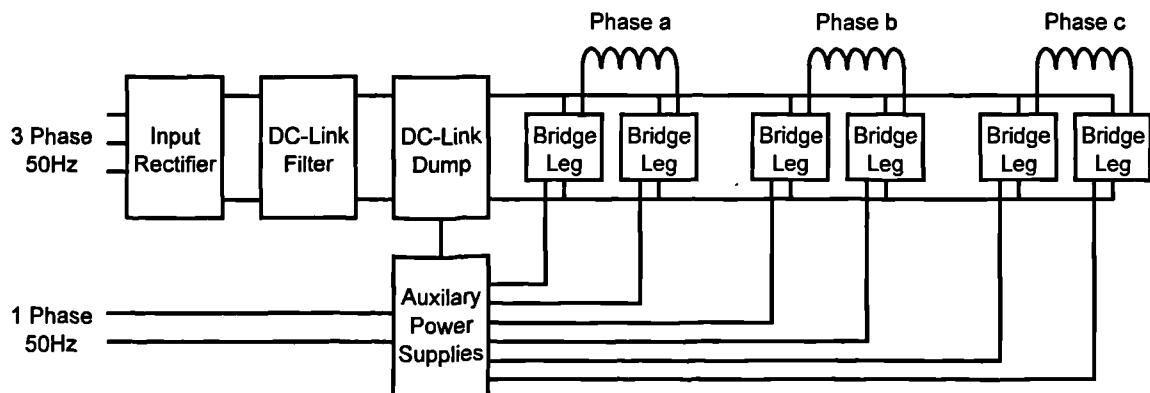


Figure C.2 Power converter block diagram.

C.2.1 Input Rectifier and DC-Link Filter

The first stage in the power converter is the input rectifier and dc-link filter, this converts the AC supply into DC, provides an energy store, and presents a low impedance to the inverter bridge legs. A three phase AC supply was chosen as this can easily be adjusted using a variac, is suitable for the design power rating, and can provide a low ripple dc-link voltage with very much smaller energy storage capacitors than required with a single phase supply. As Figure C.3 shows, the rectifier is preceded by a contactor to allow switching of the supply, and fuses to protect the rectifier in the event of a short circuit across the dc-link. The dc-link voltage was determined primarily by the bridge leg IGBT voltage rating (see section C.2.2) and set at 500V. This requires an input voltage of 355V line to line. Assuming an infinite dc-link inductance, so that the rectifier output current is DC, the rated power gives an input current of 16A RMS per phase. The three phase rectifier is actually formed from three single phase bridge

rectifiers, International Rectifier type 25MB100A. These were used as they are much cheaper than a three phase unit of the same rating, and have the advantage of better thermal performance. The dc-link inductor performs two functions: to form a second order filter with the dc-link capacitors so as to attenuate input voltage ripple, and to help limit input supply current peaks during initial charging of the capacitors and whilst running. The dc-link filter has a cut off frequency of 60Hz for nominal component values, which gives 28dB of attenuation at 300Hz the fundamental ripple frequency. As shown in Figure C.3 the dc-link capacitance is formed from four capacitors in a series-parallel combination. Series connection is required because electrolytic capacitors are not readily available with the required voltage rating. As this halves the effective capacitance, two pairs are used in parallel. Resistors are connected across each pair of capacitors to ensure equal voltage sharing, and discharge the link when the supply is switched off.

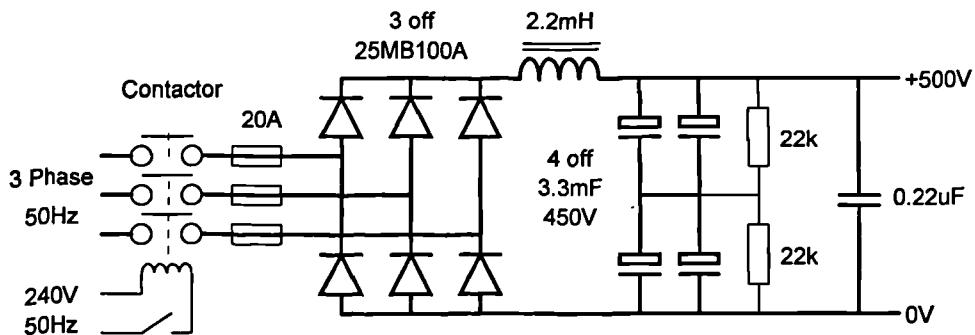


Figure C.3 Input rectifier and dc-link filter.

C.2.2 Bridge Legs

For maximum flexibility it was decided to use a 'H' bridge per phase to drive the switched reluctance machine. Each 'H' bridge is made up of two bridge legs, each of which in turn is composed of two series connected switches. The initial specifications were for a converter rated at approximately 20A per phase, with a nominal dc-link voltage of 500-750V. It was decided that for good drive performance the phase current bandwidth should be wide, and the current ripple low. This implies a relatively high switching frequency: a nominal figure of 10kHz was chosen as a good compromise between bandwidth and switching loss. Given these basic parameters, selection of the power devices is relatively straight forward. Power MOSFETs have a relatively high conduction loss for devices rated above 500V, and GTOs are too slow for the desired switching frequency. This leaves Insulated Gate Bipolar Transistors (IGBTs) as the most obvious choice. Of the devices readily available, there is a large difference in switching speed and loss, between devices rated at 600V, and those rated at 900V and above. Switching at 10kHz, speed and switching loss are important considerations, it

was therefore decided to use 600V rated devices. This puts a practical limit on the dc-link voltage of about 500V.

The IGBTs chosen are International Rectifier type IGPC50U, a 600V 55A device housed in a TO247 package. This device has no reverse blocking or conduction capability and therefore must be used with a free-wheeling diode for inductive loads. The diode chosen was a Harris RUR3060, a 600V 30A device, with a ultra-fast (55ns) soft recovery. The fast and soft recovery limits switching loss and EMI. To keep device voltages to a minimum during high speed switching power circuit inductance must be minimised, hence in addition to the high frequency de-coupling of the dc-link there are snubbers placed across the dc-link very close to the power switches. High current sections of the bridge leg circuit are shown in bold in Figure C.4. Calculations of conduction and switching loss, for the IGBT and diode using data sheet values, showed that under worst case conditions; switching 20A at 10kHz, with a dc-link voltage of 500V, the maximum loss in one IGBT is 85W, and the total loss for a 'H' bridge is 145W. Using a thermally conductive insulator (Warth K230 material) calculations showed that the heatsink temperature should be kept below 55°C so as not to exceed the maximum device junction temperature. This implies a heatsink of 0.2°C/W, which with fan cooling requires a heatsink with a free air rating of less than approximately 0.6°C/W. The heatsinks chosen have a thermal rating of 0.33°C/W in free air.

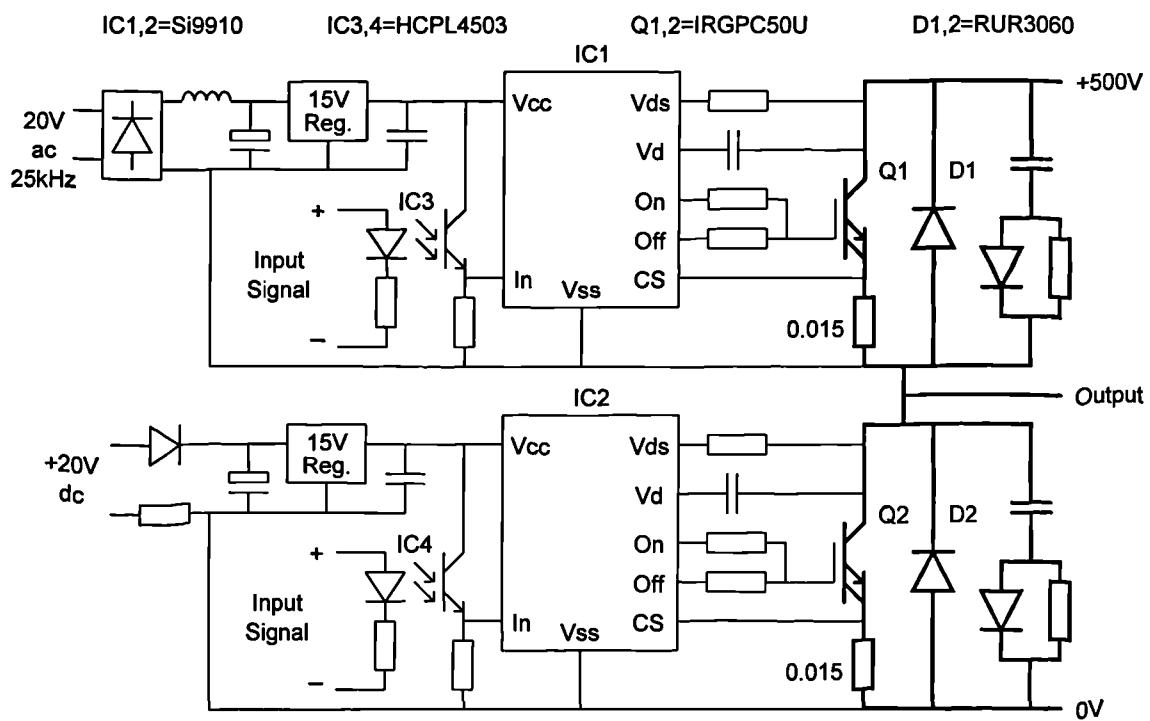


Figure C.4 Converter bridge leg circuit.

The IGBTs are driven from Si9910 drivers which, as shown in Figure C.4, have several connections to the power circuit. This driver provides protection against over current (CS input), cross conduction (Vds input), and can limit the switching speed to limit dV/dt (Vd input). It has separate 'on' and 'off' driver outputs so that the gate resistors can be optimised for each case. A 4.7Ω 'on' resistor limits turn on speed for gentle diode recovery, and a 1Ω 'off' resistor gives a fast low loss turn off. Power supply under voltage protection ensures that devices remain off whilst supplies are ramping up and down.

Requirements of noise immunity, isolation from power circuit voltages, and protection in the event of a power circuit fault, mean that control signals to the IGBT drivers should be optically or transformer isolated. For a wide variation in duty cycle optical isolation is the most appropriate method. The opto isolators used are type HPCL4503, these have a common mode immunity of $15\text{kV}/\mu\text{s}$ which makes them suitable for use in high speed power circuits.

For the low side drivers a non-isolated power supply is used as the auxiliary power circuit ground and main power circuit grounds are common. The high side supply is transformer isolated (see section C.2.4) with the rectifier and filter on the IGBT driver board. High and low side supplies have a 15V regulator to provide a stable supply to the driver circuit. On the low side supply a 10Ω resistor in the negative rail ensures that power circuit currents do not flow through the auxiliary power supply wiring.

C.2.3 DC-Link Dump

The function of the dc-link dump circuit is to limit the maximum dc-link voltage when energy is returned to the storage capacitors from the machine, whilst braking or generating. This is necessary because the single quadrant supply rectifier is unable to return energy to the AC supply. The circuit was designed to dissipate a maximum of 10kW at the nominal dc-link voltage of 500V, an average current of 20A. High current paths in the circuit are shown as bold in Figure C.5. The energy is dissipated in a 10kW load bank. An IGBT, International Rectifier type IRGBC40F, is used to switch in the dump resistor using a hysteresis control circuit. The IGBT is a medium speed device with low conduction loss, well suited to the low dump circuit switching frequency of approximately 500Hz. As the load bank and wiring is inductive, a free-wheeling diode is required. A Harris RUR3060 30A 600V ultra fast diode is used for this purpose. Worst case total dissipation, calculated from data sheet values is in the order of 40W, requiring a heatsink of $0.5^\circ\text{C}/\text{W}$. DC-link voltage is compared to a reference value by the LM311 and switches the IGBT on when the dc-link voltage reaches 515V, switching

off again when the voltage falls to 505V. A CMOS buffer is used to drive the IGBT gate, with no external gate resistor required. The auxiliary power supply provides a 20V supply, which is regulated on board to 15V.

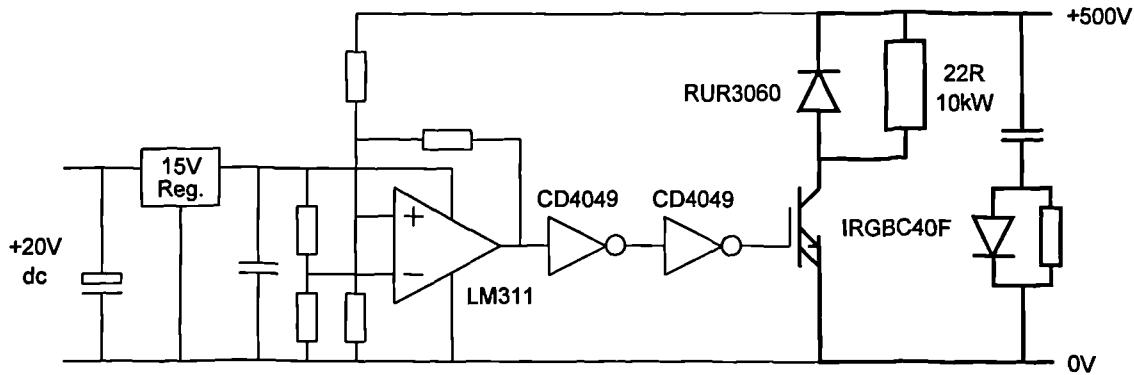


Figure C.5 DC-Link Dump Circuit.

C.2.4 Auxiliary Power Supplies

The auxiliary supply provides power to all the circuits within the power converter enclosure, including the current measurement board. As shown in Figure C.6 a line frequency transformer and rectifier is used to derive unregulated $\pm 20V$ rails. These are then regulated using standard 78 and 79 series voltage regulators, to provide the $+5V$, and $\pm 15V$ supplies required by the phase current measurement board. The $+20V$ output supplies the dc-link dump circuit, and all six low side IGBT gate drive circuits. The six isolated high side IGBT gate drives are supplied from a high frequency inverter, also powered from the $+20V$ supply. A transformer with eight windings is used to provide isolation between the six high side supplies, the primary, and the feedback winding. It has a ferrite core, type ETD34, and isolation between each winding is designed for 1kV. Note that only two of the six output windings is shown in Figure C.6. The inverter is controlled by a TL494 PWM power supply control chip, switching frequency is 25kHz, and the feedback winding is used to adjust the duty cycle of the two complementary outputs so as to keep the output voltage stable. Power is switched by an SGS Thompson L6203 full bridge driver. This chip has contains a MOSFET 'H' bridge rated at 3A, with integrated gate drives and protection. Bridge output voltage is applied to the transformer primary via a DC blocking capacitor.

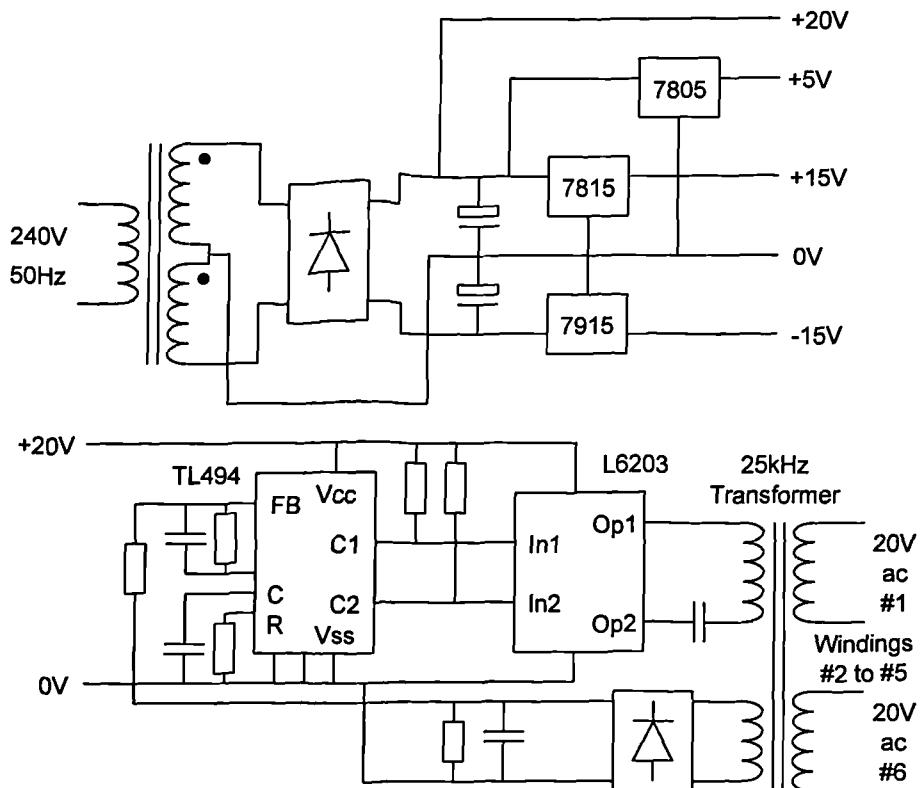


Figure C.6 Auxiliary power supplies.

C.3 Measurement Circuits

In the experimental switched reluctance drive only three parameters are measured: rotor position, phase currents, and the dc-link voltage.

C.3.1 Rotor Position Measurement

Rotor position is measured with an optical encoder, Gaebridge type 10BITCW, which has a 10bit Gray code output. This provides absolute position measurement with a resolution of 0.35° . The encoder requires only a +5V supply and gives a TTL compatible output. As the output frequency is low, and the encoder is electrically isolated from the power circuits, a simple screened cable is all that is required to connect the encoder to the control electronics.

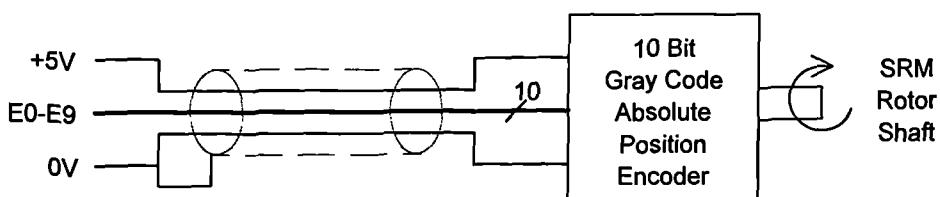


Figure C.7 Rotor Position Encoder.

C.3.2 Phase Current Measurement

Phase current is measured using a Hall effect current transducer, Heme type FP50P. This meets the basic phase current measurement requirements of a peak current rating of 25A, bandwidth of DC to at least 10kHz, and isolation from the power circuit. The FB50P has a 100kHz bandwidth, and by using a two turn primary, the scaling is conveniently altered to give 50mA output for a phase current of 25A. A precision load resistor (100Ω) converts the output current into a voltage which is then multiplied by the opamp circuit, offset correction is also added at this point. The opamp used is an AD711, this is a precision JFET input opamp with the following features; low input offset current and voltage for precision operation and low drift, sufficient gain bandwidth for the application, and an output capable of driving the analogue to digital converter input.

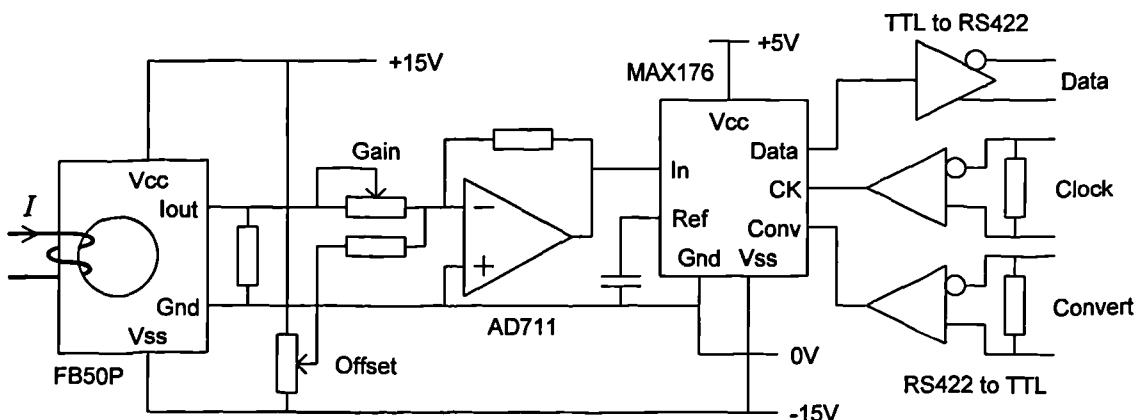


Figure C.8 Phase current measurement circuit.

In the circuit shown in Figure C.8 the left hand side is concerned with measurement of the phase current, and the right hand side the conversion to a digital value and its transmission to the controller. By placing the analogue to digital converter next to the current transducer little noise is coupled in from the power circuit which keeps errors low. The digitised value can then be transmitted to the controller with little chance of corruption due to system noise. By using a serial data link the number of connections required is substantially reduced, and RS422 differential drivers can be used without a large overhead in wiring and cost. The ADC used is a Maxim MAX176, this is a 12bit converter with built in reference, bipolar input range, serial output, and a conversion time of $4\mu\text{s}$ when clocked at 4MHz. In this application the clock frequency used is 2MHz, giving a conversion time of $8\mu\text{s}$. RS422 compatible differential drivers and receivers are used for the data, clock, and conversion start signals. These give excellent noise immunity and are suitable for driving long cables at high speed. Note that Figure C.8 only shows one of the three channels, the clock and convert RS422 receivers are shared by all three channels, conversions thus start simultaneously.

C.3.3 Over-Current Detection

Associated with the current measurement circuit is the over current detection circuit, shown in Figure C.9. This is designed to detect if the phase current has exceeded a pre-set value, if for example the control software fails. The input to this circuit comes from the opamp output, shown in Figure C.8. This voltage is filtered to remove noise and prevent spurious tripping, before being compared with positive and negative reference voltages. The comparators used are type LM393 which have an open collector output, this allows all three channels to be wire-OR-ed together (Figure C.9 only shows one channel) to give a single over current output. The current trip level is set at $\pm 24A$.

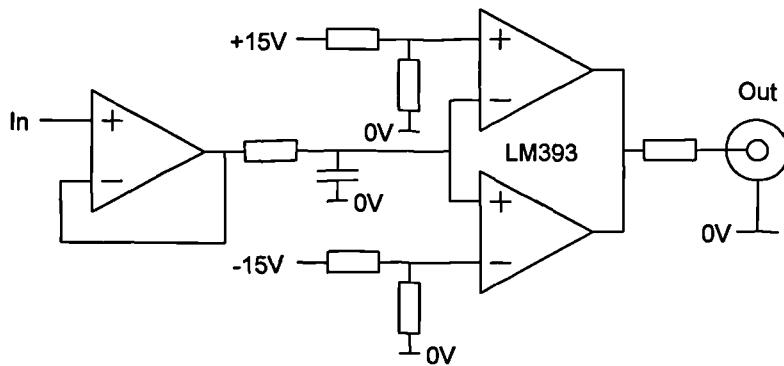


Figure C.9 Phase over-current detection circuit.

C.3.4 DC-Link Voltage Measurement

The dc-link voltage divider gives a 100:1 attenuation ratio so that the voltage presented to the MSAD12 analogue input board is within the input range of 0 to 10V. Two resistors in series are used for the top end of the divider to reduce the voltage stress and thermal dissipation in each resistor. The bottom end zener diode provides protection for the analogue input in the event of one of the top end resistors going short circuit, or a bottom end resistor going open circuit. Two bottom end resistors are used to allow easy trimming to the correct value for 100:1 attenuation.

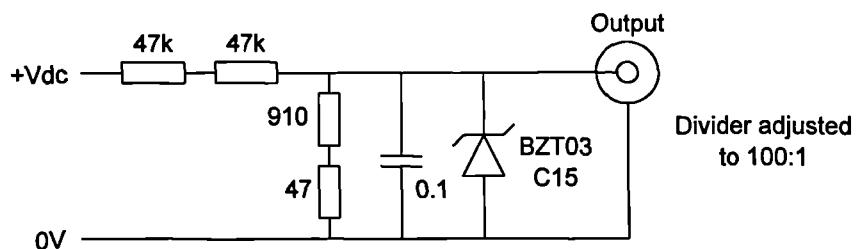


Figure C.10 DC-Link Voltage Divider.

C.4 Controller Circuits

All the drive controller boards are housed in a 6U 19" rack, fitted with power supply and cooling fans. A VME bus connects all the sub-units on the upper connector (VME P1) with additional connections on the lower connector where required. There are several control sub-units as follows; front panel controls and analogue input, analogue output, phase current input, position encoder input and system control, 68030 processor board, TMS320E14 processor board with PWM output, TMS320C31 processor board, and the MSAD12 analogue to digital converter board. Each of these will be described in the following sections.

C.4.1 Front Panel Controls and Analogue Input

To manually adjust control parameters there are two front panel rotary controls and two switches. The switches and potentiometers are arranged to give a variable input voltage to the MSAD12 analogue input board. The dc-link voltage divider input has filtering and diode clamps to protect the analogue input circuitry.

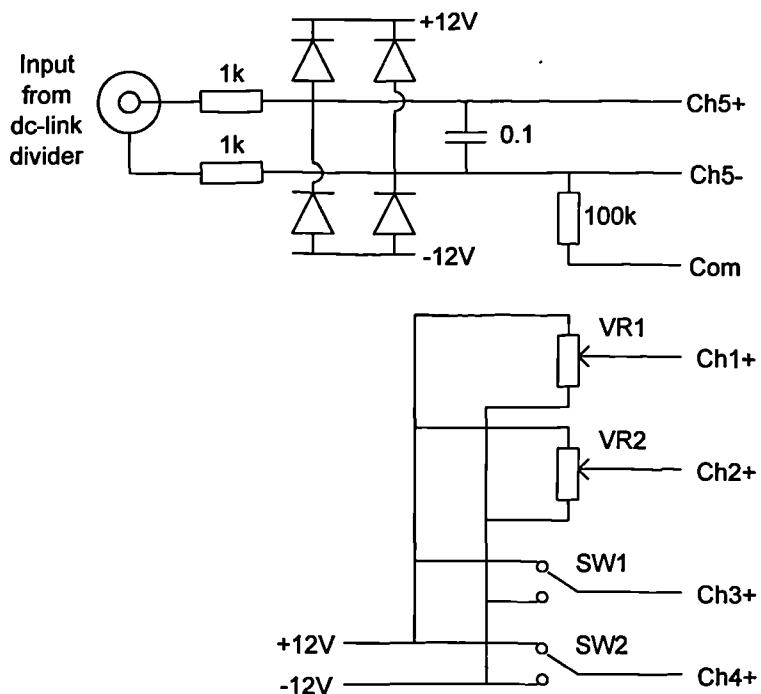


Figure C.11 Front panel controls and DC-link voltage input.

C.4.2 Analogue Output

The analogue output sub-unit provides six analogue outputs for real time monitoring of drive parameters. A 68230 Parallel Interface Timer interfaces to the VME bus, and provides parallel data outputs and control lines to the digital to analogue converters. A

16L8 PAL is used for address decoding. Figure C.12 shows two of the six output channels, only one 68230 is used to feed data to all three DAC8228 analogue to digital converters. A DAC8228 is a dual 8bit DAC with current output. The output buffer amplifiers are configured to give current to voltage conversion and add offset, giving an output range of $\pm 2.5V$. A 2.5V reference is derived by dividing down and buffering the output from a REF50Z, +5.0V reference. Precision 0.1% resistors are used for the voltage divider and output amplifiers to give good accuracy and low drift. The relationship between binary value and output voltage is as follows; 00H gives -2.5V, 80H gives 0.0V, and FFH gives +2.5V.

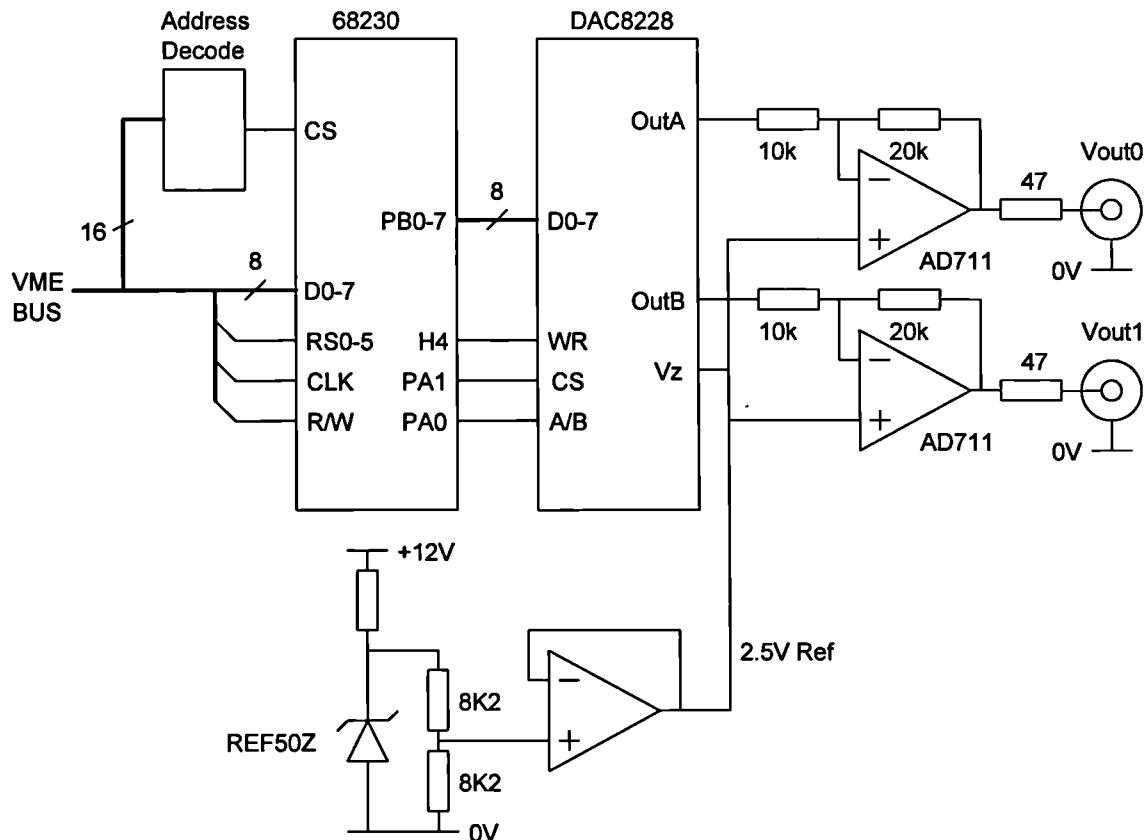


Figure C.12 Analogue output circuit.

C.4.3 Phase Current Input

Serial output analogue to digital converters used for phase current measurement are controlled by the phase current input circuit. This also converts the serial data back to parallel format for reading by the TMS320E14 processor. The ADC controller is implemented as a state machine in a 16R8 PAL. Upon receiving a conversion start signal from the processor, the correct 'clock' and 'convert' signal sequence is sent to the phase current ADCs on the phase current measurement board, and the serial data received latched into registers. Once conversion is completed the processor is informed

via the DV - data valid signal. TTL to RS422 conversion is required for the clock and convert outputs, and RS422 to TTL conversion for the incoming serial data. Two cascaded 74HC595 serial to parallel converters are used for each channel. These have separate shift and data registers which allows the processor to read the data at any time.

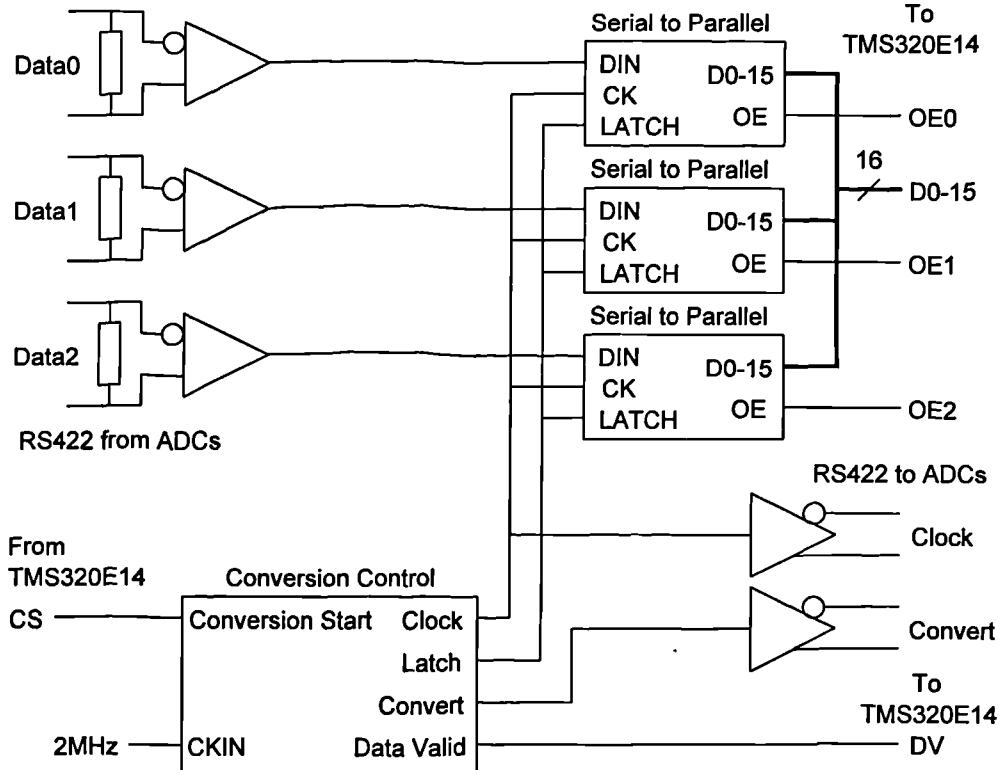


Figure C.13 Phase current input circuit.

C.4.4 Encoder Input and System Control

The main functions of this sub-unit are to enable the rotor position encoder to be accessed from the VME bus, and provide control signals to other parts of the system. A 68230 Parallel Interface Timer provides most of these functions. Address decoding is done using a 16L8 PAL. Encoder input data is first filtered to remove noise, and then buffered. Changes in rotor position are detected using a series of XOR gates, these provide one of the inputs to the latch control circuit. Each time a change in rotor position is detected the new position is latched into the encoder latch after a delay to allow the data to settle. Upon receipt of an interrupt signal from the processor, the latch control circuit ensures that the data in the encoder latch has settled before pulsing the H3 input on the 68230. The H3 input is configured to latch data on the Port A and B inputs into internal registers, and produce an interrupt request on the VME bus. The 68230 provides an interrupt vector during the VME bus interrupt acknowledge cycle. The encoder input also generates the tacho output signal for external speed

measurement, and the two encoder least significant bits go to the TMS320E14 for processor speed measurement.

Parallel data inputs and outputs from the 68230 enable the TMS320E14 to be controlled from the VME bus, and provide handshaking signals for data exchange. Two other functions performed on this board are latching of over-current and watchdog time-out signals.

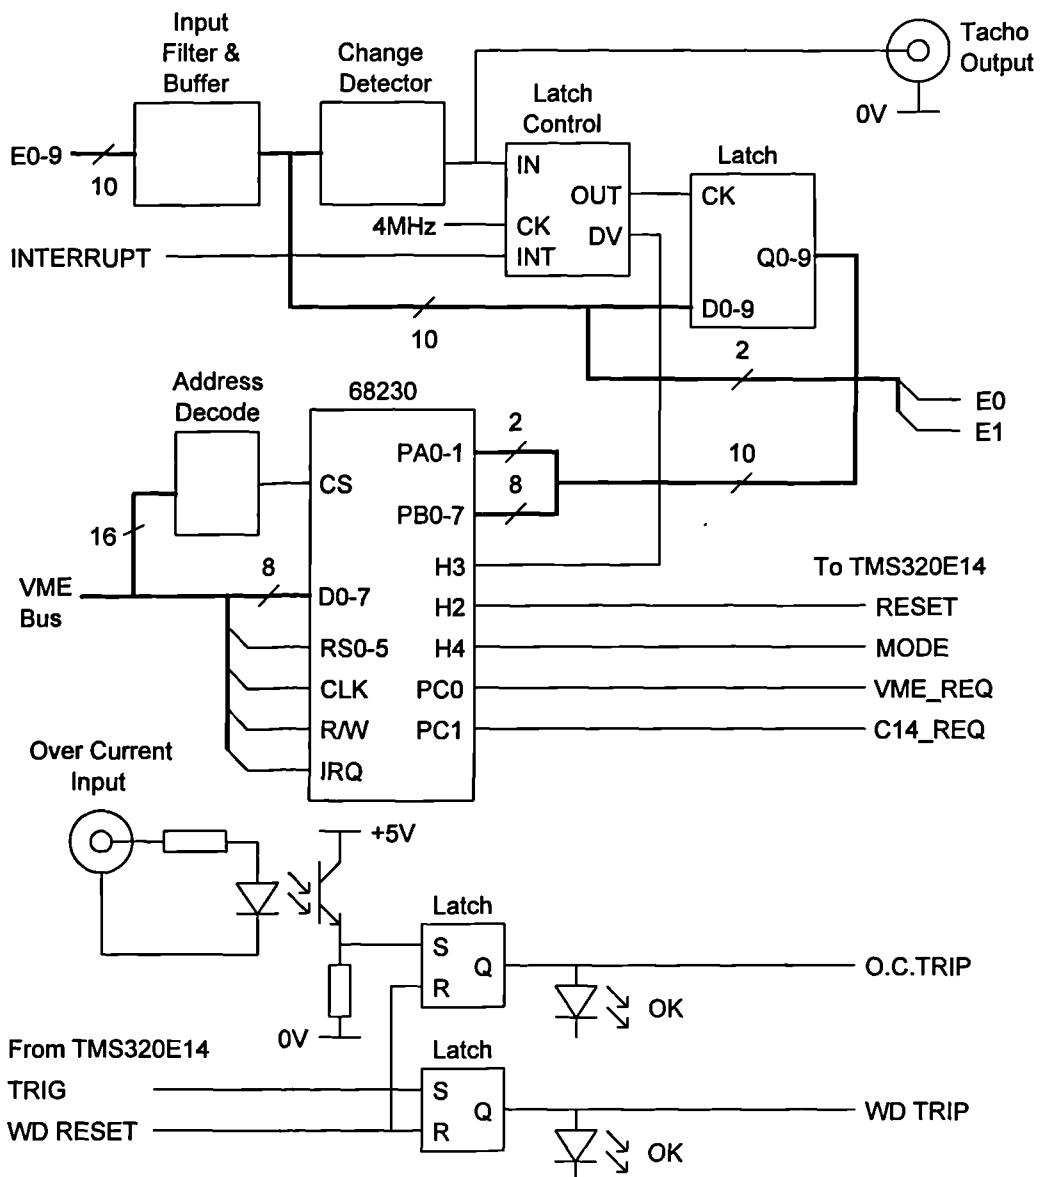


Figure C.14 Position encoder input and system control sub-unit.

C.4.5 TMS320E14 Processor Board

The TMS320E14 processor board provides enough processing power to implement digital PID current control, and has integrated timer counters for PWM generation and rotor speed measurement. Figure C.15 shows the external memory configuration, and Figure C.16 the PWM outputs and other peripheral functions.

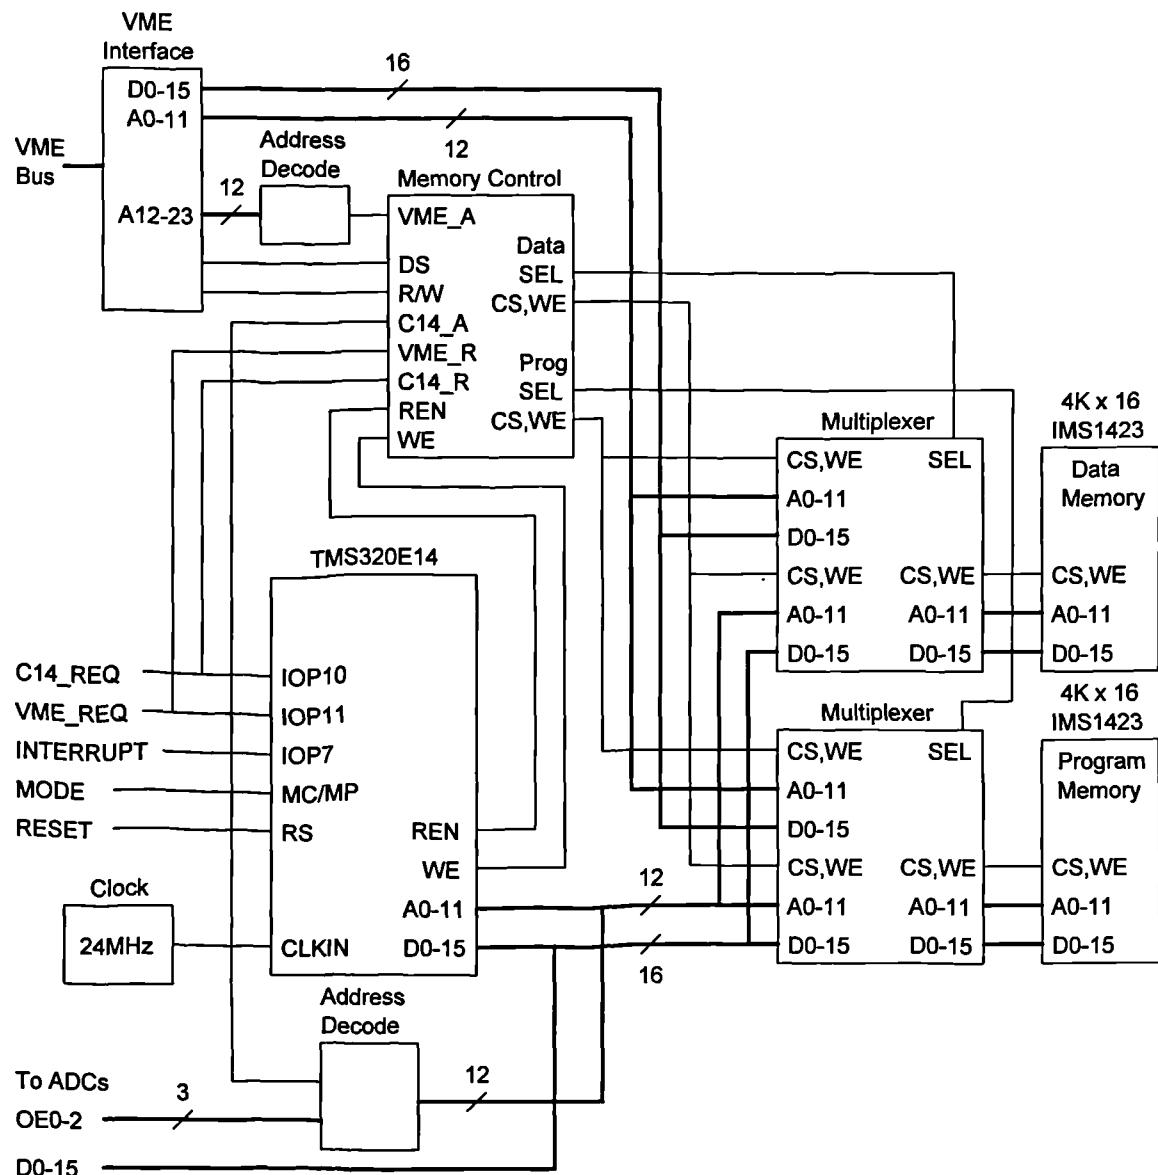


Figure C.15 TMS320E14 Memory interface circuit.

The TMS320E14 is a 16bit fixed point Digital Signal Processor, with a range of peripherals integrated on chip including several counter timers. When clocked at 24MHz, the processor cycle time is 166ns, and it executes a multiply-accumulate instruction in a single cycle. The TMS320E14 has a Harvard architecture, i.e. with separate data and code memory spaces. Although the TMS320E14 has internal code memory this is not used as it cannot be reprogrammed in situ. As configured the

processor has four address spaces. There are 4K words of external program memory, this can be accessed from the VME bus whilst the processor is held in reset so that code can be downloaded to the TMS320E14. Once out of reset the processor has sole access to this memory. The data memory area is split into three areas: one internal, and two external. The internal data area has 256 words of memory which the processor can access using single cycle data move and load instructions. The external memory area comprises of three addresses used to access the phase current registers (see section C.4.3) and 64 words of shared memory. Both these areas are accessed using IN and OUT instructions which take three cycles to execute. The shared data memory area is used for passing data to and from the VME bus, and is accessible from the TMS320E14 and the VME bus. Memory access is controlled using the C14_REQ and VME_REQ hand-shaking lines. Address decoding for the VME bus, and all TMS320E14 memory access, is controlled using 16L8 PALs.

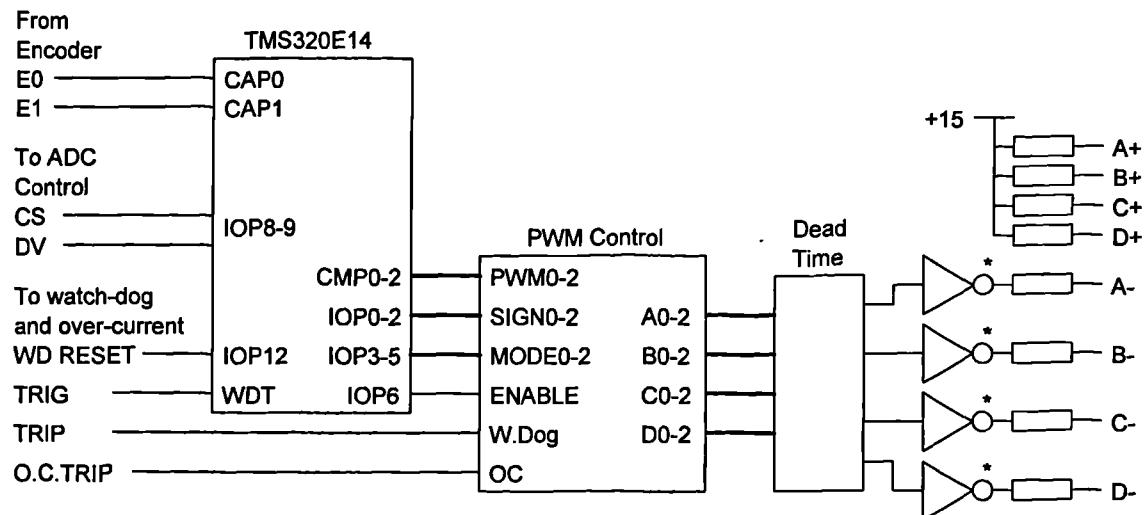


Figure C.16 TMS320E14 Control and PWM Generation.

The TMS320E14 has several counter timers which are used for three functions in this application; PWM generation, rotor speed measurement, and the processor watchdog. The CMP0-2 PWM outputs use the high resolution mode which gives a 16bit pulse width register and a minimum resolution of 40ns. For a PWM output frequency of 10kHz, a 100 μ s period, the minimum timing resolution limits the PWM value to just over 11bits, rather than the 16bits possible with a lower output frequency although this is of no consequence. The PWM period register is used to fix the system timing by generating an interrupt signal to the rest of the system on a 100 μ s cycle.

The actual switching times of the four devices in each power converter 'H' bridge are determined by the PWM control logic which is implemented in 16L8 PALs. In addition to the PWM signals, PWM0-2, there are the sign bits, SIGN0-2, which determine if the

output voltage sign, and the mode bits, MODE0-2, which determine the mode of operation, for example; bipolar or unipolar chopping. There are also two protection inputs from the watchdog and over current detector which both switch off all power devices. The PWM outputs go through a dead-time generator, and are buffered to drive the IGBT gate drive opto couplers. Note that only the output buffers for one phase are shown in Figure C.16.

Capture inputs, CAP0 and CAP1, are used for rotor speed measurement. These inputs are used to latch the capture timer value when there are changes on the encoder E0 or E1 inputs. Subtraction of successive timer values allows the speed to be determined.

The third timer function is a software watchdog: this timer is reset by writing a sequence of values to a specific location. If the timer 'times out' a pulse is generated on the WDT output, this is latched and used to disable the PWM outputs via the TRIP input.

C.4.6 Commercially Produced Boards

In addition to the circuits designed and built by the author there are three commercially produced boards, all housed within the VME rack.

MVME147 Processor Board

This board is produced by Motorola and is the VME bus master. As such it controls all data transfers on the VME bus, and services all interrupts. It houses a 32bit processor and a large amount of memory. Serial ports allow programs and data to be transferred to and from the host PC. Monitor and de-bugger firmware allows data to be transferred and registers inspected.

The processor board has the following features:-

- MC68030 32bit CISC Processor
- MC68881 Maths Co-processor
- Clock speed 25MHz
- 8MBytes of memory
- VME Bus master interface
- Four RS232 serial ports

DBV31A Processor Board

This processor board is made by Loughborough Sound Images and carries a 32bit DSP. The processor can execute most instructions, including floating point multiplication, in a single 60ns cycle. The board only has a VME bus slave interface and therefore cannot initiate data transfers on the bus. During data transfers from memory to the VME bus the processor is halted. It does however have a single bi-directional port which allows transfer of a single data value without halting the processor. A high speed emulator interface allows code to be loaded and enables symbolic de-bugging.

The processor board has the following features:-

- TMS320C31 32bit floating point Digital Signal Processor
- Clock speed 33MHz
- 128k of 32bit memory
- VME Bus slave interface
- VME Bus data transfer register
- High speed emulator serial port

MSAD12 Analogue Input Board

This board is produced by the Matrix Corporation and has a single analogue to digital converter. An input multiplexer allows one of eight channels to be selected for conversion. The board can generate a vectored interrupt upon completion of conversion, and can automatically select the next input channel for conversion.

The board has the following features:-

- 8 Differential input channels
- 12bit resolution
- 50 μ s conversion time
- $\pm 10V$ input range

Appendix D

Controller Software

D.1 Processor Configurations

There are two basic processor configurations used for the experimental switched reluctance drive: one using two control processors, and one using three control processors. In the two processor configuration a third processor may be used, but only for monitoring purposes, not control.

D.1.1 Two Processor Configuration

In the two processor configuration the TMS320E14 and 68030 are assigned drive control functions. The TMS320E14 has three main functions; current control, speed measurement, and system timing. Rotor position measurement, current reference calculation, and house keeping functions are performed by the 68030. The third processor is only there to perform the numerically intensive calculation of flux-linkage and torque for monitoring purposes. Figure D.1 shows the tasks each processor performs and the parameters passed between them.

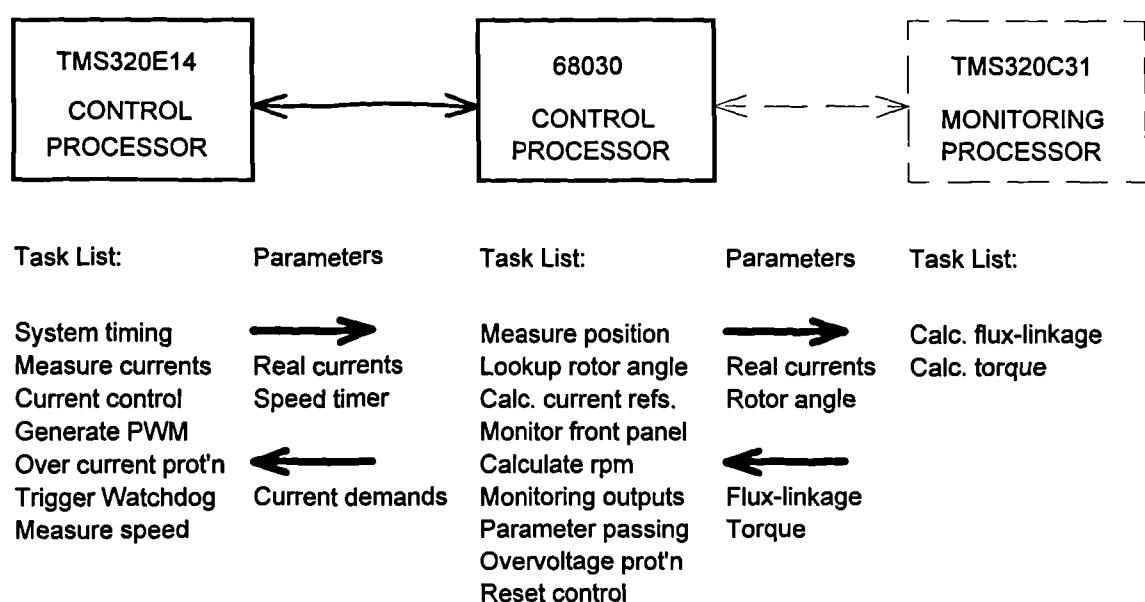


Figure D.1 Two control processor configuration.

System timing is defined by the TMS320E14 PWM timer, this generates an internal interrupt at which point the processor starts the ADC conversion for current measurement, and signals the 68030, via a VME bus interrupt, that data should be passed. When using 50µs current sampling (twice per cycle) another internal interrupt is used to start the second ADC conversion process.

When using the third processor for monitoring, interrupts are used to signal the TMS320C31 that data is ready, and the VME bus that calculations are complete. This happens as a background process and is asynchronous to the main control task.

D.1.2 Three Processor Configuration

The three processor configuration allows more complex controllers to be implemented by using the powerful TMS320C31 DSP. In this configuration the TMS320E14 measures phase currents, checking for over-current, and measures the rotor speed. The 68030 measures rotor position and performs house keeping functions. This leaves all the main control and monitoring functions to be performed by the TMS320C31. Tasks performed by the processors and parameters passed between are shown in Figure D.2. An unfortunate side effect of the 68030 being the only VME Bus Master is that the actual phase current values and the required PWM values, must be passed twice to get from the TMS320E14 to the TMS320C31.

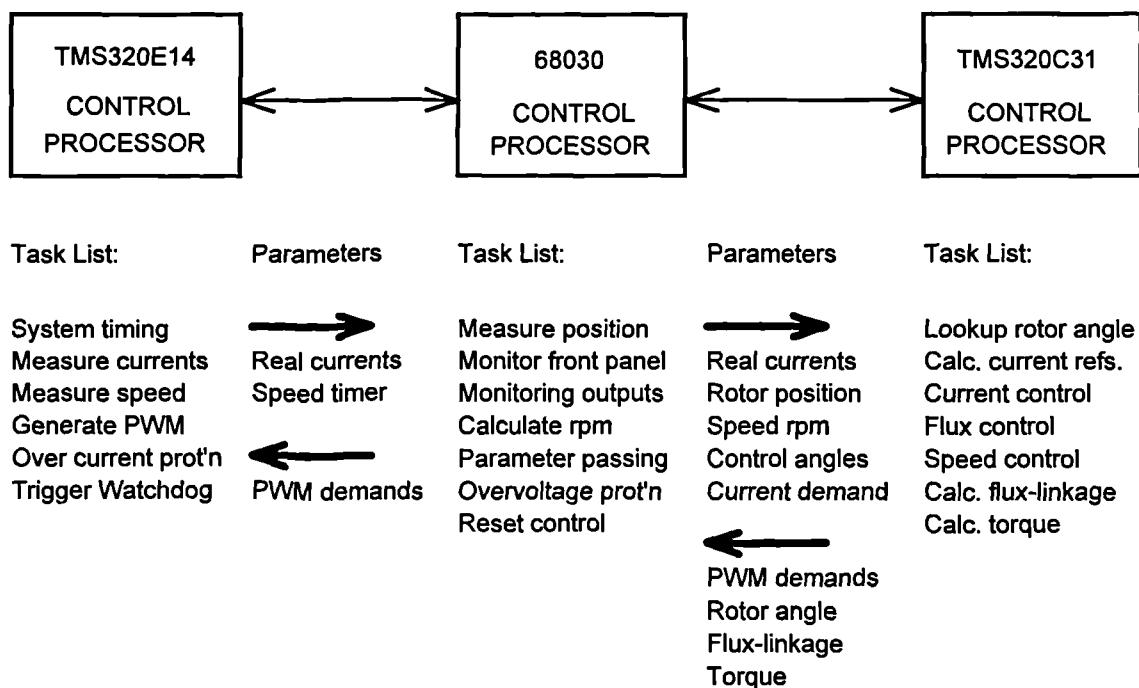


Figure D.2 Three control processor configuration.

As with the two processor configuration, system timing is defined by the TMS320E14 PWM timer interrupt. At the start of a cycle the TMS320E14 first starts the ADC conversion for current measurement, and signals the 68030, via a VME bus interrupt, that data passing is about to commence. The 68030 then gains access to the TMS320C31 data memory before transferring data to it, from the TMS320E14 and 68030. Data is also read back from the TMS320C31. When all the data has been passed the TMS320C31 bus is released, and an interrupt generated to inform it that processing can commence on the new data. This method was found to be more efficient than using interrupts to signal back and forth that data was ready, as less time is spent switching context.

D.2 Lookup Tables

Lookup tables are used to determine rotor angle, flux-linkage, torque, and phase reference values.

The rotor angle lookup table is a single dimension array, and converts from the Gray code position encoder value to a binary representation of rotor angle. Table reference is of the form:-

```
Rotor angle = Gray_to_Binary( Encoder value )
```

The table has 1024 entries, one for each possible encoder value.

Flux-linkage and torque lookup tables are used to enable real time calculation of the machine flux-linkage and shaft torque. The tables store measured values of flux and torque in two dimensional arrays. Note that the values are for the equivalent single-tooth winding parameters, so the actual fully-pitched winding values must be converted to their single-tooth winding equivalents, before looking up values in the tables. The two tables are referenced as follows:-

```
Flux-linkage = Current_to_Flux( Phase current, Rotor angle )
```

```
Torque = Current_to_Torque( Phase current, Rotor angle )
```

In each case the rotor angle index is in steps of three rotor encoder positions, (1.05° mechanical), with 21 positions corresponding to a rotation of half a rotor tooth pitch. The flux-linkage table has a phase current index of 0 to 40A in steps of 0.2A, and the torque table has a range of 0 to 24A in steps of 0.2A. The flux-linkage table therefore contains 4200 entries, and the torque table 2520 entries. Linear interpolation of table values is used to calculate the flux-linkage and torque for the actual current and

position. Symmetry of flux-linkage and torque about the aligned or unaligned position is used to calculate values for other rotor positions. No account is taken of rotor eccentricity effecting flux or torque.

Phase reference values may be in either current or flux-linkage. Tables are used for references which have a constant shape, e.g. sinusoidal, or a constant torque profile, and hold the reference value for each rotor position in a cycle. For unipolar waveforms there are therefore 128 entries (1024 positions divided by 8 cycles per revolution) and 256 entries for bipolar references. The tables are referenced as follows:-

```
Current reference = Position_to_Current( Rotor angle )
```

```
Flux reference = Position_to_Flux( Rotor angle )
```

Tables contain the reference for just one phase, other phase references uses the same table but the angles are offset by $\pm 1/3$ of a cycle.

D.3 Parameter Monitoring

There are six analogue output channels available on the experimental drive controller which can be used from monitoring a range of visible and invisible parameters. Direct measurement, numerical calculation, and lookup tables, are all used to derive the monitoring values. Table D.1 list all the monitoring values that are available, with their output range and scale.

Signal	Description	Range	Scale	Zero Value
Ia	Phase current	$\pm 20\text{A}$	10A/V	0.0A
I1	Equivalent current	$\pm 40\text{A}$	20A/V	0.0A
ψ_a	Flux-linkage	$\pm 2.5\text{Wb}$	1Wb/V	0.0Wb
ψ_1	Equivalent flux-linkage	$\pm 2.5\text{Wb}$	1Wb/V	0.0Wb
T1	Single phase torque	$\pm 125\text{Nm}$	50Nm/V	0.0Nm
Tq	Total torque	$\pm 125\text{Nm}$	50Nm/V	0.0Nm

Table D.1(a) Monitoring outputs available.

Signal	Description	Range	Scale	Zero Value
θ_{on}	On or advance angle	$\pm 360^\circ$ electrical	$140^\circ/V$	Aligned position
θ_{off}	Off angle or conduction width	$\pm 180^\circ$ electrical	$140^\circ/V$	240° Unipolar 120° Bipolar
θ_{rotor}	Rotor position	360° mechanical	$140^\circ/V$	Aligned position
ω	Rotor speed	2500rpm	1000rpm/V	0.0rpm
T _{run}	Processor run time	250 μ s	100 μ s/V	0.0 μ s

Table D.1(b) Monitoring outputs available.

D.4 Reference Calculation

Current or flux-linkage references for each phase can be calculated from lookup tables or equations. Lookup tables are used for waveforms which have a fixed shape, or which may be time consuming to calculate in real time. Equations allow the phase angle and conduction width to be varied with ease for basic unipolar and bipolar current references.

Typical code structure for lookup table current references is as follows:-

```
angle = Gray_to_Binary(encoder value)
angle = angle + advance angle
phase a ref = Position_to_Current(angle)
angle = angle + phase offset
phase b ref = Position_to_Current(angle)
angle = angle + phase offset
phase c ref = Position_to_Current(angle)
```

After looking up the rotor angle from the position encoder value, the overall phase advance angle is added. Each phase reference is then looked up, with phases offset in angle from each other.

When using equations there are usually three variables; the overall advance angle, current pulse width, and phase current magnitude. Typical code structure is as follows:-

```

angle = Gray_to_Binary(encoder value)
angle = angle + advance angle
phase a ref = 0
phase b ref = 0
phase c ref = 0
IF(angle < pulse width) THEN phase a ref = current demand
angle = angle + phase offset
IF(angle < pulse width) THEN phase b ref = current demand
angle = angle + phase offset
IF(angle < pulse width) THEN phase c ref = current demand

```

With this scheme the phase current references default to zero, and are made equal to the demand value if the phase angle is less than the required pulse width. Note that all angles are calculated modulo 128 or 256, for unipolar and bipolar waveforms respectively.

D.5 Current Controllers

There are two implementations of the adaptive PID current controller used for the experimental switched reluctance drive: one written in assembler for the TMS320E14, and one written in 'C' for the TMS320C31. Both perform the same basic functions which can be outlined as follows:-

```

read real phase currents
get phase current demands
IF( current demand > current limit ) THEN
    current demand = current limit
calculate current error
IF( current error > limit ) THEN
    hysteresis control:
        IF(current error > 0 ) THEN Vout = -maximum
        IF(current error < 0 ) THEN Vout = +maximum
IF( current error < limit AND output = maximum ) THEN
    proportional control:
        Vout = Kp × Ierr
IF( current error < limit AND output < maximum ) THEN
    PID control:

```

```

Vout = PID( Ierr )
convert voltage into PWM value
limit PWM value to the range 1 to (maximum PWM - 1)
set PWM registers and output sign bits
IF( phase current > current limit ) THEN disable output devices

```

The PID controller function (`PID(Ierr)`) is the discrete time implementation of the 's-domain' transfer function:-

$$D(s) = K_p \left[1 + \frac{1}{sT_i} + sT_d \right] \quad (\text{D.1})$$

Chapter 4, section 4.1.3 gives further details of this function, and the derivation of the discrete time difference equation implementation which is:-

$$y_n = y_{n-1} + K_p(a_2e_n + a_1e_{n-1} + a_0e_{n-2}) \quad (\text{D.2})$$

Terms are as follows; 'y' is the output, 'e' is the phase current error, 'K' and 'a' are gain terms. The subscripts relate to their sample time, 'n' being the current sample, and 'n-1' the previous sample etc.

The difference equation is implemented in TMS320E14 assembler as follows:-

```

LARK      0, PID_GAINS      ;LOAD POINTER TO GAIN VALUES
ZALH      Y_MSW            ;RELOAD LAST VALUE ACC=Y(N-1)
ADD5      Y_LSW
LT        ERROR2            ;T=E(N-2)
MPY       *+, 0              ;P=K2*E(N-2)
LTD       ERROR1            ;ACC=Y(N-1)+K2*E(N-2)
MPY       *+, 0              ;P=K1*E(N-1)
LTD       ERROR0            ;ACC=Y(N-1)+K2*E(N-2)+K1*E(N-1)
MPY       *+, 0              ;P=K0*E(N)
APAC      ;ACC=Y(N-1)+K2*E(N-2)+K1*E(N-1)+K0*E(N)
SACH      Y_MSW            ;STORE NEW VALUE Y(N)
SACL      Y_LSW

```

The twelve instructions required only take $2\mu\text{s}$ to complete. Extensive use is made of the LTD instruction which loads one of the multiplication registers, adds the last result to the accumulator, and shifts data values in memory.

The 'C' code version used for the TMS320C31 has the following form:-

```
output_volts := last_output;
for( n=0 ; n < 3 ; n++ )
    output_volts += current_error[n] * pid_gain[n]; /*calc. output*/
for( n=2 ; n > 0 ; n-- )
    current_error[n+1] := current_error[n]; /*shift data*/
last_output := output_volts;
```

The first loop performs the actual PID calculation, with the second loop moving the data in memory to the next sample position. When compiled the actual code has several more instructions than the TMS320E14 hand assembled version. However as floating point values are used, no pre- and post-scaling of values is required which compensates for the extra instructions in the difference equation.

D.6 Flux-linkage Controllers

The principles of flux-linkage control are described in Chapter 7, section 7.1.2, from which it is evident that flux-linkage control is more computationally intensive than PID current control, although the implementation uses basic mathematical functions.

Flux-linkage control code for a single-tooth winding switched reluctance machine drive has the following structure:-

```
get real phase currents
get rotor angle
calculate real flux-linkage from current and position
get phase flux-linkage demands
calculate flux-linkage errors
calculate voltage required to correct flux error
add IR compensation voltage
add last output voltage compensation
calculate required PWM duty cycle for desired output voltage
IF( PWM value > maximum PWM ) THEN limit PWM value to maximum
save PWM value
IF( phase current > current limit ) THEN disable output devices
```

The code required for a fully-pitched winding switched reluctance drive is essentially the same, but has additional code to transform from fully-pitched winding to single-tooth winding parameters, and visa-versa. The resulting code has the following form:-

```

get real phase currents
transform real current to equivalent single-tooth winding current
get rotor angle
calculate equivalent flux-linkage from current and position
get single-tooth winding flux-linkage demands
calculate single-tooth winding flux-linkage errors
calculate voltage required to correct flux error
transform single-tooth winding voltage back to real phase voltage
add IR compensation voltage
add last output voltage compensation
calculate required PWM duty cycle for desired output voltage
IF( PWM value > maximum PWM ) THEN limit PWM value to maximum
save PWM value
IF( phase current > current limit ) THEN disable output devices

```

D.7 Speed Measurement

The optical encoder has a 10bit Gray code output which gives 1024 position per revolution. However the two least significant bits (LSBs) only change 512 times per revolution, 256 complete cycles per revolution. The period of one of the LSBs is therefore:-

$$\tau = \frac{60}{256 \times rpm} \quad (D.3)$$

The encoder LSB is connected to one of the TMS320E14 capture timer inputs which is programmed to store the timer value on the rising edge of the input only. The timer counter is clocked at 6.0MHz, and therefore the number of counts between successive rising edges on the input are given by:-

$$Count = \tau \times 6.0 \times 10^6 \quad (D.4)$$

Which substituting the value of τ is:-

$$Count = \frac{60 \times 6.0 \times 10^6}{256 \times rpm} = \frac{1,406,250}{rpm} \quad (D.5)$$

As it is the count which is the known value, speed can be calculated using:-

$$rpm = \frac{1,406,250}{Count} \quad (D.6)$$

With a maximum counter value of 65536 the minimum speed without counter overflow is 21.4rpm. By counting the number of times the counter overflows between successive input transitions the speed range can be extended down to one revolution every 51 minutes using 32bit arithmetic.

D.8 Speed Controllers

The speed measurement low pass filter has a discrete time difference equation, as described in Chapter 8 section 8.5, of :-

$$y_n = e^{-\omega T} \cdot y_{n-1} + (1 - e^{-\omega T}) \cdot x_n \quad (D.7)$$

This is implemented in 'C' as follows:-

```
filtered_speed := Ksf * filtered_speed + (1 - Ksf) * measured_speed;
```

Where the speed filter gain term Ksf, is equal to the constant $e^{-\omega T}$

The adaptive speed controller, also described in Chapter 8 section 8.5, is a PI controller with high frequency roll off for normal operation, and a higher gain proportional only controller if the speed error is large, for a fast response. Speed control code has the following structure:-

```
get actual speed
get required speed
calculate speed error
IF( speed error > limit ) THEN
    proportional gain:
    torque demand = Kp × speed error
IF( speed error < limit ) THEN
    PI controller:
    torque demand = PI( speed error )
```

The PI controller function calculates the torque demand from current error using the following discrete time difference equation:-

$$y_n = y_{n-1} \cdot (1 + K_b) - y_{n-2} \cdot K_b + (x_n - x_{n-1} \cdot K_a) \cdot \left(\frac{K_p \cdot (1 - K_b)}{1 - K_a} \right) \quad (D.8)$$

This is implemented in 'C' as follows:-

```
torque_demand[0] = (1-Ka) / (1-Kb) * (rpm_err[0] - Ka * rpm_err[1]);
torque_demand[0] *= Kp;
torque_demand[0] += (1+Kb) * torque_demand[1];
torque_demand[0] -= Kb * torque_demand[2];
```

The indexes indicate the sample interval where 0 is the n'th sample, and 1 the (n-1)'th sample etc.

D.9 Random Number Calculation

The 'C' language random number function "random(N)", which returns a value between 0 and N, is usually implemented in two steps. First a new random number is generated by multiplying the seed by a constant and taking the modulus of the result. The new random number is then divided down to give a result in the range 0 to N. Although the TMS320C31 is a fast processor it is not optimised for division, and with the large volume of random numbers required for GA optimisation this would occupy a lot of processor time. A much less processor intensive algorithm is one based on the operation of a digital pseudo-random bit sequence generator. In hardware this is formed from a shift register with exclusive OR feedback which is clocked to produce new bits. This can be implemented in software to produce a 32bit pseudo-random number. Feedback taps at bits 28 and 31 give the maximum length sequence of $(2^{32} - 1)$ numbers. Numbers in the range 0 to $(2^n - 1)$ are produced by masking off the upper $(32 - n)$ bits.

Appendix E

An Introduction to Genetic Algorithms

E.1 Basics

Genetic Algorithms (GAs) originated from work by John Holland, first published in 1975 (Holland [E.1]). Since then there has been a steady growth in interest, and applications have been implemented covering a broad range; from job scheduling, to image recognition and neural net training (Buckles *et al* [E.2]).

Genetic Algorithms are essentially search algorithms, based on the principles of natural selection. Many terms from natural genetics, such as population, generation, chromosome, breeding and mutation are used. A GA attempts to find the optimum solution within a defined search space by evaluating a set of possible solutions. The best solutions are then combined to try and find a better solution. This process is then repeated until an acceptable solution is found.

E.2 Terms and Search Strategy

A possible solution to a particular problem can be represented by an *Organism* consisting of a single *Chromosome*, and which can be defined as a vector thus:-

$$U = \{x_1, x_2, \dots, x_N\}$$

Each element of the chromosome is called an *Allele* or *Gene*, and can be any member of the *Alphabet* which is valid for that allele. Each allele may be a single element, or another vector. The alphabet can have a limited number of members; for example, two elements for the binary alphabet, or 26 for the letters a to z; or it may have an infinite number of members, for example all complex numbers. Organisms which possess a particular allele are called a *Schemata*, for example, given organisms with an alphabet of the letters 'a' to 'z', a particular schema are all those organisms with a chromosome $\{#, b, c, #, #, #\}$, where '#' indicates "don't care". A *Population* is made up of a fixed number of organisms, and is defined at a time t as $P(t)$. Successive populations are called *Generations*, with the next population, $P(t+1)$, generated from $P(t)$ by a combination of *Breeding*, *Mutation* and *Selection*.

An overview of GA operation can be expressed as follows:-

```
Time t=0
Generate initial population P(0)
Evaluate P(0)
REPEAT
    Time t=t+1
    Generate P(t) from P(t-1)
    Evaluate P(t)
UNTIL Suitable solution is found.
```

Each of these processes is implemented as follows:-

Generation of initial population - Each allele, in each member of the population, $P(0)$, must be defined at $t=0$. Each allele can be selected at random from the alphabet, or from initial estimates of the solution.

Evaluation of the population - Each member of the population, $U(n)$, must have its *Fitness* evaluated. This is any means by which the members of a population can be compared to each other, by evaluation of some function which returns a value $F(n)$. Fitness may be evaluated, for example from simple numerical size, or from the number of alleles which match the desired solution

Generation of the next population - During this process the organisms in the present generation are combined, using a combination of selection, breeding, and mutation to from the next generation. Each of these processes is implemented as follows:-

Selection - This is the process by which members of the population are selected based on their fitness. This can be done in several ways. Proportional selection randomly selects an organism for breeding with a probability proportional to its fitness. A linear rank selection strategy sorts organisms into order by fitness, and then randomly selects organisms with a probability proportional to their position in the list. With an elitist strategy the fittest organism always survives into the next generation unchanged.

Breeding - This is the means by which a new generation is produced from the last generation by *Crossover* of alleles in the chromosomes. Two parents are combined such that the offspring contains alleles from each of the parents. The crossover point is selected at random, and there may be more than one crossover point.

Example:	Parent 1	{a,b,c,d,e,f}
	Parent 2	{u,v,w,x,y,z}
	Offspring	{a,b,c,x,y,z}

Mutation - This models the process of random mutation of alleles in nature. One or more alleles, selected randomly, are replaced by another allele randomly selected from the alphabet.

Example:	Parent	{a,b,c,d,e,f}
	Mutant	{a,b,x,d,e,f}

Stopping Conditions - This dictates at what point the breeding process is halted so that no further generations of organisms are produced. There are several popular stopping criteria as described by Buckles *et al* [E.2] and Buckles *et al* [E.3]. These are; when the desired solution is found, when all the organisms, $U(n)$, in a population have identical fitness, when the fitness of organisms varies by an amount less than some tolerance value, and after a fixed number of generations.

E.3 Theorems

The fundamental theory of GAs (the schema theorem, Holland [E.1]) is that the instances of certain fitter schema will increase exponentially with each generation. This theorem is slightly modified by the effects of mutation and the particular selection strategy implemented (Buckles *et al* [E.3], Grefenstette *et al* [E.4] and Goldberg [E.5]) but still maintains the property that fitter organisms will increase in number with each generation.

Another important property of GAs is that of "implicit parallelism", this means that GAs effectively perform an extensive search of hyperplanes within the search space, without actually testing all possible values.

The search for solutions involves the use of random operators and probability, thus the result is not deterministic and is dependant on the starting point.

By using population which comprises of a set of possible solutions information is gathered from many points in the search space concurrently. A GA is able to use this information to direct the search, and give better solutions. This property also makes optimisation with GAs less susceptible to the problems of local maxima and minima, and system noise.

Appendix F

List of Symbols

F.1 Abbreviations

SR Switched Reluctance

SRM Switched Reluctance Machine

MMF Magneto Motive Force

GA Genetic Algorithm

F.2 Matrix and Vector Symbols

i_{abc} The matrix column vector
$$\begin{pmatrix} i_a \\ i_b \\ i_c \end{pmatrix}$$

i_{123} The matrix column vector
$$\begin{pmatrix} i_1 \\ i_2 \\ i_3 \end{pmatrix}$$

Ψ_{abc} The matrix column vector
$$\begin{pmatrix} \Psi_a \\ \Psi_b \\ \Psi_c \end{pmatrix}$$

Ψ_{123} The matrix column vector
$$\begin{pmatrix} \Psi_1 \\ \Psi_2 \\ \Psi_3 \end{pmatrix}$$

\mathbf{C} The transformation matrix
$$\begin{pmatrix} +1 & -1 & -1 \\ -1 & +1 & -1 \\ -1 & -1 & +1 \end{pmatrix}$$

F.3 General Symbols

abc Subscripts used to denote the three phases 'a', 'b' and 'c' of the experimental fully-pitched winding SRM

123 Subscripts used to denote the three phases '1', '2' and '3' of the equivalent single-tooth winding SRM

m Number of electrical phases

n_r, n_s Number of rotor and stator teeth respectively

N_x Number of turns on phase 'x'

I_x Current in phase 'x'

V_x Voltage applied to phase 'x'

\mathcal{Z}_x MMF produced by phase 'x'

Φ Magnetic flux

ψ_x Magnetic flux-linkage in phase 'x'

W' Co-energy

M_{xy} Mutual inductance between phase 'x' and phase 'y'

L_x Inductance of phase 'x'

T Torque

T Sample interval in discrete time system

T_x Time constant 'x'

ω Rotor angular velocity

ω_x Angular frequency 'x'

θ Rotor position

K_x Gain Constant 'x'

P Real Power

References and Publications

References used in this Thesis.

Miller [1.1] Switched Reluctance Motors and their Control, Miller TJE, Clarendon Press, Oxford, 1993.

Bedford [1.2] US Patent No.3678352 and 3679953, Bedford BD, 1972.

Bausch *et al* [1.3] Speed and Torque Control of Thyristor-fed Reluctance Motors, Bausch H, Rieke B, ICEM Proceedings, Part 1, p128.1-128.10, 1976.

Lawrenson *et al* [1.4] Variable-Speed Switched Reluctance Motors, Lawrenson PJ, Stephenson JM, Blenkinsop PT, Corda J, Fulton NN, IEE Proceedings, Part B, Vol.127, No.4, p253-265, 1980.

Byrne *et al* [1.5] Characteristics of Saturable Stepper and Reluctance Motors, Byrne JV, Lacy JG, IEE Conference Publication 136, Small Electrical Machines, p93-69, 1976.

Hughes *et al* [1.6] Prediction of Stepping Motor Performance, Hughes A, Lawrenson PJ, Stelle ME, Stephenson JM, Proceedings of the International Conference on Stepping Motors and Systems, p67-76, 1974.

Mecrow [1.7] New winding arrangements for doubly salient reluctance machines, Mecrow BC, IEEE IAS Conference Proceedings, Vol.1, p249, 1992.

Scharf [1.8] Scharf A, Power Conversion and Intelligent Motion, Sep/Oct 1994, p218.

Drury [1.9] Drury W, Power Conversion and Intelligent Motion, Mar/Apr 1995, p56.

Finch *et al* [1.10] Variable Speed Drives using Multi-tooth per pole Switched Reluctance Motors, Finch JW, Harris MR, Musoke A, Metwally HMB, 13th Incremental Motion Control Systems Society Symposium, p293-302, 1984.

Miller *et al* [1.11] Design of a Synchronous Reluctance Motor Drive, Miller TJE, Hutton A, Cossar C, Staton DA, IEEE Transactions on Industry Applications, Vol.27, No.4, p741, 1991.

Horst [1.12] US Patent No.5122687, Horst G, 1992.

Chan [1.13] Low-Cost Electronic Controlled Variable-Speed Reluctance Motors, Chan CC, IEEE Transactions on Industrial Electronics, Vol.34, No.1, p95, 1987.

Byrne [1.14] US Patent : Controlled Saturation, Byrne JV.

Michaelides *et al* [1.15] A New Magnetic Flux Pattern to Improve the Efficiency of the Switched Reluctance Motor, Michaelides AM, Pollock C, IEEE IAS Conference Proceedings, p226, 1992.

Blake *et al* [1.16] Blake RJ, Lawrenson PJ, PCIM conference Nuremberg, 1992.

West [1.17] DC, Induction, Reluctance and PM Motors for Electric Vehicles, West JGW, Power Engineering Journal, Vol.8, No.2, p77-88, 1994.

Chang [1.18] Comparison of AC Drives for Electric Vehicles - A Report on Experts' Opinion Survey, Chang L, IEEE Aerospace and Electronic Systems Magazine, Vol.9, No.8, p7-11, 1994.

Radun [1.19] High-Power Density Switched Reluctance Motor Drive for Aerospace Applications, Radun AV, IEEE Transactions on Industry Applications, Vol.28, No.1, p113-119, 1992.

Ferreira *et al* [1.20] Detailed Design of a 30kW Switched Reluctance Starter/Generator System for a Gas Turbine Engine Application, Ferreira CA, Jones SR, Heglund WS, Jones WD, IEEE Industry Applications, Vol.31, No.3, p553, 1995.

Corda *et al* [1.21] Linear Switched Reluctance Actuator, Corda J, Skopljak E, IEE Conference Publication 376, Electrical Machines and Drives, p535, 1993.

Liang *et al* [1.22] New Variable Reluctance Motor Utilising An Auxiliary Commutation Winding, Liang F, Lipo TA, IEEE Transactions on Industry Applications, Vol.30, No.2, p423-432, 1994.

Li *et al* [1.23] Doubly Salient Doubly Excited Variable Reluctance Motor, Li SHY, Liang F, Zhao Y, Lipo TA, IEEE Transactions on Industry Applications, No.1, p99-106, 1995.

Ray *et al* [1.24] High Performance Switched Reluctance Brushless Drives, Ray WF, Lawrenson PJ, Davis RM, Stephenson JM, Fulton NN, Blake RJ, IEEE IAS Conference Record, p1769-1776, 1985.

Pollock *et al* [1.25] Power Converter Circuits for Switched Reluctance Motors with the Minimum Number of Switches, Pollock C, Williams BW, IEE Proceedings Part B, Vol.137, No.6, p373, 1990.

Krishnan *et al* [1.26] Design of Single Switch Per-Phase Converter for Switched Reluctance Motor Drives, Krishnan R, Materu PN, IEEE Transactions on Industrial Electronics, Vol.37, No.6, p469, 1990.

Le-Huy *et al* [1.27] A Novel Unipolar Converter for Switched Reluctance Motor, Le-Huy H, Viarouge P, Francoeur B, IEEE Transactions on Power Electronics, Vol.5, No.4, p469-475, 1990.

Taylor [1.28] An Experimental Study On Composite Control Of Switched Reluctance Motors, Taylor DG, IEEE Control Systems Magazine, Vol.11, No.2, p31-36, 1991.

Amor *et al* [1.29] Adaptive Non-linear Torque Control Of A Switched Reluctance Motor, Amor LB, Akhrif O, Dessaint LA, Olivier G, IEEE American Control Conference, p2831-2834, 1993.

Reay *et al* [1.30] Application Of Associative Memory Neural Networks To The Control A Switched Reluctance Motor, Reay DS, Green TC, Williams, BW, IEEE Conference Proceedings of the 19th Annual International Conference on Industrial Electronics, Control, and Instrumentation, Vol.1, p200-206, 1993.

Reay *et al* [1.31] Switched Reluctance Motor Control Via Fuzzy Adaptive Systems, Reay DS, MirkazemiMoud M, Green TC, Williams BW, IEEE Control Systems Magazine, Vol.15, No.3, p8-15, 1995.

Buja *et al* [1.32] Variable Structure Control of an SRM Drive, Buja GS, Menis R, Valla MI, IEEE Transactions on Industrial Electronics, Vol.40, No.1, p56-63, 1993.

Sugden *et al* [1.33] The Control of SR Drives: Review and Current Status, Sugden DM, Webster PD, Stephenson JM, EPE Conference Proceedings, p35, 1989.

Elmas *et al* [1.34] Position Sensorless Operation of a Switched Reluctance Drive Based on Observer, Elmas C, ZelayaDeLaParra H, IEE Conference Publication No.377, Vol.6, p82-87, 1993.

Rochford *et al* [1.35] Development Of Smooth Torque In Switched Reluctance Motors Using Self-Learning Techniques, Rochford C, Kavanagh RC, Egan MG, Murphy JMD, IEE Conference Publication No.377, Vol.6, p14-19, 1993.

Miller *et al* [1.36] New Control IC for Switched Reluctance Motor Drives, Miller TJE, Cossar C, Anderson D, IEE Conference Publication No.324, p331-335, 1990.

Wallace *et al* [1.37] Current Harmonics and Acoustic Noise in AC Adjustable-Speed Drives, Wallace AK, Spee R, Martin LG, IEEE Transactions on Industry Applications, Vol.26, No.2, p267-273, 1990.

Garcia-Otero *et al* [1.38] Reducing Acoustic Noise by Variable Speed PWM at Low Switching Frequency, Garcia-Otero S, Devaney M, ICEM Conference Proceedings, Vol.1, p210, 1992.

Kim *et al* [1.39] Resonant Link Bi-directional Power Converter: Part 2- Application To Bi-directional AC Motor Drive Without Electrolytic Capacitor, Kim JS, Sul SK, IEEE Transactions on Power Electronics, Vol.10, No.4, p485, 1995.

Cameron *et al* [1.40] The Origin and Reduction of Acoustic Noise in Doubly Salient Variable-Reluctance Motors, Cameron DE, Lang JH, Umans SD, IEEE Transactions on Industry Applications, Vol.28, No.6, p1250-1255, 1992.

Wu *et al* [1.41] Analysis and Reduction of Vibration and Acoustic Noise in The Switched Reluctance Drive, Wu CY, Pollock C, IEEE IAS Conference Record, Vol.1, p106-113, 1993.

Wu *et al* [1.42] Time Domain Analysis of Vibration and Acoustic Noise in the Switched Reluctance Drive, Wu CY, Pollock C, IEE Conference Publication No.377, p558, 1993.

Ray *et al* [1.43] A Sensorless Method For Determining Rotor Position For Switched Reluctance Motors, Ray WF, Al-Bahadly IH, Proceedings of the 5th IEE Power Electronics and Variable Speed Drives Conference, 1994.

Ehsani *et al* [1.44] Elimination of Discrete Position Sensor and Current Sensor in Switched Reluctance Motor Drives, Ehsani M, Husain I, Kulkarni AB, IEEE Transactions on Industry Applications, Vol.28, No.1, p128-135, 1992.

Husain *et al* [1.45] Rotor Position Sensing in Switched Reluctance Motor Drives by Measuring Mutually Induced Voltages, Husain I, Ehsani M, IEEE Transactions on Industry Applications, Vol.30, No.3, p665-672, 1994.

Panda *et al* [1.46] Waveform Detection Technique For Indirect Rotor-Position Sensing Of Switched-Reluctance Motor Drives Part 1: Analysis, Panda SK, Amaratunga GAJ, IEE Proceedings Part B: Electric Power Applications, Vol.140, No.1, p80-88, 1993.

Mvungi *et al* [1.47] Accurate Sensorless Rotor Position Detector in a SR Motor, Mvungi NM, Stephenson JM, EPE Conference Proceedings, Vol.1, p390, 1991.

Lumsdaine *et al* [1.48] State Observers for Variable Reluctance Motors, Lumsdaine A, Lang JH, IEEE Transactions on Industrial Electronics, Vol.37, No.2, p133, 1990.

Ertugral *et al* [1.49] A New Algorithm for Sensorless Operation of Permanent Magnet Motors, Ertugral N, Acarnley PP, IEEE Transactions on Industry Applications, Vol.30, p126, 1994.

Leonhard [1.50] Field-orientation for controlling AC-machines - principle and application, Leonhard W, IEE Conference Publication No.291, p277-282, 1988.

ABB [1.51] Direct Torque Control, Technical Guide No.1, ABB Industry Oy, 1995.

Mitsubishi [1.52] A140 - 3 Phase Inverters, Product Catalogue, Mitsubishi Electric, 1993.

Casadei [1.53] Study and Implementation of a Simplified and Efficient Digital Vector Controller for Induction Motors, Casadei D, Grandi G, Serra G, IEE Conference Publication 376, Electrical machines and Drives, p196, 1993.

Du *et al* [1.54] Shaft Speed, Load Torque and Rotor Flux Estimation of Induction Motor Drive using an Extended Luenberger Observer, Du T, Brdys MA, IEE Conference Publication 376, Electrical machines and Drives, p179, 1993.

Burnett *et al* [1.55] The Detection and Location of Faults within 3 Phase Induction Motors, Burnett R, Watson JF, Elder S, ICEM Conference Proceedings, Vol.2, p288, 1994.

Penman *et al* [1.56] Machine Diagnostics with ANN's : Possibilities and Problems, Penman J, Stavrou A, Yin CM, Jiang H, Hatzipantelis E, ICEM Conference Proceedings, Vol.2, p363, 1994.

French *et al* [1.57] Optimal Torque Control of Permanent Magnet Motors, French CD, Acarnley PP, Jack AG, ICEM Conference Proceedings, Vol.3, p720, 1994.

French *et al* [1.58] Real Time Current Estimation in Brushless DC Drives Using a Single DC Link Current Sensor, French CD, Acarnley PP, Jack AG, IEE Conference Publication No.377, Vol.4, p445-450, 1993.

French [1.59] Real Time Control of Electric Drives, French CD, PhD Thesis, University of Newcastle-upon-Tyne, UK, 1995.

Lessmeier *et al* [1.60] Microprocessor Controlled AC Servo Drives with Synchronous or Induction Motors : Which is Preferable?, Lessmeier R, Schumavher W, Leonhard W, IEEE Transactions on Industry Applications, Vol.22, No.5, 1986.

Mecrow [2.1] Fully-pitched Winding Switched Reluctance and Stepping Motor Arrangements, Mecrow BC, IEE Proceedings Part B, Vol.140, p61-70, 1993.

Preston *et al* [2.2] A Switched Reluctance Motor Model With Mutual Coupling and Multiphase Excitation, Preston MA, Lyons JP, IEEE Transactions on Magnetics, Vol.27, No.6/2, p5423-5425, 1991.

Kron [2.3] A Short Course in Tensor Analysis for Electrical Engineers, Kron G, John Wiley & Sons, Inc. 1942.

Mecrow [2.4] New winding arrangements for doubly salient reluctance machines, Mecrow BC, IEEE IAS Conference Proceedings, Vol.1, p249, 1992.

Barrass *et al* [2.5] The Unipolar Operation of Fully Pitched Winding Switched Reluctance Drives, Barrass PG, Mecrow BC, Clothier AC, ICEM Conference Proceedings, p71-76, 1994.

Miller [2.6] Converter Volt-Ampere Requirements Of The Switched Reluctance Motor Drive, Miller TJE, IEEE IAS Conference Record, p813-819, 1984.

Harris *et al* [2.7] A Review of the Integral Horsepower Switched Reluctance Drive, Harris MR, Finch JW, Mallick JA, Miller TJE, IEEE Transactions on Industry Applications, Vol.22, No.4, p716, 1986.

Fulton *et al* [2.8] Recent Developments In High-Performance Switched Reluctance Drives, Fulton NN, Lawrenson PJ, Stephenson JM, Blake RJ, Davis RM, Ray WF, IEE Conference Publication No.254, p130-133, 1985.

Ray *et al* [2.9] High Performance Switched Reluctance Brushless Drives, Ray WF, Lawrenson PJ, Davis RM, Stephenson JM, Fulton NN, Blake RJ, IEEE IAS Conference Record, p1769-1776, 1985.

Torrey *et al* [4.1] Optimal-Efficiency Excitation Of Variable-Reluctance Motor Drives, Torrey DA, Lang JH, IEE Proceedings Part B: Electric Power Applications, Vol.138, No.1, p1-14, 1991.

Ray *et al* [4.2] Industrial Switched Reluctance Drives - Concepts And Performance, Ray WF, Davis RM, Blake RJ, Fulton NN, Stephenson JM, Lawrenson PJ, IEE Conference Publication No.234, p357-360, 1984.

Sugden *et al* [4.3] The Control of SR Drives: Review and Current Status, Sugden DM, Webster PD, Stephenson JM, EPE Conference Proceedings, p35, 1989.

Sugden *et al* [5.1] The Control of SR Drives: Review and Current Status, Sugden DM, Webster PD, Stephenson JM, EPE Conference Proceedings, p35, 1989.

Mecrow [5.2] New winding arrangements for doubly salient reluctance machines, Mecrow BC, IEEE IAS Conference Proceedings, Vol.1, p249, 1992.

Lovatt [7.1] Flux Controller That Improves Drive System Performance By Accounting For Magnetic Circuit Saturation, Lovatt HC, IEE Conference Publication No.377, Vol.4, p163-167, 1993.

Matsui *et al* [8.1] High-Precision Torque Control of Reluctance Motors, Matsui N, Akao N, Wakino T, IEEE Transactions on Industry Applications, Vol.27, No.5, p902-907, 1991.

Amor *et al* [8.2] Adaptive Non-linear Torque Control of a Switched Reluctance Motor, Amor LB, Akhrif O, Dessaint LA, Olivier G, American Control Conference, Part 3, p2831-2834, 1993.

Jiang *et al* [8.3] Control Strategy to Minimise Torque Ripples of Switched Reluctance Machine, Jiang Q, Ye J, Zhou E, ICEM Conference Proceedings, Vol.2, p475, 1992.

Kavanagh *et al* [8.4] Torque Ripple Minimisation In Switched Reluctance Drives Using Self-Learning Techniques, Kavanagh RC, Murphy JMD, Egan MG, IECON Conference Proceedings, Vol.1 p289-294, 1991.

Rochford *et al* [8.5] Development of Smooth Torque in Switched Reluctance Motors Using Self-Learning Techniques, Rochford C, Kavanagh RC, Egan MG, Murphy JMD, IEE Conference Publication No.377, Vol.6, p14-19, 1993.

Schramm *et al* [8.6] Optimum Commutation - Current Profile on Torque Linearisation of Switched Reluctance Motors, Schramm DS, Williams BW, Green TC, ICEM Conference Proceedings, Vol.2, p484, 1992.

Taylor [8.7] An Experimental Study on Composite Control of Switched Reluctance Motors, Taylor DG, IEEE Control Systems Magazine, Vol.11, No.2 p31-36, 1991.

Ilic'-Spong *et al* [8.8] Instantaneous Torque Control of Electric Motor Drives, Ilic'-Spong M, Miller TJE, MacMinn SE, Thorp JS, IEEE PESC Conference Record, p42-48, 1985.

Wallace *et al* [8.9] A Balanced Commutator For Switched Reluctance Motors To Reduce Torque Ripple, Wallace RS, Taylor DG, IEEE Transactions on Power Electronics, Vol.7, No.4, p617-626, 1992.

Wallace *et al* [8.10] Low-Torque-Ripple Switched Reluctance Motors For Direct-Drive Robotics, Wallace RS, Taylor DG, IEEE Transactions on Robotics and Automation, Vol.7, No.6, p733-742, 1991.

Steiert *et al* [8.11] Torque Control Of The Doubly-Salient Reluctance Motor, Steiert U, Spath H, European Transactions on Electrical Power Engineering/ETEP, Vol.3, No.4, p265-272, 1993.

Reay *et al* [8.12] Minimisation Of Torque Ripple In A Switched Reluctance Motor Using A Neural Network, Reay DS, Green TC, Williams BW, IEE Conference Publication, No.372, 3rd International Conference on Artificial Neural Networks, p224-226, 1993.

Reay *et al* [8.13] Application Of Associative Memory Neural Networks To The Control A Switched Reluctance Motor, Reay DS, Green TC, Williams BW, IEEE IECON Conference Proceedings, p200-206, 1993.

Cameron *et al* [8.14] The Origin And Reduction Of Acoustic Noise In Doubly Salient Variable-Reluctance Motors, Cameron DE, Lang JH, Umans SD, IEEE Transactions on Industry Applications, Vol.28, No.6, p1250-1255, 1992.

Wu *et al* [8.15] Analysis and Reduction of Vibration and Acoustic Noise in the Switched Reluctance Drive, Wu CY, Pollock C, IEEE IAS Conference Record, Vol.1, p106-113, 1993.

Holland [8.16] Adaptation in Natural and Artificial Systems, Holland JH, University of Michigan Press, 1975.

Rawlins [8.17] Foundations of genetic algorithms, Rawlins GJE, M. Kaufmann Publishers, 1991.

Whitley [8.18] Foundations of genetic algorithms 2, Whitley LD, M. Kaufmann Publishers, 1993.

Buckles *et al* [8.19] Genetic Algorithms, Buckles BP, Petry FE, IEEE Computer Society Press, 1992.

Grefenstette [8.20] Optimisation of Control parameters for Genetic Algorithms, Grefenstette JJ, IEEE Transactions on Systems, Man and Cybernetics, p1222-128, 1986.

Goldberg [8.21] Sizing Populations for Serial and Parallel Genetic Algorithms, Goldberg DE, Proceedings of the 3rd International Conference on Genetic Algorithms, p70-79, 1989.

Grefenstette *et al* [8.22] How Genetic Algorithms Work : A Critical Look at Implicit Parallelism, Proceedings of the 3rd International Conference on Genetic Algorithms, p20-27, 1989.

Buckles *et al* [8.23] Schema Survival Rates and Heuristic Search in Genetic Algorithms, Buckles BP, Petry FE, Kuester RL, Proceedings : Tools for Artificial Intelligence, p322-327, 1990.

Cameron *et al* [9.1] Origin Of Acoustic Noise In Variable-Reluctance Motors, Cameron DE, Lang JH, Umans SD, IEEE IAS Conference Record, p108-115, 1989.

Wu *et al* [9.2] Time Domain Analysis of Vibration and Acoustic Noise in the Switched Reluctance Motor, Wu CY, Pollock C, IEE Conference Publication No.377, p558, 1993.

Holland [E.1] Adaptation in Natural and Artificial Systems, Holland JH, University of Michigan Press, 1975.

Buckles *et al* [E.2] Genetic algorithms, Buckles BP, Petry F, IEEE Computer Society Press, 1992.

Buckles *et al* [E.3] Schema Survival Rates and Heuristic Search in Genetic Algorithms, Buckles BP, Petry FE, Kuester RL, Proceedings : Tools for Artificial Intelligence, p322-327, 1990.

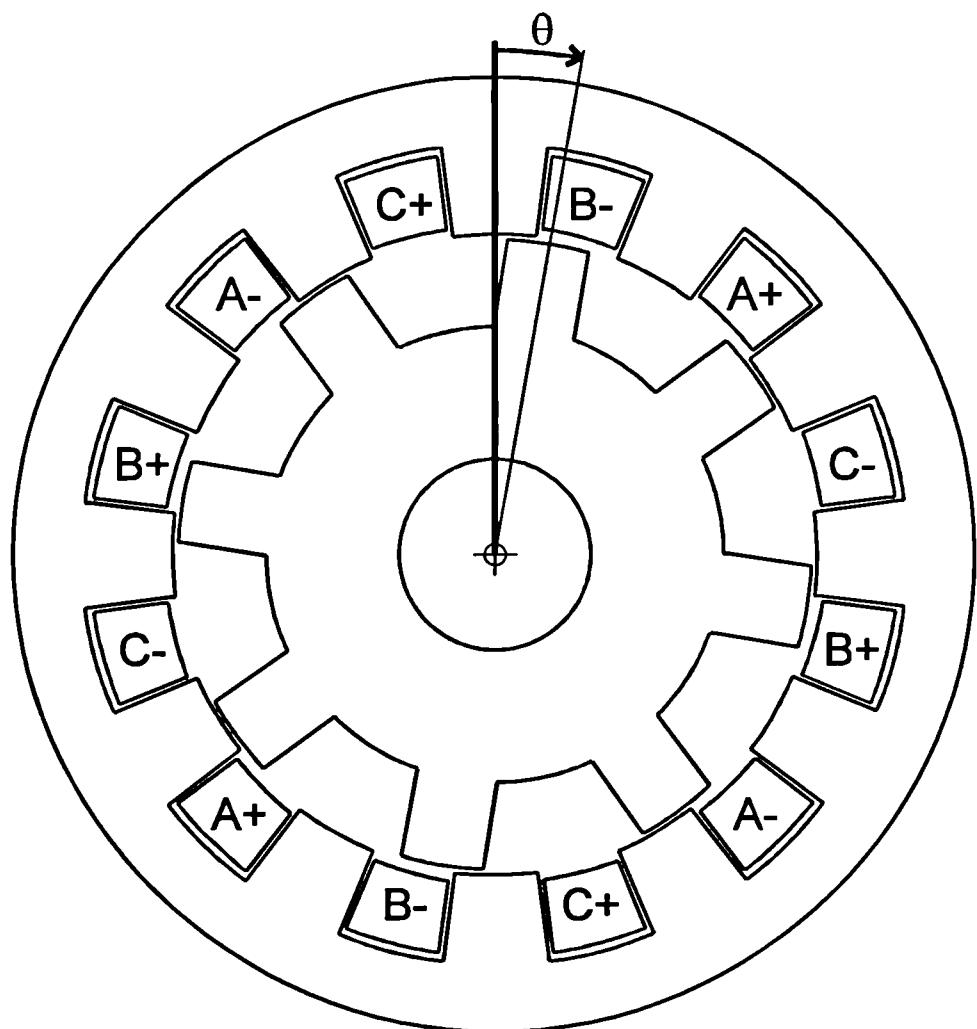
Grefenstette *et al* [E.4] How Genetic Algorithms Work : A Critical Look at Implicit Parallelism, Proceedings of the 3rd International Conference on Genetic Algorithms, p20-27, 1989.

Goldberg [E.5] Sizing Populations for Serial and Parallel Genetic Algorithms, Goldberg DE, Proceedings of the 3rd International Conference on Genetic Algorithms, p70-79, 1989.

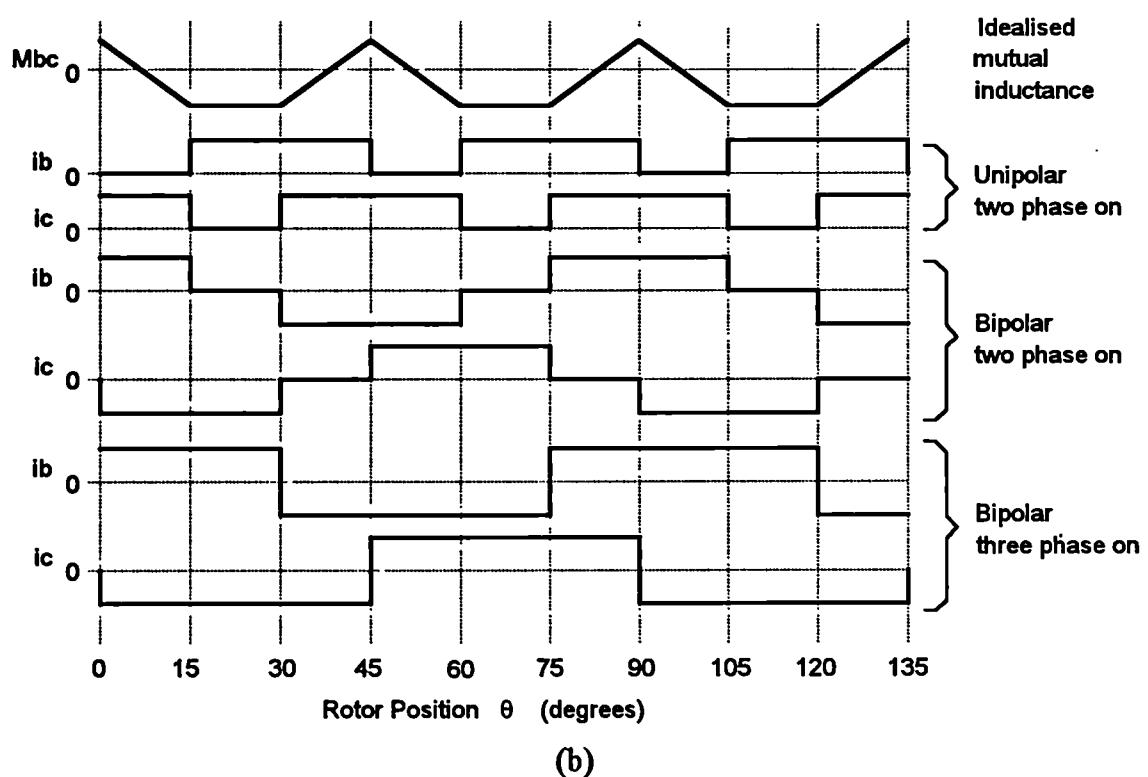
Publications by the Author

The Unipolar Operation of Fully-Pitched Winding Switched Reluctance Drives, Barrass PG, Mecrow BC, Clothier AC, ICEM Conference Proceedings, Vol.1, p71, 1994.

Bipolar Operation of Fully-Pitched Winding Switched Reluctance Drives, Barrass PG, Mecrow BC, Clothier AC, IEE EMD Conference Proceedings, p252, 1995.



(a)



(b)

Figure 5.4.1 (a) - Cross section of a three phase 12:8 machine showing winding positions and rotor position angle θ measured from the mechanically aligned position. (b) - Idealised mutual inductance and current waveforms for unipolar and bipolar operation.