Real-time Model-based Loss Minimisation Control for Electric Vehicle Drives

Dave Winterborne

Department of Electrical and Electronic Engineering

Newcastle University

A thesis submitted for the degree of

Doctor of Philosophy

September 2014

Abstract

Environmental concern and the opportunity for commercial gain are two factors driving the expansion of the electric vehicle (EV) market. Due to the limitations of current battery technology, the efficiency of the traction drive, which includes the electric motor and power electronic converter, is of prime importance. Whilst electric machines utilising permanent magnets (PMs) are popular due to their high energy density, industry concerns about the security of supply have led to interest in magnet-free solutions. Induction machines (IMs) offer such an option.

Control of IMs is a mature but complex field. Many techniques for optimising the efficiency of the drive system have been proposed. The vast majority of these methods involve an analytical study of the system to reveal relationships between the controlled variable and efficiency, allowing the latter to be optimised. This inevitably involves simplifications of the problem to arrive at a practically-implementable control scheme. What has not been investigated is real-time calculation of the system losses in order to optimise the efficiency, and the work presented in this thesis attempts to achieve this.

The conventional control scheme is examined and a new structure implemented where a model of the system loss is able to directly influence the switching action of the inverter, thus reducing loss. The need to maintain performance alongside loss minimisation is recognised and a cost function-based solution proposed. The validation of this structure is performed both in simulation and on a practical test platform.

A model of the principle losses in the drive system is derived, taking into account the processing power typically available for this application, and implemented in the structure outlined. The effect of the new control scheme on efficiency is investigated and results show gains of up to 3%-points are achievable under certain conditions.

To no-one in particular.

Acknowledgements

First and foremost I'd like to thank Volker for putting his faith in my ability to take on this project, from start to finish, and for giving me the space to develop my own ideas, and help when I needed it. I'm also grateful to the rest of the academic and research staff in what is now the Electrical Power group, particularly Barrie as my second supervisor and Dave for lending a patient ear and sharing his huge depth of knowledge in matters both theoretical and practical.

This work was supported financially by EPSRC, and without their generous funding, would not have been possible.

I received excellent technical support from every angle. Thanks go to James Richardson and Steve Mitchell for their valuable help in the lab; Darren in the electronics workshop; Jack and all the lads in the mechanical workshop; and Ian, Dave and Ian for dealing with IT issues of every variety.

Thanks also to the office staff, particularly Gill, for keeping everything running smoothly in the background.

Life as a PhD student is not all hard work, and I have met many people in the last four years who have made it a pleasure. There are too many of you to name and I'm loath to miss anyone out: you know who you are. Many of you are now valued colleagues as well as friends.

I must, however, specifically mention a few people: Simon for general advice, guidance and friendship; Chris for sharing his wealth of practical experience; Dave for sharing my perverse interest in the minutiae of typesetting, and other matters of similar importance; and Andy for providing words of wisdom during the early days, which are with me still.

Finally, and most importantly, I am indebted to two very special people for their love and support. My ability to see this process through is no doubt in large part due to qualities instilled in me by my mum. Whilst she may not understand the content of this work, Dani has provided my most compelling reason to do it, as she does with everything I do.

Contents

List of Figures

List of Tables

Nomenclature

Roman Symbols

A	Cost function weighting factor
C	Capacitance
D	Deviation in discretised modulation scheme
E	Energy
f	Frequency
g	Cost function result
i	Electric current
i_c	IGBT collector current
I_f	Diode forward current
i_r	Rotor current
J	Inertia
L	Inductance
L'_{lr}	Rotor leakage inductance referred to stator
L_{ls}	Stator leakage inductance
L_m	Magnetising inductance
L'_m	Magnetising inductance referred to the stator
m	Modulation index
p	Pole pairs

P_{3ph}	3-phase power at inverter output/motor input
$P_{\text{cu-r}}$	Rotor copper loss
$P_{\text{cu-s}}$	Stator copper loss
P_{dc}	Input power at DC link
P_{iron}	Iron loss
P_{mech}	Mechanical power at the machine shaft
R	Resistance
R_{ce}	IGBT on-state collector emitter resistance
R_f	Diode on-state forward resistance
R'_f	Iron loss resistance referred to the stator
r	Resolution
R'_r	Rotor resistance referred to the stator
R_s	Stator resistance
S	Standard deviation
T_0	Time period for which zero vector(s) applied
T_0	Time period of zero vector(s)
T_1	Time period for which non-zero vector(s) applied
T_a	Time period of fundamental vector A
T_b	Time period of fundamental vector B
T_e	Electromagnetic torque
T_L	Load torque
T_{pwm}	PWM switching period
T_r	Rotor time constant
$V*$	Demand voltage vector

$V_0 ... V_7$ Fundamental voltage vectors for a six-switch inverter

V_a Component of voltage vector co-incident with vector A

V_b Component of voltage vector co-incident with vector B

v_{ce} IGBT collctor emitter voltage drop

V_{ce0} IGBT collector-emitter threshold voltage

V_{dc} DC link voltage

V_f Diode forward voltage drop

V_{f0} Diode forward threshold voltage

V_{min} Voltage vector which minimises the cost function

Greek Symbols

α Angle of demand vector within current sector

ΔP_{Q-D} Loss difference between IGBT and diode

ε_m Vector amplitude error

ε_θ Vector angle error

ε Relative uncertainty

η_{drive} Drive efficiency

η_{inv} Inverter efficiency

η_{mot} Motor efficiency

λ Flux linkage

ω Angular velocity

θ Angle

θ_{app} Angle of applied vector

θ_e Electric field angle

θ_s Slip angle

P_{loss} Total power loss

Superscripts

$*$ Demand

$'$ Quantity referred to machine stator

^{e} e-frame

$^{e'}$ e-frame

^{s} s-frame

Subscripts

$_{a,b}$ Adjacent vectors

$_{A,B,C}$ Phase A, B, C

$_{\text{cond}}$ Conduction loss

$_{D}$ Diode loss

$_{d}$ Direct axis

$_{e}$ Electric field

$_{\text{equiv}}$ Equivalent vector component in discretised modulation scheme

$_{Q}$ IGBT loss

$_{q}$ Quadrature axis

$_{r}$ Rotor

$_{\text{sw}}$ Switching loss

Other Symbols

e'-frame Excitation reference frame with phase-shift

e-frame Excitation reference frame

s-frame Stator reference frame

Acronyms

ADC Analogue-to-Digital Converter

BEV Battery electric vehicle

DSP Digital Signal Processor

EV Electric vehicle

FEA Finite element analysis

FFT Fast Fourier Transform

FIR Finite Impulse Response

FOC Field Oriented Control

HEV Hybrid electric vehicle

ICE Internal Combustion Engine

IGBT Insulated Gate Bipolar Transistor

IM Induction machine

LMC Loss minimisation control

DPWM Discontinuous Pulse Width Modulation

MI Modulation Index

MPC Model predictive control

PI Proportional-integral (controller)

PM Permanent magnet

PMSM Permanent magnet synchronous machines

PWM Pulse width modulation

SVM Space Vector Modulation

THD Total Harmonic Distortion

Chapter 1

Introduction

Climate scientists are now in almost unanimous agreement that the temperature of the Earth's atmosphere is increasing, and furthermore at an increasing rate. This has already caused far-reaching changes to the global climate, affecting patterns of weather on both large and small scales. The effects can be seen on ecosystems at every level, including human beings, and are often destructive. What's more, increased scientific research has revealed subtler effects that anticipate greater problems in the future [1].

More controversial, but generally supported by the available evidence, is the hypothesis that global warming is being caused, or at least exacerbated, by human activity. Certainly the rise in temperature correlates with the explosion in the use of fossil fuels in the post-industrial era since the mid eighteenth century, and the associated release of previously trapped greenhouse gases, principally carbon dioxide, into the atmosphere. Since gaining widespread popularity, this idea has driven the development of technologies that attempt to reduce man's impact on the environment in many aspects of life.

One such major area is transportation, of both people and goods, which accounted for 62% of world oil consumption in 2011 [2]. Thanks to the convenience they offer, this is dominated by use of road vehicle, despite their relative inefficiency compared with other methods [3]. The consumption of finite resources of petroleum-derived fuels by these vehicles represents a significant proportion of the world's energy use and emission of greenhouse gases. The source of these gases is the internal combustion engine (ICE) which provides traction in almost all road vehicles, and this has led to alternative methods of propulsion being pursued. Foremost among them are vehicles propelled by electric motors, known as electric vehicles (EVs).

Practical limits on the efficiency of ICEs mean that large gains in the overall efficiency of the whole system, from energy production to vehicle propulsion—its

so-called "well-to-wheels" efficiency—can be realised. The greatly larger efficiency of electric motors compared with ICEs—on the order of 70 to 95% under most operating conditions compared with typically 25% for a petrol engine and 30% for diesel engine—is tempered by the losses incurred in generating, transmitting, distributing, storing and converting electrical energy along its journey to the motor. Similarly, the lack of direct greenhouse gas emissions from the electric drivetrain is clouded by the emissions from the electricity generation process, which depend on the fuel source. Currently, this is largely fossil fuels, although the share of renewable sources is growing [4].

Despite these drawbacks, electric traction motors represent a crucial shift in automotive design both technologically and paradigmatically. Conventional ICE powertrains are constrained to the model of burning a fuel to directly convert the energy to mechanical work. By contrast, electric motors are usually supplied with energy stored on-board in a battery. The stored energy may be generated from a variety of sources both renewable and non-renewable, which may occur some distance from the point of use. Alternatively, the electricity may be generated on-board from a fuel source using an ICE, gas turbine or hydrogen fuel cell, giving a variety of options. It is this flexibility that gives electric motors a considerable advantage in the development of new solutions to the overall issues of vehicle propulsion.

Thanks to these advantages, EVs have been intensively researched over the last two to three decades and many mass-manufactured passenger cars with electric powertrains are now available.The flexibility of a motor-based solution discussed above has led to a number of different implementations. The main distinction is between battery electric vehicles (BEVs) which are driven exclusively by the electrical system and hybrid electric vehicles (HEVs) which combine power from a motor and ICE.

A number of different HEV topologies exist, with variations on the mechanical coupling of power sources, and the flow of power between them. HEVs usually use the electric machine as a generator to recover power during braking in a process known as regeneration, but the ICE may also be used as a prime mover in the generation of electrical energy all or some of the time. Even HEVs where the electrical energy is generated on-board may provide the facility to charge the battery from the mains supply while stationary; it is then referred to as a plug-in hybrid (PHEV).

The two basic HEV topologies are series and parallel. In a series HEV, the ICE is used with an electric machine to generate electrical energy, which is then used by an electrical machine to drive the wheels. There is no direct connection between the ICE and the load. By contrast, in a parallel HEV, both the ICE and an electrical machine are able to provide power to the wheels. The main components of these two types and the power flows between them are illustrated in Figure 1.1. There are many variations on these topologies, using different combinations of electric

machines (often the same machine is used for both motoring and generating) and ICEs via different transmission elements and clutches. HEV topologies and power management within them is a broad and expanding topic and will not be covered in more depth here.

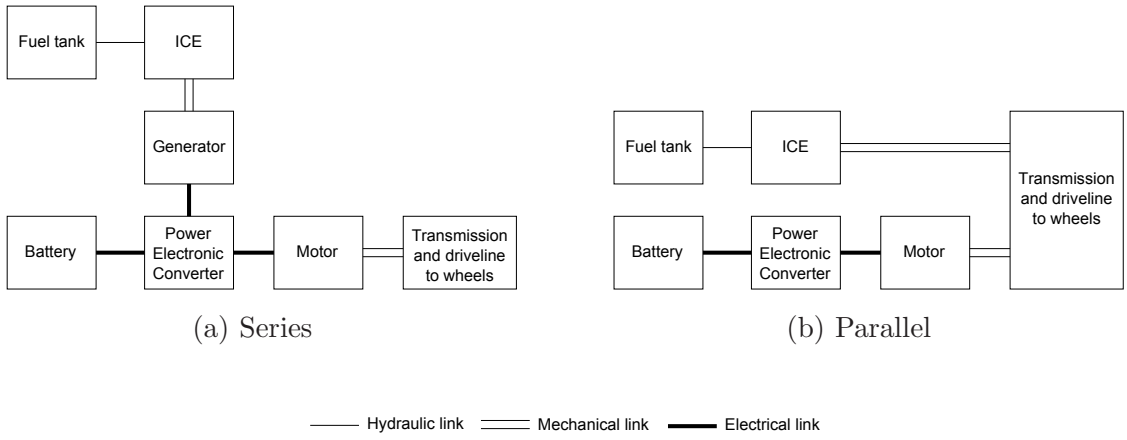

(a) Series (b) Parallel

Figure 1.1: Basic HEV topologies

The proportions of electric and conventional power use in an HEV can be varied to achieve different aims. These range from micro hybrids which use the motor only to increase vehicle efficiency by starting and stopping the ICE automatically, through mild hybrids which add electric assistance and recover braking energy, to strong hybrids which can operate under solely electrical power [5]. The strongest hybrids are known as range extended EVs, as the ICE is only used occasionally when it is necessary to charge the battery. The principle categories of EV are shown in Figure 1.2.

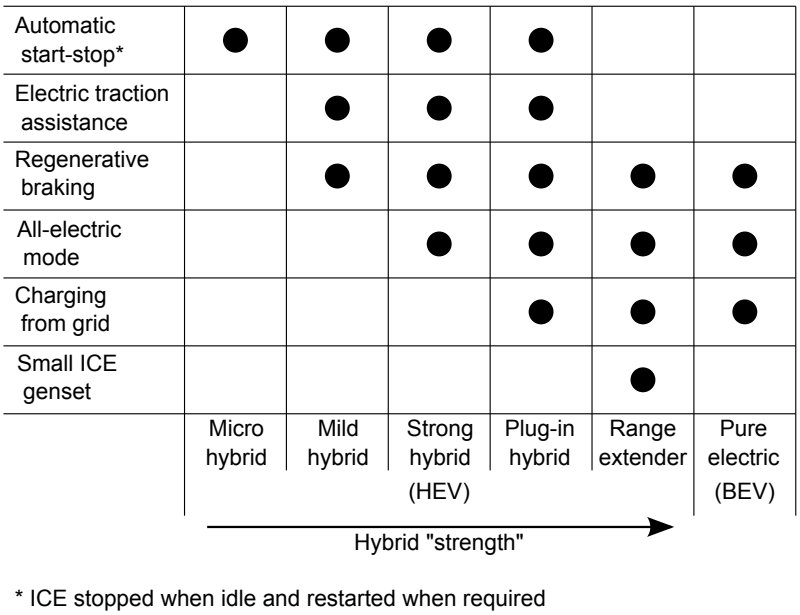

Figure 1.2: Classfication of EV types by degree of hybridisation

Rising fuel prices and increased environmental concern have led to EVs making up a small but growing part of the automotive market. HEVs made up over 3% of new car sales in the USA in 2012; more than double that in 2012 [5]. As the number of vehicles in use increases, the total amount of energy and greenhouse gas emissions saved by small gains in efficiency becomes significant, and much research effort is focussed on this goal. The move towards stronger hybrids and pure BEVs is being slowed by what is termed "range anxiety": the consumer's reluctance to adopt all-electric technology for fear of not being able to reach their destination without charging their vehicle, which is often time-consuming or unavailable, and potentially becoming stranded. This provides further incentive to increase the efficiency of electric drivetrains to overcome this, which in turn will increase the EV market share and bolster the overall energy consumption and emission reductions described above.

Early examples of EVs featured DC machines, largely due to their simplicity of control. However, their low efficiency, low power density and need for slip rings and brushes to commutate the machine make this an unattractive proposition for a mass-market EV [6].

Most commercially-available EVs, such as the Nissan LEAF, Toyota Prius and Honda Civic, Citroen DS5 and Peugeot 508 hybrids, today use permanent magnet synchronous machines (PMSMs) [5, 7]. These offer power density unrivalled by any other motor type and are very efficient. The excitation of the rotor is provided by permanent magnets rather than induced currents and therefore there are no resistive losses caused by currents in the rotor.

Induction machines (IMs) do offer some advantages over PMSMs however. Their simple construction and lack of permanent magnets makes them both robust and inexpensive to build. The lack of permanent magnets also enables field weakening to be implemented, extending the speed range of the motor. Additionally, the safety of the vehicle is improved under certain conditions since no electric field can be generated without external excitation. Due to their lower power and torque density, and efficiency, IMs are not as prevalent as PMSMs in EV applications. The popular Tesla Model S and Roadster, which feature copper rotor IMs [8], are notable examples. They have also featured recently in vehicles from Toyota, Honda, Ford and Fiat [9].

Traditionally seen as a rather primitive design, and dismissed due to high acoustic noise and torque ripple, the switched reluctance machine (SRM) has gained interest in recent years due to its fundamentally very simple, and hence low-cost, and robust construction. Research effort has addressed the aforementioned problems and improved power density, making it competitive.

The potential advantages of IMs as an alternative to PMSMs as outlined above are part of the reason this work is focussed on improving the efficiency of induction machine drives: one of the key drawbacks of IMs.

An important factor in the comparison of IM and PMSM motors for EVs is not technical but political. There are concerns over availability of the rare earth elements needed to manufacture permanent magnets due to finite deposits of the containing minerals and the location of these. The vast majority of rare earth material is currently mined in China and western carmakers are reluctant to be reliant on this source due to the historical tension between China and many western nations, particularly the United States [10]. For this reason, the automotive industry is keen to explore magnet-free electric drivetrain solutions, and the focus on IMs in this work is mainly a response to this.

Another simple but compelling reason to pursue IM drives as a solution for future EVs is that of cost. The mass-market automotive industry, like any other volume-manufacturing industry, is highly driven by cost. Thanks mainly to their lack of permanent magnet material, but also to simple and well-practised construction methods, IMs represent a considerably lower-cost solution than PMSMs.

The stator of the IM consists of a stack of steel laminations containing slots which hold the copper conductors of the phase windings. The rotor is usually of the type termed "squirrel-cage": a number of copper or aluminium bars embedded in laminated steel core and joined at both ends by conducting rings. When the stator is fed with a set of rotating electric currents, the resultant magnetic field induces currents in the rotor since it is short-circuited. The interaction between these two electric fields produces torque.

In an industrial setting, an IM may be operated directly from the 3-phase supply, since this consists of the set of rotating currents needed to produce a rotating magnetic field. This results in fixed-speed operation, dictated by the frequency of the supply and the design of the motor. In many cases, variable speed operation is desirable, and traction is one such example. Although direct AC-AC conversion is possible, the reduction in the cost of silicon power devices has made indirect conversion, in which the supply is rectified to DC before being inverted back to AC at the required frequency, almost universal. The latter process can also be used in isolation where a DC source, such as a battery in an EV is available.

IMs present a more complex control problem than DC and PM machines due to the difference in frequency between the stator and rotor electric fields, called slip, which determines the torque produced. The low cost of processing power today means that the computation necessary to produce high quality control of induction motors is not a limiting factor in EV applications. Many approaches to this control problem exist and will be studied in the course of this work. The process of modulating a DC

source to impose a given voltage waveform on the stator windings is an interrelated issue and will also be investigated thoroughly.

Particular attention will be paid to those control schemes which attempt to reduce losses in the drive system and thereby increase efficiency. These are often based on detailed mathematical analysis of the system in order to establish relationships between the controlled variables and the efficiency and determine rules or patterns for the controller. This work aims to take a fundamentally different approach by developing a model of the losses which can be used to assess the impact of a particular control action on the efficiency. In doing so, the work hopes to avoid the pitfalls of trying to precisely define the operation of parts of the system which may require simplification. In essence, the idea is to take a high-level view of the whole system, looking only at the variables that are directly controllable, and the quantities relevant to the stated objectives: in this case the losses. This thesis describes the investigation of the hypothesis that this will increase the system efficiency, from background study, through the development and testing of novel methods to the analysis of the results of experimental work.

The objectives of this project were to:

- Examine and compare published approaches to the modelling of various sources of loss in IM drives, and thus
- Determine a model of suitable detail for implementation in a practical real-time system, balancing the needs for accuracy and computational speed;
- Study the structures of the most common existing techniques for control of IM drives, and subsequently
- Identify a method by which these structures can be modified to allow the incorporation of the loss model;
- Verify the functionality and performance of the modified structure through simulation;
- Establish a test platform suitable for practical testing of the conventional and modified control structures in an environment representative of an EV;
- Perform a comparison of the efficiency and performance of the modified control structure incorporating the loss model with the conventional methods.

The contributions of this work which are believed to be novel and not previously published are:

- A vector control strategy based on explicit discretisation at the voltage control level to allow another control objective to be integrated, and based on this,

- A control scheme which incorporates a loss model and hence improves efficiency whilst maintaining performance.

Chapter 2 contains an exploration of topics relevant to the project. A full description of currently used control techniques, to which the outcome of the research will be compared, is given. This is followed by a review of more specific methods for minimising losses in induction motor drives to gain an insight into methods already investigated, and to demonstrate the novelty of this work. A broad study of predictive control shows the origin of some concepts that will feature later in the new control schemes. Existing techniques for loss modelling are examined to determine the available levels of accuracy and complexity, and the compromise between these factors.

Having identified the conventional control methods in Chapter 2 and thoroughly studied their structure, changes to these which will enable the aims of the work to be achieved are proposed and developed in Chapter 3. A consideration of the proposed approach with respect to the current techniques is presented, resulting in a rationale for the development of a control structure which allows multiple objectives to be met in the manner discussed previously in this chapter. Three different schemes are proposed, though some features are shared between them. Simulation results for each are presented, analysed and compared, giving reasons for the choice of one method for further investigation.

Chapter 4 discusses the need for practical testing of the proposed control scheme and hence the setup of a platform for this testing. The conventional control scheme presented in Chapter 2 is implemented on this platform and the efficiency of the inverter, motor and complete drive measured to demonstrate the test platform and provide a point of comparison for the new control scheme. The new control scheme developed in Chapter 3 is tested without the loss minimisation aspect of its operation. This proves the correct operation of the control structure, laying the groundwork for the implementation of the control structure in the next chapter.

Chapter 5 introduces the loss model in its proposed implementation for the new control scheme. It is evaluated using data from practical tests in Chapter 4 in order to validate it. Its ability to reduce losses when combined with the modulation scheme developed in Chapter 3 is estimated with analytical calculations, informing the final choice of loss components to be included. Practical results showing the efficiency of the new control scheme incorporating the loss model, in comparison with conventional methods, are presented. Several tests are performed to assess the impact of different parameters of the new control scheme. Finally, an analysis of the computational requirements of the new scheme is given to show that implementation in a real-world application is feasible.

The final chapter brings together the work presented in the previous chapters and conclusions drawn from the whole project. Ideas for the further development of the themes seen in this thesis are discussed.

Chapter 2

State of the Art

This chapter presents background research on a number of topics relating to the project. The state of the art of several areas which are considered and developed by this work are investigated by studying relevant literature. The aim is to give context to the novel concepts presented later by showing how these relate to existing work.

Classical control techniques are described as these provide a benchmark against which the proposed control scheme, which is the ultimate product of this work, is evaluated.

Structured approaches to minimisation of drive losses through control are presented, to show the ways this problem has been addressed previously and to demonstrate the novelty of the new scheme.

A summary of the field of model predictive control is given, as key aspects are shared with the proposed scheme.

Finally, a breakdown of the power losses in the drive system is shown and various ways of modelling them investigated.

2.1 Conventional IM Drive Control Techniques

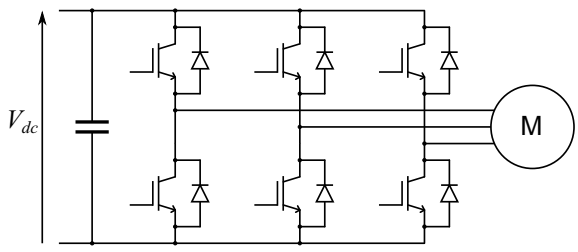

Figure 2.1: Topology of a conventional three-phase two-level inverter

A conventional three-phase voltage source inverter, as is most commonly used in EV applications, consists of six power switching devices with anti-parallel diodes, in the configuration shown in Figure 2.1. The power devices will henceforth be assumed to be insulated gate bipolar transistors (IGBTs) since this is the most common choice for the range of power usually encountered in this application. These six IGBTs, controlled by the potential difference across the gate and source terminals, ultimately offer the sole means of controlling the motor.

The IGBTs switch the DC link voltage across the stator windings of the motor, driving current into the windings. The IGBTs are operated digitally—i.e. they are either fully on or fully off—as their losses are lower in these conditions [11]. Since the windings are inductive, low frequency current waveforms can be produced by the high frequency switching of the IGBTs. The effective voltage, and hence current, is controlled by the duty cycle of the switching signal. This is generally known as pulse width modulation (PWM). There are a variety of different modulation schemes, but the purpose is to achieve voltage control.

In the open loop, this is the minimum level of control which is necessary to operate the motor. Greater flexibility can be achieved by closing the current control loop. This is most simply achieved by hysteresis control, or more commonly, with proportional-integral (PI) controllers.

Controlling the current effectively allows control of the torque and flux. Torque control is analogous to throttle control in an ICE and so is usually used in EV motor control, but speed control is desirable in some applications. This can be achieved with an outer loop controlling the torque to produce a demanded speed.

This hierarchy and the general structure of the cascaded controller are shown in Figure 2.2.

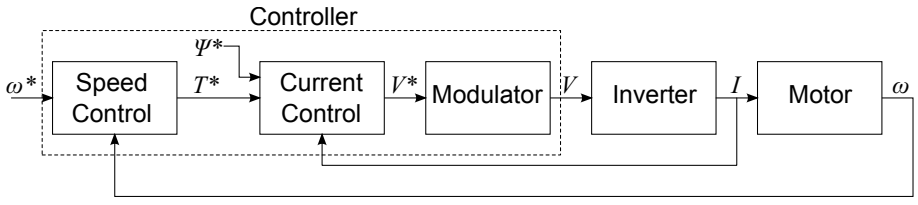

Figure 2.2: Basic elements of generic cascaded motor controller

2.1.1 Voltage Modulation Strategies

2.1.1.1 Pulse Width Modulation

The most basic form of modulation compares a fixed high frequency triangular carrier waveform, usually in the order of 1–20 kHz [12], with a modulating signal

representing the desired current in the phase, which is usually sinusoidal. The carrier and modulating signals are compared and this determines whether the top or bottom IGBT in the leg is switched on. The top and bottom IGBTs are switched in a complementary manner; switching both on would short-circuit the DC link, a condition known as shoot-through. To avoid this, dead time is inserted between the switch-off of one IGBT and the switch-on of the other. Each leg is controlled separately, and has its own modulating signal. PWM is used in a wide variety of power electronics applications.

The Modulation Index (MI), the ratio of modulating signal to carrier signal amplitude, is an important concept as it controls the effective amplitude of the voltage, and ultimately the current.

The modulating signal is most commonly a sinusoid with the required fundamental electrical frequency according to desired speed.

However, various other non-sinusoidal waveforms have been proposed. A distinction may be made between continuous and discontinuous modulating waveforms. Sinusoidal PWM is considered a continuous modulation method, since each leg is switched in every PWM period. In discontinuous methods one leg is not switched during a certain part of the waveform, with a period greater than the PWM period, with the result that the leg in question is clamped to one of the DC link rails. Such a modulating signal may be produced by adding a zero-sequence signal to the sinusoidal waveform. The zero sequence signal consists of the difference between the modulating signal and one of the DC link rails.

An example is shown in Figure 2.3 of a generalised discontinuous PWM (DPWM) scheme [13]. The zero sequence signal V_0 is added to the fundamental modulating signals to produce new voltage demands V_a^* (highlighted in blue), V_b^*, V_c^*. The angle ψ determines the phase relationship of the zero sequence signal and fundamental signals. This angle may be 0 as in DPWM0 [14], $\frac{\pi}{6}$ as in DPWM1 [15] or $\frac{\pi}{3}$ as in DPWM2 [14]. The optimum angle is dependent on the load angle [16] and may also be optimised online [13].

Eliminating switching in one leg for a significant proportion of the fundamental period increases the efficiency of the inverter since there is no switching loss in that leg during that period [16]. Analysis in [13] shows switching loss reduction of up to 50%.

Consideration of the output waveform quality is also presented in [13] where it is quantified by means of a harmonic distortion factor (HDF). The analysis shows that whilst continuous modulation techniques exhibit an HDF which increases monotonically with modulation index, DPWM methods show a peak HDF below a MI of 1. For low MIs, DWPM shows more harmonic distortion than continuous modulation,

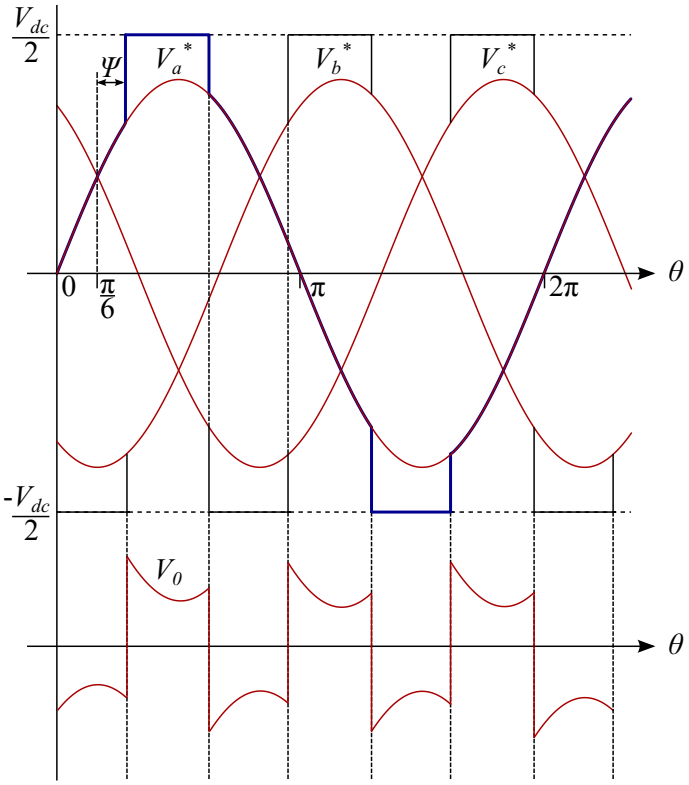

Figure 2.3: Addition of a zero sequence signal to sinusoidal modulating waveforms in a discontinuous PWM scheme

but at higher demanded voltages the situation is reversed. Hence [13] proposes a scheme which switches between continuous and discontinuous schemes depending on the MI.

2.1.1.2 Space Vector Modulation

Space vector pulse width modulation, often shortened to space vector modulation (SVM), has become the de facto standard for IM drives, aided by the proliferation of high-speed, low-cost digital controllers [12, 17].

SVM is a form of pulse width modulation which considers the effect of the switching action on all three phases of the machine in order to maximise utilisation of the DC bus voltage [12]. The maximum achievable MI is increased from $\frac{\sqrt{3}}{2}$ to 1, an improvement of 15.7 %. The voltage demand from the outer control loop (discussed in Section 2.1.2) and the voltages resulting from each possible switching state of the inverter are represented by vectors on a 2D plane. This plane represents a cross-section of the motor, with the origin on the axis of rotation and the horizontal axis, called the direct (d-) axis, aligned with the flux produced by positive current flowing in the stator phase winding labelled A, as shown in Figure 2.4 [18]. The vertical

axis is called the quadrature (q-) axis. This is referred to as the stator (s-) reference frame, indicated by the superscript ^{s}.

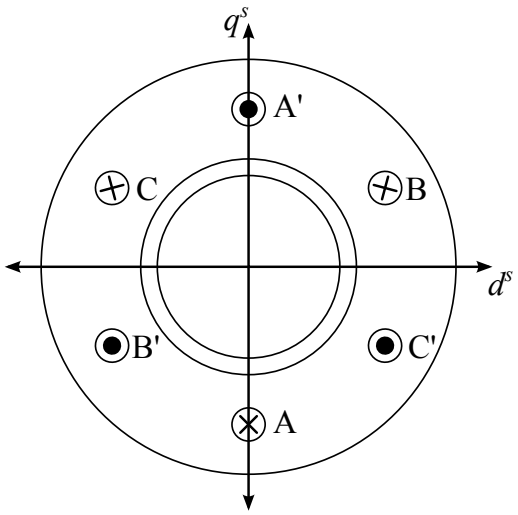

Figure 2.4: The stator reference frame of the space vector system

Using vector mathematics, a combination of the fundamental state vectors is calculated which recreates the demand vector.

There are eight possible switching states: six 'active' or 'non-zero' states , in which current is supplied to the stator from the DC link, and two 'zero' states, in which the stator windings are isolated from one side of the DC link. The voltage vectors produced by these are shown in Figure 2.5, expressed in the s-frame. Vectors V_1-V_6 represent the active states and V_0 and V_7 the zero states. This diagram also shows the numbers of the six sectors, with the sector containing the demand shaded.

The SVM scheme takes a voltage demand as its input, expressed as d and q components in the stator reference frame. The two fundamental states closest to this demand vector are determined by comparing angles, as it is these states which will be used to recreate it. The angle α of the demand vector from the adjacent clockwise fundamental vector, and which of the six sectors it lies in, must be determined, for which there are a variety of methods. This angle is shown for an example, where the demand vector V^* lies in the third sector, in Figure 2.5.

A combination of the two adjacent fundamental vectors can be found which is vectorially equivalent to the demand. The deconstruction of the demand into these components, denoted V_a and V_b, is shown in Figure 2.6. The calculation of the fundamental components V_a and V_b for the example demand vector in this figure are shown in (2.1).

$$\begin{bmatrix} V_a \\ V_b \end{bmatrix} = \frac{2}{\sqrt{3}} m \begin{bmatrix} \sin(\frac{\pi}{3} - \alpha) \\ \sin\alpha \end{bmatrix} \tag{2.1}$$

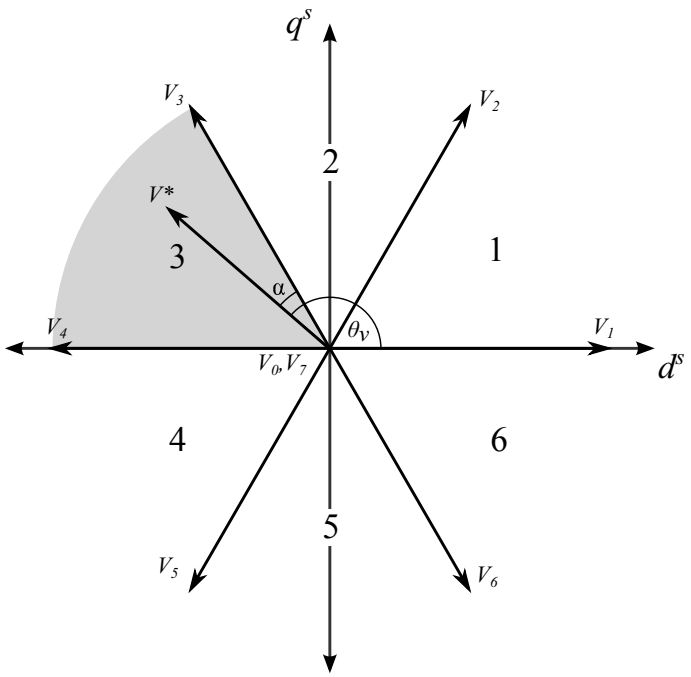

Figure 2.5: Fundamental state vectors V_0–V_7 in the stator reference frame, with example demand vector V^*

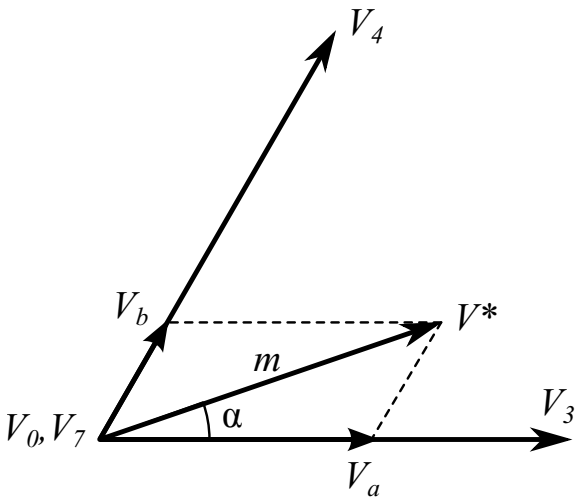

Figure 2.6: Synthesis of demand vector using fundamental vector components

As with sinusoidal PWM, time is divided into discrete periods of T_{pwm}, giving a fixed frequency. The voltages calculated in (2.1) are synthesised by dividing this time appropriately between the fundamental vectors. The periods T_a and T_b for which the fundamental vectors clockwise and anti-clockwise of the demand, respectively, are to be applied are calculated in (2.2) where V_{dc} is the DC link voltage.

A combination, which varies according to the scheme used, of the zero vectors V_0 and V_7 is applied for the remainder of the PWM period T_0. The ratio of the sum of T_a and T_b to T_0 controls the effective amplitude or MI.

$$\begin{bmatrix} T_a \\ T_b \end{bmatrix} = \frac{3}{2V_{\text{dc}}} \begin{bmatrix} V_a \\ V_b \end{bmatrix}$$
$$T_0 = T_{\text{pwm}} - T_a - T_b \tag{2.2}$$

There are many different ways of sequencing the vectors within the PWM period [19]. The most common is the centre-aligned method where the pattern is symmetrical within the PWM period and both zero vectors are used. The symmetrical waveform results in lower harmonic distortion [20]. Using both the zero vectors results in equal utilisation of all the devices [12]. This can be implemented by means of comparison of a reference value for each phase to a triangular carrier signal, similar to classical PWM. These reference values are related to the times calculated previously, but their calculation varies depending on the sector. The calculation of the reference values for the example is shown in (2.3), where $m_{\max}$ is the maximum value of the carrier waveform.

The resultant switching pattern during one PWM period for the example demand vector, which lies in sector 3, is shown Figure 2.7.

$$\begin{aligned} m_A &= \frac{m_{\max}}{T_{\text{pwm}}} \frac{T_0}{4} \\ m_C &= m_B + \frac{m_{\max}}{T_{\text{pwm}}} \frac{T_a}{2} \\ m_A &= m_C + \frac{m_{\max}}{T_{\text{pwm}}} \frac{T_b}{2} \end{aligned} \tag{2.3}$$

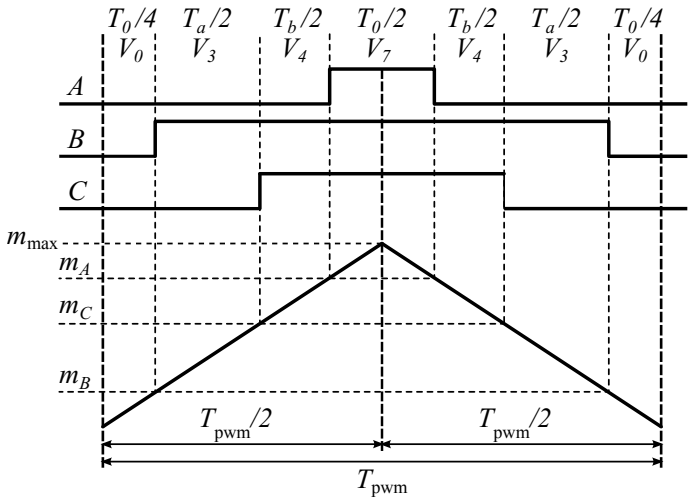

Figure 2.7: Production of switching signals using triangular carrier signal and calculated reference values in a centre-aligned SVM scheme

2.1.2 Speed and Torque Control

This section deals with techniques to control speed or torque, either directly, or through a current controller, as in the hierarchy identified in Figure 2.2.

2.1.2.1 Constant Volts per Hertz

The most basic form of speed control is an open loop technique known as Constant Volts per Hertz. The name arises from the principle of keeping the applied voltage in proportion to the frequency, thus maintaining a constant flux. In practice, the voltage must be "boosted" at low speed to account for the significance of the voltage drop across the stator resistance.

2.1.2.2 Field-Oriented Control

One of the most common techniques for controlling variable speed induction motor drives where high performance is required and control complexity is not severely restricted, such as in an electric vehicle drive, is field oriented control (FOC), also known as vector control [21].

Field-oriented control is a method of controlling the stator current in relation to the rotor flux, ultimately allowing control of the torque [22]. By decomposing the stator current into components aligned with and orthogonal to the rotor flux, the controlled values become DC quantities. This allows PI controllers to be used and thus fast, accurate control to be achieved [12]. The result is a drive with excellent dynamic response.

The rotor flux position can be obtained directly using sensors in the air-gap, but this is undesirable for cost and reliability reasons. Instead, the rotor flux is usually inferred from the stator winding currents and the rotor mechanical position.

An induction machine produces torque when the electric field created by the stator windings rotates faster than the mechanical rotation of the rotor. The difference between these speeds, expressed as a per unit quantity relative to the electrical frequency, is known as the slip s and is defined in (2.4). ω_s and ω_r are the angular velocities of the stator electric field and the rotor respectively, f the electrical supply frequency and p the number of pole pairs.

$$s = \frac{\omega_s - \omega_r}{\omega_s} \tag{2.4}$$

$$\omega_s = \frac{2\pi f}{p} \tag{2.5}$$

The slip can be calculated from the position of the stator current vector and physical knowledge of the machine: specifically, the rotor time constant, the ratio of rotor inductance to resistance. This is a disadvantage of FOC: the controller must be tuned to the individual machine it is used with, and it is sensitive to variations in this parameter, particularly due to the resistance variation with temperature [23].

The complete structure of a field-oriented controller can be seen in Figure 2.8.

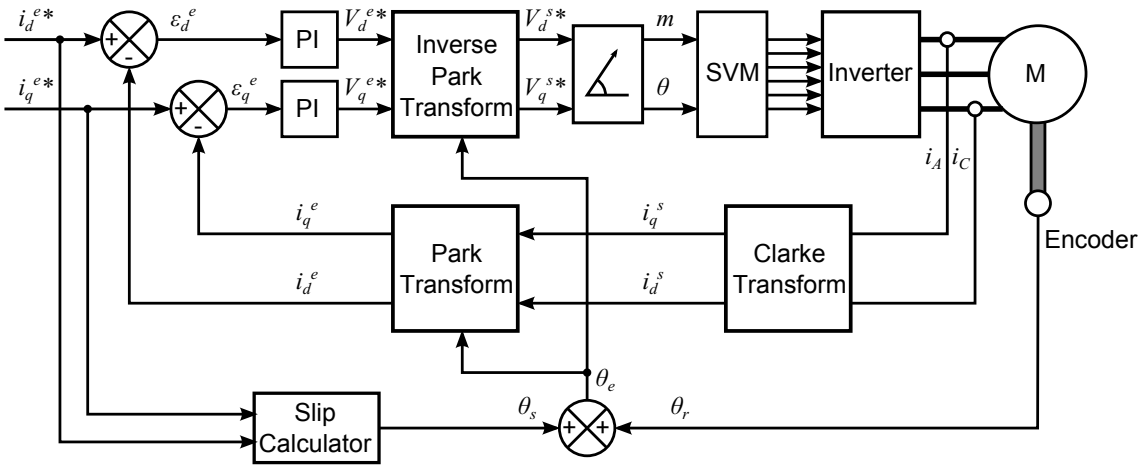

Figure 2.8: Structure of a field-oriented controller

From the rotor time constant T_r and stator currents, the slip angle θ_s can be estimated, which, with the measured rotor position θ_r allows calculation of the orientation of the electrical field, which is coincident with the rotor flux and is denoted by the angle θ_e. The demanded current, and hence voltage, is resolved into components aligned with and orthogonal to the rotor flux. If this is accurately achieved and maintained, the aligned, or direct, component controls the magnetisation of the machine and the orthogonal, or quadrature, current the torque produced. The magnitudes of these components are constant under steady state conditions and so PI controllers may be used.

The estimation of slip and electrical field angles are shown in (2.6) and (2.7) for a discrete-time system where $\theta_{s(n-1)}$ is the slip in the previous sample period and F_{pwm} is the PWM frequency. This is the same frequency at which the currents and rotor position are sampled. The demanded values of the stator current components i_d^{e*} and i_q^{e*} are used under the assumption that the controller is accurately reproducing these.

$$\theta_{s(n)} = \theta_{s(n-1)} + \frac{1}{F_{\text{pwm}} T_r} \frac{i_q^{e*}}{i_d^{e*}} \tag{2.6}$$

$$\theta_e = \theta_r + \theta_s \tag{2.7}$$

The desired flux and torque determine the direct (d-) axis and quadrature (q-) axis current demands, respectively, when expressed in the excitation reference frame (e-frame), the d-axis of which is aligned with the rotor flux. The flux demand is normally the rated flux of the machine, up to base speed, whereafter it is reduced to achieve field weakening. In an EV application, the torque demand is representative of the accelerator pedal position, which makes the operation of the vehicle similar to that of an internal combustion engine.

Two of the phase currents are measured, allowing the third to be calculated from the knowledge that the three phase currents must sum to zero in the absence of a neutral connection, according to Kirchoff's current law. The currents are then expressed in the stator reference frame (s-frame), in which the d-axis is aligned with the flux from stator phase A winding, using the Clarke transform, (2.8).

$$\begin{bmatrix} i_d^s \\ i_q^s \end{bmatrix} = \begin{bmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & \frac{\sqrt{3}}{2} & \frac{\sqrt{3}}{2} \end{bmatrix} \begin{bmatrix} i_A \\ i_B \\ i_C \end{bmatrix} \tag{2.8}$$

The two-axis currents are then rotated into the excitation reference frame using the electric field angle θ_e obtained from (2.7) and the Park transform shown in (2.9).

$$\begin{bmatrix} i_d^e \\ i_q^e \end{bmatrix} = \begin{bmatrix} \cos\theta_e & \sin\theta_e \\ -\sin\theta_e & \cos\theta_e \end{bmatrix} \begin{bmatrix} i_d^s \\ i_q^s \end{bmatrix} \tag{2.9}$$

Expressed in the rotating e-frame, the stator currents are DC values in the steady state and error values are produced by comparing with the demand values, as in (2.10). These errors form the inputs to two PI controllers which produce a voltage demand.

$$\begin{bmatrix} e_d \\ e_q \end{bmatrix} = \begin{bmatrix} i_d^e* \\ i_q^e* \end{bmatrix} - \begin{bmatrix} i_d^e \\ i_q^e \end{bmatrix} \tag{2.10}$$

The inverse transform of (2.9), shown in (2.11), is applied to rotate the voltage demands back into the s-frame, where the modulation scheme will determine the timings of the gate signals necessary to implement them.

$$\begin{bmatrix} v_d^s \\ v_q^s \end{bmatrix} = \begin{bmatrix} \cos\theta_e & -\sin\theta_e \\ \sin\theta_e & \cos\theta_e \end{bmatrix} \begin{bmatrix} v_d^e \\ v_q^e \end{bmatrix} \tag{2.11}$$

2.1.2.3 Direct Torque Control

Direct Torque Control (DTC) [24] was developed as an alternative to FOC, but shares the same aim. Torque and stator flux control are achieved by hysteresis controllers, without an intermediate current controller. The hysteresis controller

Table 2.1: DTC Vector Table [24]

Sector		1	2	3	4	5	6
$c_\phi = -1$	$c_T = -1$	V_2	V_3	V_4	V_5	V_6	V_1
	$c_T = 0$	V_7	V_0	V_7	V_0	V_7	V_0
	$c_T = 1$	V_6	V_1	V_2	V_3	V_4	V_5
$c_\phi = 1$	$c_T = -1$	V_3	V_4	V_5	V_6	V_1	V_2
	$c_T = 0$	V_0	V_7	V_0	V_7	V_0	V_7
	$c_T = 1$	V_5	V_6	V_1	V_2	V_3	V_4

output states determine which of the fundamental voltage vectors will be applied, according to a simple logic contained in a lookup table, which is dependent on which sector the stator flux lies in. The classical DTC lookup table is shown in Table 2.1. Note that in DTC, the sector refers to the range of angles 30 degrees either side of the fundamental vectors; they are phase shifted 30 degrees clockwise from the conventional sector definitions in SVM shown in Figure 2.5.

DTC is inherently shaft-sensorless (encoderless) as it relies only on measurements of voltage and current for flux and torque estimation.

In the original DTC scheme, the voltage vectors are applied asynchronously, and so the switching frequency is variable and the average frequency is dependent on the hysteresis band widths. Variations on this have since been proposed which allow the vectors to be applied using SVM, allowing a constant switching frequency.

Comparing the schematic diagram of DTC in Figure 2.9 with the FOC equivalent in Figure 2.8, it can be seen that the DTC is simpler to implement.

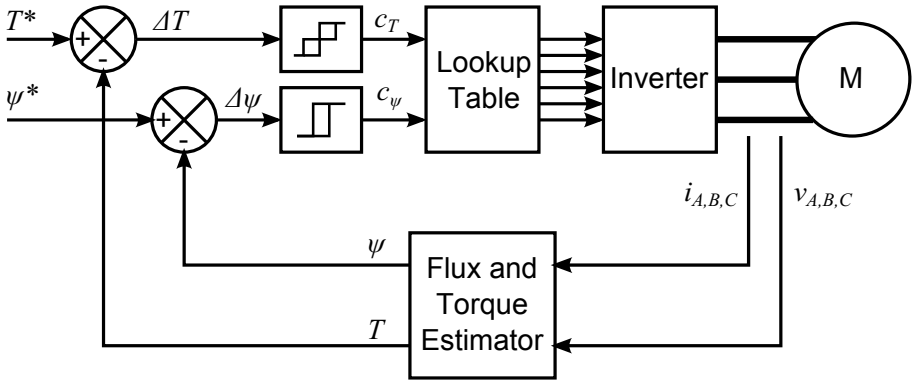

Figure 2.9: Structure of a Direct Torque Controller

Apart from its simpler implementation, requiring no position feedback, it is less sensitive to parameter variation than FOC and provides a better dynamic response

for similar average switching frequencies. However, without additional modulation, it has a variable switching frequency and it also has accuracy issues at very low speed [25]. Most importantly for this application, DTC exhibits significant current and torque ripple in the steady state [26, 27]. Since EV motor operation is characterised by long periods of steady-state operation and slow transients, the benefits of DTC are outweighed by this, and FOC is usually regarded as a better solution for this application. For applications where high dynamic response is critical, such as machine tool and robotics applications [28], DTC may be preferable.

2.2 Efficiency Optimisation Techniques

In developing a new control strategy for efficiency optimisation of an IM drive, it is instructive to study existing techniques in order to understand what has been tried and achieved previously. This field is commonly termed loss minimisation control (LMC). This section examines various schemes and identifies common features in order to categorise them. Much of this is derived from a recent review of this subject [29]. This study firstly sub-divides LMC techniques into offline—where the optimisation is carried out at the design or commissioning stage–and online—where the optimisation is carried out in real-time whilst operating—methods. Further categories are shown in Figure 2.10.

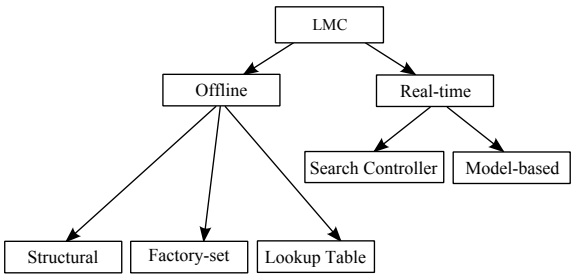

Figure 2.10: Categorisation of loss minimisation techniques [29]

2.2.1 Offline Techniques

Three approaches to offline loss minimisation are considered here. Some studies have attempted to find an optimum switching strategy, whilst others have investigated the effect of physical machine construction, and still more have optimised control parameters to create lookup tables for use online. In all cases, the losses are determined under different conditions (different PWM types, motor geometries and control parameters) in order to find those which optimise efficiency. The effect

can be assessed either by simulation or measurement, which can be done directly, as in [30], in which a calorimeter is used, or indirectly, using the input-output method.

An attempt to find a PWM switching scheme that improves overall drive efficiency is presented in [31]. The authors point out that conventional sinusoidal PWM has the disadvantages of a limited output voltage range, and creates very short pulses near the peaks of the modulating waveform, which cause switching loss but do not contribute to torque production [32]. This is compared with two more advanced strategies: SVM and discontinuous PWM. These schemes add a different common mode signal to the sinusoidal modulating waveform. However, the tests performed on two small (375 W and 2.2 kW) IMs did not show any significant effect on efficiency, and the conclusion is drawn that effort is better focussed on improving motor efficiency.

An example of offline loss minimisation by varying motor design is given in [33]. This study focusses on the geometry of the rotor slots, which are considered in fine detail. However, the same approach could be applied to other similarly detailed areas of motor design, or to the motor as a whole in less detail. The rotor slot is parameterised and a model of the motor developed using a combination of an equivalent circuit and finite element analysis (FEA). An iterative process is then used to find the set of parameter values which maximises the calculated efficiency for given constraints based on the intended application of the motor. This process results in some small increases in efficiency, differing by application, but is of interest here only as a practical example of an offline LMC technique.

In the past, when there were greater limitations on processor power, optimisation was more likely to be performed offline to reduce the amount of online computation needed. An example of this is presented in [34]. The slip which gives the greatest efficiency at each point over a range of speeds is determined by measuring the input and output power of the motor in a test procedure. This information is compiled into a lookup table and used to impose the optimal slip for maximum efficiency, depending on the speed.

2.2.2 Real-time Techniques

Two distinct approaches to minimising losses online in real time are used in modern LMC systems.

A search controller, referred to as a physics-based [29] or "perturb-and-observe" technique, iteratively varies a control variable, such as flux or current, and measures its impact on a controlled variable, commonly input power, to find a minimum for the controlled variable.

A model-based controller minimises loss by calculating the value of a control variable which minimises the input power or loss according to a model based on motor parameters.

2.2.2.1 Search Controllers

A search controller for loss minimisation is developed in [35]. Like most search controllers with this objective, the goal is to minimise the input power, and this is achieved by finding an optimal value of air gap flux. In order to calculate the input power, a loss model is developed that includes the copper, stator iron and stray losses. The input power can then be determined by summing this loss and the output power, calculated from the product of mechanical speed and torque. This value is not very accurate, and moreover, it has a low sensitivity with respect to the air gap flux, and so it is difficult for the controller to find the minimum. This results in oscillations in the air gap flux which cause undesirable torque variations.

An alternative is to minimise the DC link current, which is proportional to the input power if the DC link voltage is assumed to be constant; a reasonable assumption in practice. This reduces the computation time required in comparison with the loss model, but it is still relatively insensitive to air gap flux variation. As such, it is concluded that using the input power or DC link current as the variable to be minimised is only effective on small, inefficient IMs.

Furthermore, it is then shown that the stator current is much more sensitive to variations in air gap flux and that it is minimised when input power is minimised, and therefore makes a more effective and practical controlled variable. In addition, oscillations of the air gap flux are reduced, as are inverter losses. Experimental and simulated results are presented showing that minimising the stator current results in optimal efficiency.

The search controller presented in [36] illustrates well the difference between online and offline optimisation. Similarly to the offline scheme in [34], this algorithm is based on the principle that efficiency can be maximised by operating at an optimal value of slip, dependent on the speed and torque. However, in this case the determination of the optimum slip value is performed in operation, with the lookup table being populated as new operating points are encountered. When the system is operated at a speed and torque–calculated from stator voltage and current measurements—for which an optimal slip has not yet been determined, the slip is varied by a small amount and the input power calculated until a minimum is found. When operating at this point in the future, the optimal slip value will be read from the lookup table, as shown in Figure 2.11.

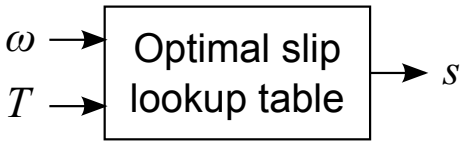

Figure 2.11: Lookup table of optimal slip from [36]

Another method that could be considered of this category is that of dynamic frequency scaling. An investigation of this method is presented in [37] in the context of a single-phase DC-DC converter. The principle of the scheme is to reduce the switching frequency where necessary to reduce switching loss. When a higher switching frequency is required to achieve a faster response to transient conditions, the switching frequency is increased, effectively increasing the bandwidth of the control system. The error between demanded and applied voltage is continuously monitored to determine when a change of frequency is necessary. No application of this method to a motor drive has been found, but it is not inconceivable.

2.2.2.2 Model-Based Controllers

A LMC based on a model of the IM losses is developed in [38]. A model of the IM based on the conventional d-q axis equivalent circuit is used and the losses are calculated from this. The model is manipulated and the result obtained that the loss is minimised when the losses in the d and q axes are equal. This condition forms the basis of the control algorithm. Results from simulations and experiments at a variety of operating points are given, showing a marked improvement on the 'conventional' LMC scheme in [39], a search controller which is often used as a benchmark.

2.2.2.3 Search Controllers vs. Model-Based Controllers

The main disadvantages of the search controller are accuracy [35] and speed [29]. The system's ability to accurately find the optimum operating point is limited by the accuracy of the measurement of the controlled variable. This is usually the input power, which is particularly difficult to minimise, although better results can be achieved by controlling stator current as described in [35]. Search controllers take longer to converge on the optimum point than model-based controllers, due to the time taken to make each change to the control variable and measure its effect.

Another disadvantage of search controllers is their tendency to oscillate around the optimal point. This can cause torque ripple according to [40]. This can be reduced by varying the perturbation step so that it is smaller around the optimum [29].

Model-based controllers generally have a shorter settling time than search controllers [29]. They can be more accurate, as they do not rely as heavily on real-time measurements, although they are dependent on an accurate loss model. Since the model relies on knowledge of motor parameters such as resistances and inductances, inaccuracies in determining these and variations during operation reduce the effectiveness of the controller.

This effect is avoided in [41], however, where the loss minimisation condition is derived from the IM model, but does not require knowledge of the parameters for its implementation.

2.2.2.4 Hybrid Techniques

An LMC technique that can be classified as neither model-based or a search controller, but features aspects of both, is identified in [29], and so is described as a hybrid LMC. Ripple correlation control (RCC) has been used in a variety of power electronics control applications, including maximum power point tracking (MPPT) in photovoltaic systems and static VAR compensation [42,43]. This method exploits the inherent current ripple caused by switching converters. The ripple is used as a constant perturbation which can be used to find a maximum (such as in MPPT) or minimum (as in LMC).

Another hybrid between model-based controllers and SCs is presented in [44]. This uses a simplified loss model which reduces the sensitivity to parameter changes and fuzzy logic to implement the search function, which claims to reduce the settling time.

2.2.2.5 Dynamic Response

It is important to consider the impact of loss minimisation on controller performance. The effect of most LMCs is to reduce the flux, whether directly through explicit control of the flux-producing current, or indirectly as a result of minimisation of the input power, for example. Operating at a lower flux level adversely affects the dynamic response of the drive [35].

This problem can be avoided by supplying nominal flux to the motor when a torque demand is detected and greater performance is required, reverting to LMC when operating in the steady state.

Although a simple solution, this approach means the losses are not minimised during the period of the transient. If the application demands significant periods of transient response, the efficacy of the controller may be significantly reduced if this problem is not addressed. A solution is proposed in [45] which builds on the flux

control-based method developed in [46] for steady-state loss minimisation, by estimating the flux during transients.

2.3 Predictive Control

Generically, the approach of using a mathematical model of a system to predict its operation under a given set of conditions, and the subsequent use of this information to control one or more variables to desired values, is known as predictive control. This is an intuitive method of control, and yet it usually requires more calculations than conventional control techniques [47]. The vastly reduced cost of processing power in recent years has resulted in control schemes of this type becoming cost-effective in a greater range of applications, and research interest has led to a number of different schemes being proposed.

2.3.1 General Predictive Control Methods

At the top level, there are several basic approaches to IM drive control, of which predictive control is one. Linear controls, such as PI controllers, are a common component of FOC systems [48]. Even simpler are hysteresis current controllers.

The defining principle of predictive control is, as the name suggests, use of a prediction of the motor's future state to influence its control. Within this broad definition, further distinct methods exist. The most successful of these is model predictive control (MPC), but others have also been proposed.

2.3.1.1 Hysteresis-based Control

The simplest of these is the hysteresis-based predictive current controller [49]. This is very similar in structure to the classic hysteresis current controller, except that instead of the measured current, a prediction of current in the next switching period for each possible switching state is used to determine the chosen state. The time taken for the current to reach the edge of the hysteresis band (or the analogous boundary circle for a two-dimensional space vector) is predicted and the state that results in the longest time is applied.

This method is illustrated in Figure 2.12 where the present ($i^*_{s_n}$) and predicted ($i^*_{s_{n+1}}$) current demand vectors are shown. The hysteresis boundaries associated with each of these is marked. The predicted trajectory of the current vector for two possible switching states is shown, with time to reach the hysteresis boundary again calculated. In this example, the switching state for which this time is t_1 will be applied.

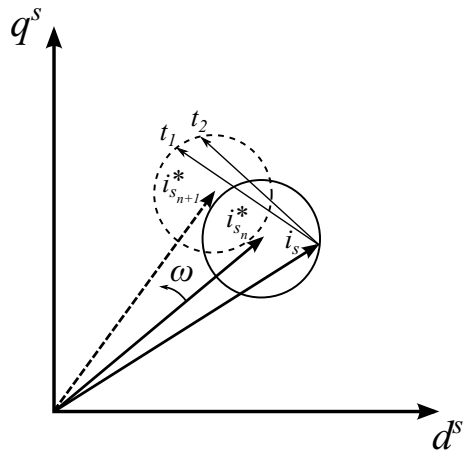

Figure 2.12: Principle of hysteresis-based predictive control

2.3.1.2 Deadbeat Control

Deadbeat current control is an intuitive way of obtaining a fast current response. The principle is simply to calculate the voltage which will produce the demanded current if applied in the next switching period. Since the current response to a given voltage can be predicted, it follows that the necessary voltage can be calculated to give a certain current, as long as the parameters of the machine are known. This voltage is then applied with a modulator, and the structure is very similar to a linear controller where the deadbeat controller has replaced the PI controller [50].

2.3.1.3 Trajectory-based Control

One advantage to predictive control is that it offers the opportunity to reduce response times for 'outer loop' variables (see Section 2.1.2) by directly predicting their behaviour and controlling the switching state accordingly. Whilst most predictive control schemes predict absolute values of the controlled variable, trajectory-based predictive control schemes consider the future trajectory of the controlled variable.

Direct speed control [51] is an example of this type of controller. The switching states are characterised in terms of their influence on acceleration a and speed error e, resulting in a set of trajectories on the a-e plane. The system applies the switching states which result in the operating point on this plane remaining as close as possible, within the constraints set by the switching frequency, to the desired point, in this case the origin. In this way, the system has an area of hysteresis.

Direct torque control [24] shows elements of both hysteresis- and trajectory-based control, and again demonstrates direct control of an outer loop variable. The torque and flux are controlled within hysteresis bands, but the controller also considers direction and proximity to the demand value when choosing a voltage vector to provide a damped response, similar to trajectory-based systems.

2.3.2 Model Predictive Control

Model predictive control is the term given to a class of predictive control strategies with the same basic structure which can be adapted to various objectives. The controlled variable or variables—for instance, torque—are modelled for each possible switching state. The results are compared using a cost function, which calculates the error between each modelled variable and its demanded value. The state which minimises the cost function is then applied. Where multiple variables are considered, the relative importance may be adjusted by introducing weighting factors. Constraints such as current limits can also be introduced by saturating the cost function output.

2.3.2.1 Finite-Set MPC

Since the inverter inherently has a discrete number of states—eight in the case of a two-level three-phase inverter—the required processing can be minimised by only considering these states, and directly applying the minimising state. However, this creates a variable switching frequency which results in torque pulsations [52].

2.3.2.2 Continuous-Set MPC

In order to achieve the desirable flexibility of finite-set MPC with a fixed switching frequency, it is necessary to consider a larger set of possible voltage vectors. The cost function minimum will then not necessarily represent a fundamental vector, but a vector which can be applied by a conventional modulator, such as SVM. The number of voltage vectors determines the resolution of the voltage control. Between 12 and 16 bits is considered appropriate [53].

2.4 Loss Modelling

In order to implement a model-based control scheme which maximises efficiency in a system, it is desirable to consider as many sources of loss as possible over the range of possible operating conditions and how they might be effected by the control scheme. This requires a model of the system which is able to predict the resulting loss given a set of state variables which describes the current operating point.

A hierarchy of the sources of loss in an electric drive can be seen in Figure 2.13. This breakdown informs the following discussion on methods of estimating and predicting losses.

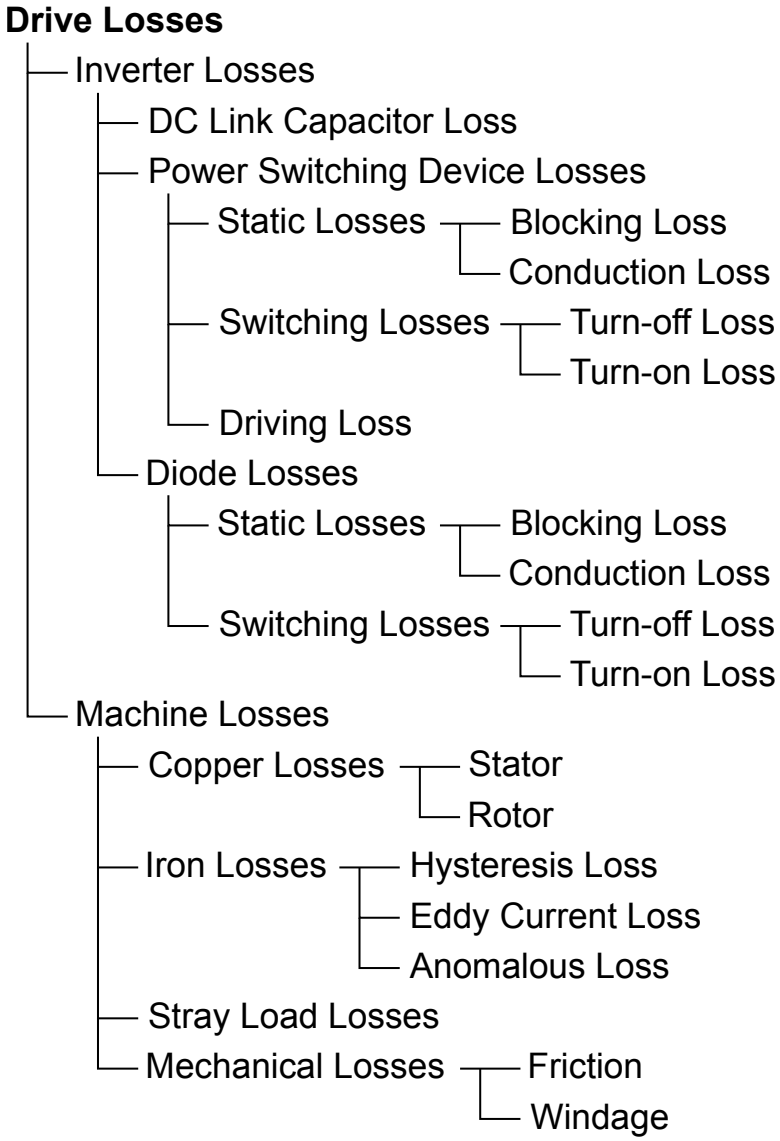

Figure 2.13: Breakdown of power losses in an electric drive

2.4.1 Inverter Losses

Inverter losses are typically divided into static and switching losses. These losses need be considered for both the power switching devices and their anti-parallel diodes.

Practical tests in [54] indicate that inverter losses in a PWM-controlled induction machine drive typically account for between a quarter and a half of total losses.

Static losses occur in the steady state - i.e. when the device is either fully on or fully off - and so are split into blocking (off-state) and conduction (on-state) losses. The conduction loss represents the ohmic loss caused by current flowing through the device, which in practice has some resistance. Blocking loss is caused by a very small

leakage current flowing in the device in the off-state, and is often ignored because it is very small in modern power devices [55, 56].

Whilst the ideal device switches between states instantaneously, in practice this is not the case and the device voltages and currents have finite rates of change which are limited by numerous physical processes. When the IGBT is turning on, the current rises at a non-zero rate, and only after it has reached its steady-state value does the drain-source voltage begin to fall to zero, leading to losses during the process. Similarly, when the IGBT turns off, the current does not begin to fall until the drain-source voltage has risen to its final value. Turn-on loss in modern diodes is very small and often ignored, but the phenomenon of reverse recovery during diode turn-off, in which a significant reverse-bias current flows for a short period of time before returning to zero, causes loss in the diode. Furthermore, where the current is commutated from a diode to an IGBT, such as in a complementary-switched inverter leg, the reverse recovery current flows in the IGBT and thus affects the loss in its turn-on transient.

There are a number of other incidental losses associated with the drive system. These include losses in the gate drive circuits; busbars and other conductors; relays or contactors; and power supplied to control circuits, etc. These losses are usually insignificant compared to the power component losses and, in any case, are of little interest to the control scheme designer, since they are unavoidable and unaffected by the control scheme.

2.4.1.1 Modelling

Examples of techniques and calculations presented here are taken in the main from three academic studies, which approach the problem in different ways and with varying levels of detail. The inverters in [55] and [56] are intended for HEV applications. [57] is intended to model inverters used in power transmission systems, but the concepts seen are applicable to other inverter applications, including EVs. [56] models the inverter as part of a complete vehicle model, including an IM, battery, internal-combustion engine, transmission, etc. The inverter model in [55] is also intended to be used as part of a whole-vehicle simulation, but concentrates on the development of motor and inverter models to be used within an existing vehicle model. Because of its application, the inverter model in [57] is developed independently of other systems, though a thermal model of the inverter is designed alongside it.

The approaches taken can basically be summarised as follows:

- Physical model [55]
- Simple analytical model [56]
- Complex empirical model [57]

Each of these approaches is outlined here, but further details are provided in Appendix B.

2.4.1.2 Physical Model

One approach to loss modelling is to use constants that describe the physical attributes of the devices with the state variables (voltage, current) to calculate the loss. This is the approach taken in [55].

IGBT Switching Loss
The IGBT switching loss is expressed as an average value over the total output period, which simplifies the calculation over multiple switching periods if a constant switching frequency is used. Note that the averaging is conceptually separate from the modelling approach; the same modelling approach could be used to calculate energy loss for a single event or period.

In this model, the turn-on and turn-off losses are expressed using constants which take into account physical aspects of the IGBT, such as transconductance, breakdown voltage and semiconductor permittivity.

IGBT Static Loss
The IGBT conduction loss is based on the collector-emitter resistance, demonstrating its physical origin as an ohmic loss. The loss is expressed in terms of the peak line current, averaged over a period according to the modulation index. The IGBT blocking loss is not considered in this study.

Diode Switching Loss
The primary factor affecting switching loss in the diodes is the process of reverse recovery, and in this model is considered to be the only source of loss. The parameters used to quantify the reverse recovery loss are the physical parameters snappiness factor and reverse recovery time. Again, the loss is averaged over a sample period.

Diode Static Loss
The diode blocking loss is ignored here, but the conduction loss is determined using the diode forward resistance and forward voltage drop.

2.4.1.3 Empirical Model

Whilst the previous model considers the actual physical processes which result in power loss, the empirical model presented in [57] takes advantage of the fact that the instantaneous power loss can always be determined as the product of voltage and current at that instant. Thus switching losses can be estimated by parameterising the voltage and current during switching by examining typical waveforms. The method presented in [57] is an example of this, taken to a very high level of detail. Whilst this would not be practical in the application under discussion here, it is nevertheless instructive.

IGBT Switching Loss

The IGBT turn-on loss model in [57] is obtained by studying the collector current and collector-emitter voltage characteristics during a typical turn-on event, shown in Figure 2.14. Distinct periods in the waveform are identified to enable an analytical expression of loss for each stage of the switching event. In essence, these describe the voltage and current waveforms, and taking the product and integrating over the period of the turn-on event gives the energy lost. The resulting equation contains many parameters, but solving this is computationally less expensive than simulating numerically with a time-step small enough to accurately capture the features of the very short transient.

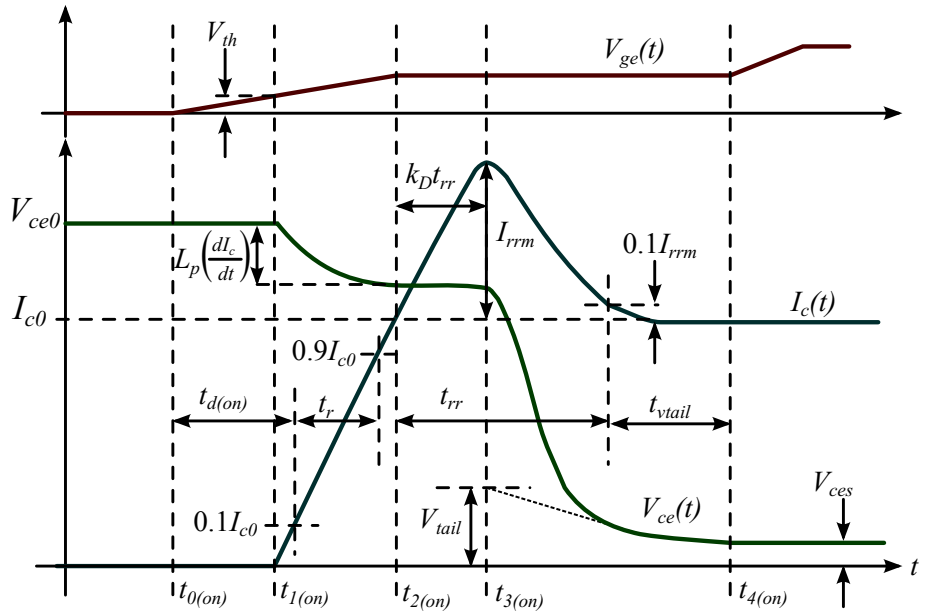

Figure 2.14: Characteristic IGBT turn-on transient waveforms showing quantities used to calculate loss [57]

The fitted waveform is shown along side an actual measured waveform in Figure 2.15.

Detailed methods for calculating the many parameters in the above equations are explained in the paper. They could also be obtained by fitting curves to measured

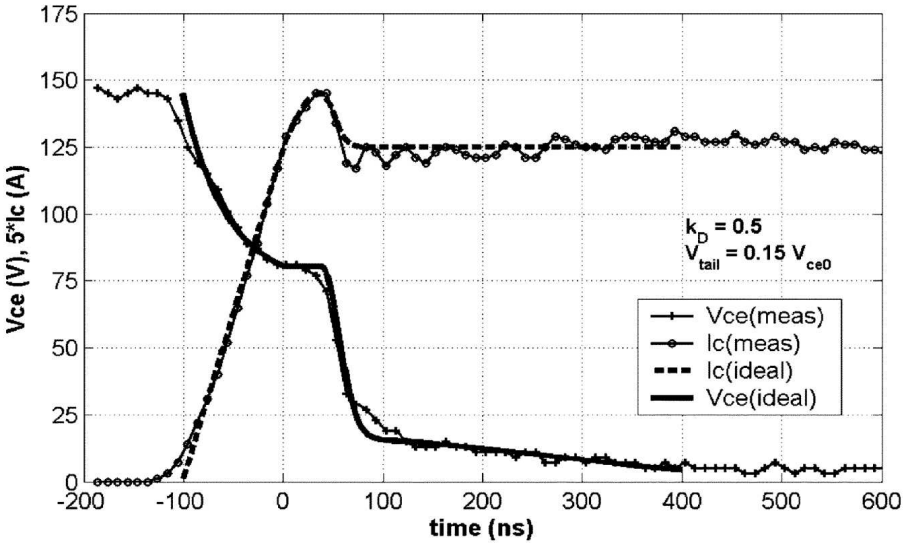

Figure 2.15: Comparison of measured and fitted IGBT turn-on transients [57]

characteristics.

One important point to note is the dependence of the IGBT turn-on loss on the reverse recovery current of the opposing diode, which is taken into account in this model.

The IGBT turn-off loss model is obtained in a very similar way, and so the derivation and results will not be reproduced here.

IGBT Static Loss

The IGBT conduction loss in this model is simply the product of device current and forward saturation voltage. This model also considers blocking loss, given as the product of blocking voltage and leakage current, since it is intended for high power transmission applications where this is not insignificant.

Diode Switching Loss

Only the diode turn-off loss is included in this model, with the turn-on loss ignored. The process of reverse recovery is again the key factor in this loss, but the whole turn-off characteristic is modelled.

The diode turn-off loss equation is based on the on-state current, on-state forward voltage drop, peak reverse recovery current, post-switching blocking voltage and the time taken for the reverse recovery current to decay from its peak value to 10% of that value. Again, these values are found by studying measured waveforms.

Diode Static Loss

The static losses in the diode are calculated in exactly the same way as in the IGBTs: as the product of current and forward voltage.

2.4.1.4 Analytical Model

A model of inverter losses for use in the simulation of a plug-in HEV is developed in [56]. This is a fairly simple model which develops analytical expressions to represent empirical data provided by the device manufacturer.

The inverter model in [56] defines the power loss of each cell (one IGBT and its anti-parallel diode) as the sum of the IGBT switching loss, IGBT forward conduction loss, diode reverse recovery loss, and diode forward conduction loss. Blocking losses and diode turn-on loss are not considered.

IGBT Switching Loss
The IGBT turn-on and turn-off energies are obtained by fitting quadratic curves to the loss characteristics given by the manufacturer. The total switching loss is then given by the sum of the turn-on and turn-off energies averaged over the switching period:

IGBT Static Loss
The IGBT conduction loss is calculated from the forward characteristic: the variation of collector current with collector emitter voltage. Again, a quadratic function is fitted to the characteristics, and the collector-emitter voltage multiplied by the collector current at each sample point to give the conduction loss.

Diode Switching Loss
The diode loss model is simplified by ignoring turn-on loss and only considering reverse recovery in the turn-off transient. The reverse recovery energy, averaged over the switching period to give the turn-off power, is expressed as a function of the diode forward current by fitting to the manufacturer's characteristic in the same way as the IGBT.

Diode Static Loss
The diode conduction loss model is very similar to the IGBT conduction loss model due to the similarities in the forward characteristics between the two devices.

2.4.2 Motor Losses

The main components of loss in an IM are the copper loss due to electric currents flowing through the windings (which are usually made of copper) and iron, or core, loss due to magnetic effects in the core material (which is usually steel). Additional loss is caused by mechanical losses and stray load loss. Stray load loss accounts for several smaller loss components with miscellaneous causes.

In general, the copper loss is dominant in smaller machines, whereas the proportion of core and stray load losses increases with machine size [58].

Many models for motor losses are based on the assumption that the supply is purely sinusoidal. This is in practice only true for industrial drives fed from the three-phase mains. However, when fed from an inverter, as is the case in EVs, there are significant harmonic current components and these must be considered. These harmonics do not contribute significantly to the mechanical output power of the machine, but create additional losses [54].

2.4.2.1 Copper Loss

The term copper loss refers to the ohmic losses caused by currents following through the resistive windings of the machine. In an induction machine, currents flowing in the stator induce currents to flow in the rotor windings (or bars), as the name implies. Hence it is common to refer to the copper loss in two parts: the stator and rotor copper loss. However, the rotor current, and hence loss, is often referred to the stator by considering the machine as a transformer and calculating the equivalent stator current. This is a common simplification in IM modelling.

It is common to use an equivalent circuit approach when modelling electric machines, and these are often split into separate circuits for the d- and q- axes. A dq-axis model of an IM with rotor quantities referred to the stator is shown in Figure 2.16.

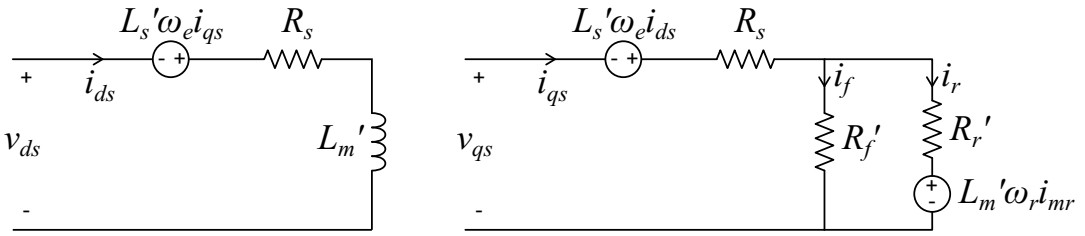

Figure 2.16: Steady-state equivalent model of an IM (a) d-axis, (b) q-axis [59]

The equivalent circuit approach is particularly appropriate for modelling copper loss in both the stator and rotor as their physical origin is an ohmic loss due to electric current flowing and they can be represented by series resistors. These are labelled R_s in Figure 2.16.

In the IM model, the stator copper loss may be calculated as

$$P_{cu-s} = R_s(i_{ds}^2 + i_{qs}^2) \tag{2.12}$$

since the stator current components are state variables [59]. Similarly, the rotor copper loss can be expressed as

$$P_{cu-r} = R_r^{'} i_r^2. \tag{2.13}$$

However, the rotor current, i_r, must be calculated from the stator current as it cannot be directly measured:

$$i_r = \frac{R_f^{'}}{R_f^{'} + R_r^{'}} i_{qs} - \omega \frac{L_m^{'}}{R_f^{'} + R_r^{'}} i_{ds}. \tag{2.14}$$

$R_r^{'}$ is the rotor resistance referred to the stator during the simplification of the equivalent circuit. It can be calculated from the actual rotor winding resistance, R_r using (2.15) or obtained from a locked rotor test [60].

$$R_r^{'} = R_r \left(\frac{L_m}{L_r} \right)^2 \tag{2.15}$$

$L_m^{'}$ is the magnetising inductance referred to the stator by

$$L_m^{'} = \frac{L_m^2}{L_r}. \tag{2.16}$$

2.4.2.2 Iron Loss

When a magnetic field is applied to an iron core, such as by current flowing in the winding of an electrical machine, magnetic domains within the iron align and the core is said to be magnetised. If the field is removed, by reducing the current to zero, the flux density does not return to zero and the iron retains a certain amount of 'residual magnetism,' which varies according to the material. To remove the residual magnetism, a magnetic field must be applied in the opposite direction by reversing the current.

This phenomenon is known as magnetic hysteresis and is classically illustrated by plotting the magnetic flux density, B, against magnetic field strength, H, as in Figure 2.17. During this process, energy is dissipated as heat.

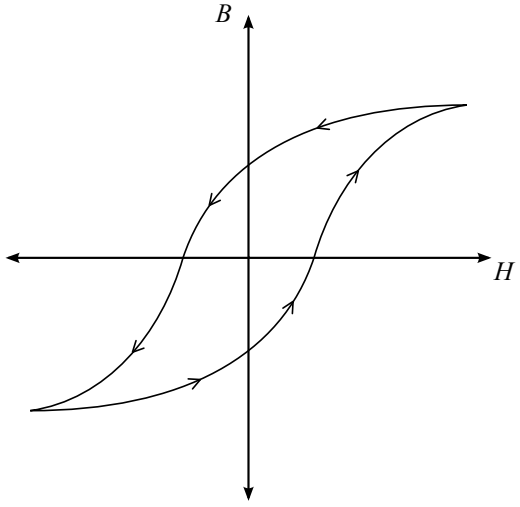

Figure 2.17: Typical magnetisation curve for steel showing magnetic hysteresis [61]

Further loss is caused by eddy currents induced by rapid flux changes [62]. Together these losses are referred to as the iron or core loss.

A simple method of accounting for the iron loss is by inclusion of a resistor in the equivalent circuit. In the IM equivalent circuit in Figure 2.16, the iron loss is represented by the resistor $R_f^{'}$ [59]. The power loss is then calculated as

$$P_{iron} = R_f^{'}(i_{sq} - i_r)^2 \tag{2.17}$$

in conjunction with (2.14).

This model is a simplification however, as the physical mechanisms are much more complex than a simple ohmic loss. More detailed models are studied in Appendix B.

2.4.2.3 Stray Load Loss

Losses not accounted for in the calculation of the iron and copper losses are usually lumped together under the term stray load loss. The physical origins of these losses in IMs were first described in [63] and this paper is often cited in discussions of stray load loss [41, 64]. These are described in Appendix B.

A simpler approximation of the stray load loss for an IM in a control application is presented in [41], where it is proportional to the square of speed and the square of rotor current as shown in (2.18). The rotor current is referred to the stator since it cannot be directly measured in a squirrel-cage IM.

$$P_{str} = c_{str}\omega^2 I_r^{'2}. \tag{2.18}$$

Proportional Approximation

In an even coarser approximation, the standard IEEE 112-2004 ("Standard Test Procedure for Polyphase Induction Motors and Generators") simply assigns difference assumptions for the proportion of output power attributed to stray load loss based on the rated power of the motor for certain test procedures, as shown in Table 2.2 [65].

Table 2.2: Assigned allowance for stray load loss in IEEE 112-2004 test procedures E1, F1 and E1/F1 [65]

Motor ratings (hp)	Stray loss in output power (%)
1-125	1.8
126-500	1.5
501-2499	1.2
$\geqslant$ 2500	0.9

2.4.2.4 Mechanical Losses

Electrical machines suffer from mechanical losses due to friction in the bearings and air moving over and through the machine for cooling purposes (known as windage loss) [66].

Whilst it is necessary to consider the mechanical losses to accurately calculate the efficiency of the system, and when designing a machine, they are largely irrelevant in a control context, since they are dependent only on speed. They may, however, be considered by a HEV system controller when determining the desired operating point of the electric motor.

2.4.2.5 Summary

The key dependencies of each component of the motor loss based on the models investigated here are summarised in Table 2.3.

Table 2.3: Summary of key dependencies for motor loss components

Loss Component	**Dependencies**
Copper	I, R, T
Iron	B, ω
Stray Load	I, ω
Mechanical	ω

The copper losses are ohmic in origin and so are dependent on current and resistance. The current may be calculated separately for each axis in an orthogonal axis model, or for the stator and rotor. The rotor current cannot be directly measured and so is calculated from the stator current: see (2.14). The rotor current is usually referred to the stator by adjusting its value according to the ratio of magnetising and rotor leakage inductances: see (2.15). The resistance has a further dependence on temperature. This can be taken into account by adjusting the resistance according to (B.16).

The iron loss is generally dependent on frequency, air gap flux and slip. The air gap flux can be determined from the frequency, speed, rotor resistance and rotor current using (B.23). The rotor current can be calculated from the stator current according to (2.15).

The stray load loss is a function of current and speed and the mechanical loss is a function of speed and the volume of air in the machine.

Chapter 3

Multi-Objective Controller Structure

The aim of this work as a whole is to investigate the effect on drive system efficiency of the direct inclusion of a loss model in the control loop. To include the loss model in the control system, some changes are made to the structure of the controller. These changes are presented and discussed here and simulation and practical results demonstrating the operation of the proposed structures, without the loss model, are shown. The implementation of the loss model is discussed in Chapter 5, whereas this chapter deals only with the implementation of the new control structure.

3.1 Modification of Control Structure

3.1.1 Introduction

Thanks to the proliferation of low-cost, high-speed microcontrollers, FOC with SVM is widely used in industrial drives in all but the most cost-sensitive situations [67]. These techniques are explained in detail in Section 2.1.2.2 and Section 2.1.1.2. The exponential increase in affordable processing power [68] will be exploited in the novel control scheme proposed in this work. FOC with SVM will be used as a benchmark against which this is measured. As noted in Section 2.1, DTC offers an even faster dynamic response than FOC. However, because of inertia in the mechanical drivetrain and loading, the time constants encountered in EV systems are longer than some industrial applications. Hence a very fast dynamic response is not a major priority.

In order to include the loss model—introduced in Chapter 5—changes have been made to the conventional control methods. The predictive control techniques de-

scribed in Section 2.3.2 have heavily influenced the design of this control scheme. Reasons for and details of changes to the conventional control methods are presented in this section.

Existing predictive control techniques either replace the current controller of the classic controller—as in hysteresis-based control (Section 2.3.1.1) and deadbeat control (2.3.1.2)—or directly control an outer loop variable, such as speed or torque—as in trajectory-based control (Section 2.3.1.3)—thus implicitly controlling current to achieve an input demand.

Power loss differs from these outer loop variables in that controlling it does not imply control of the input demand (in this case torque). To demonstrate this, if the power loss was to be minimised, with no other objective, no power would be supplied to the machine at all, as this represents the lowest possible loss condition. Obviously this does not represent useful control, and so other objectives need to be integrated within the loss minimisation controller.

There are examples of other variables which cannot be controlled in isolation in [69], including DC link capacitor voltages, which may be controlled for the purpose of balancing. These are used alongside current control using a cost function to achieve multiple-objective optimisation. The objectives are combined with a cost function, allowing weighting of the different objectives. This technique will be used in the controller proposed in this work.

3.1.1.1 Aim

The fundamental aim of restructuring the controller is to manipulate the input demands and range of possible output such that a decision may be made on the optimal output. This can be achieved by reducing the problem to a finite number of possible outcomes. This number must be low enough to allow evaluation of each, since whatever computation is required to determine the optimum must be repeated for all cases. However, enough possibilities must be considered to give the freedom required to meet all the control objectives, which in this work are accurate torque control and loss minimisation.

3.1.1.2 Voltage Control

In the proposed control scheme, the functional operation (i.e. torque control) is achieved by minimising the voltage error during the modulation process. The modelled loss can be incorporated at this point using a cost function, thus achieving multiple-objective control. Implementing the control at the voltage stage results in a finite number of states—the eight fundamental switching states—that the voltage

error and loss model need to be evaluated for; this is the key reason for choosing this method over a modification of the current loop.

Existing schemes reduce the current control to a discrete problem, but only by use of a machine model to predict the current for each fundamental voltage vector. Estimating the current in this way is computationally expensive, and so its elimination is desirable.

The conventional SVM scheme does not allow for the inclusion of a secondary objective; the chosen vector is based entirely on a demanded input voltage, which the modulator seeks to reproduce as accurately as possible. The input is a demand voltage in two axes and the output is a set of gate signals which imposes a voltage vector which recreates this. The modulation methods proposed in this chapter aim to introduce a secondary influence: the loss caused by a given output vector. As a result, the accuracy of reproduction of the demand vector is not the sole objective; some accuracy may be sacrificed, in a controlled manner, in order to improve the efficiency. This is the common aim of all the schemes demonstrated here.

3.1.2 Proposed Modulation Methods

The main change to the conventional control scheme described in Section 2.1 is in the way gate signals are produced from the demand values by the modulator. The process of generating gate signals by comparison of reference values to a carrier signal is unchanged, but new techniques are proposed for obtaining these reference values. The processes involved in calculating the reference values in conventional SVM are shown in Figure 3.1. The alternative methods are discussed below, and then key aspects of them, which represent a significant proportion of the novelty of this work, are examined in detail.

Each method will be shown in the same format as Figure 3.1 to allow comparison and aid understanding of the new schemes. The width of signals (i.e. the number of separate values being communicated) is shown by a number next to a diagonal line across the signal path, where it is greater than one, in each of these diagrams. For example, there are three reference values, one for each phase, so this signal has a width of 3.

Three possible designs are presented here which attempt to meet the aims described above. Each design is believed to be novel and has not previously been published. They will be compared to existing techniques to put them in the context of the state of the art. Later in this chapter, after a full and detailed description of each scheme, simulation results demonstrating their operation will be presented.

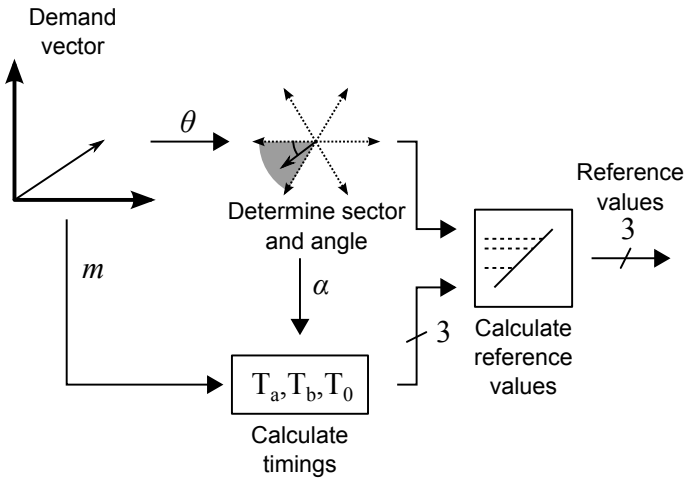

Figure 3.1: Block diagram showing conventional SVM scheme

The number of ways the operation can be varied, referred to as degrees of freedom, is assessed for each method, since this determines how the loss may be reduced with the later introduction of the loss model.

Finite set and continuous set MPC are discussed in Section 2.3.2. To summarise, finite set MPC [52] considers a discrete (finite) number of possible states, which is essential for the evaluation of the modelled control variable, but leads to a non-deterministic switching pattern with a variable switching frequency and hence harmonic distortion of the output current. Continuous set MPC [53] aims to improve on this by starting with a continuous range of possible states. These are then discretised, leading to finer control and a deterministic switching pattern, at the expense of greater computation time.

The single-state modulation technique proposed below is similar to finite set MPC with a single period horizon, with added filtering. The discretised modulation method is more comparable to finite set MPC with a horizon greater than one, or continuous set MPC, but has key differences from both.

Cascaded polar modulation aims to combine the advantages of both these types. This is achieved by use of a cascaded control structure in which the demand vector is considered as a magnitude (or MI) and angle rather than Cartesian co-ordinates, with these components treated separately. The demand MI is calculated and applied precisely, and a further choice of vector made based on the angle.

3.1.3 Single-State Modulation

One of the simplest solutions is to consider how well each switching period meets all the objectives individually, and apply the optimum state for the whole period. This method exploits the inherently discrete nature of the inverter, which, as previously

noted, has eight possible states. The number of computations can be reduced to four if only the two adjacent active states and the two zero states are considered. Moreover, the zero states may be considered as one, since they represent the same voltage vector, and the decision of which to apply based on some pre-determined criterion, or left to the influence of the loss model.

This technique on its own would result in the same active vector being applied for an extended period - i.e. the 60 electrical degrees for which it is closest to the demand. The modulation index would also only be able to take a value of 0 or 100% for each switching period, and the average modulation index over multiple switching periods would not be considered: i.e. there would be no modulation of voltage.

These shortcomings can be alleviated by filtering the applied vector in order to effectively achieve modulation over a number of switching periods, determined by the length of the filter. In other words, instead of the voltage being modulated within a switching period, modulation takes place over longer period consisting of multiple switching periods. This principle is illustrated in Figure 3.2 where the top signal shows the vector that would be applied in an unfiltered single-state modulation scheme, and the bottom signal shows how filtering might allow modulation.

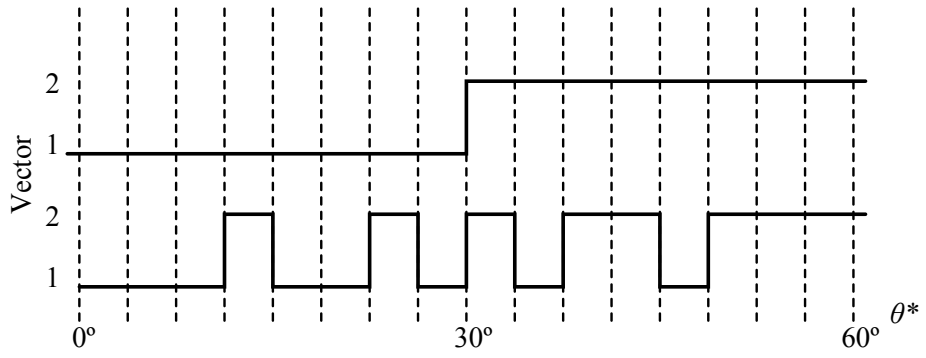

Figure 3.2: Illustrative vector choices for unfiltered (top) and filtered single-state modulation

Even using filtering, this method still gives a lower 'resolution' of control over the same period than conventional SVM patterns: i.e. the number of available voltage vectors is limited rather than being a continuous range as in SVM. This could cause large deviations from the demanded current if not properly controlled. The non-deterministic nature of the resultant switching sequence may also increase harmonic content, though as an additional benefit, the radiated electromagnetic interference (EMI) may be reduced since the effective switching frequency is varied.

3.1.3.1 Filtering

Only six angles, those of the six fundamental non-zero vectors, can actually be applied by the inverter. However, if the angle error is calculated only between these fundamental angles and the demand, the same fundamental vector will be applied for a continuous 60° of electrical field revolution. Unlike conventional SVM where the proportions of two applied non-zero vectors are varied smoothly from period to period, this will result in a poor quality current waveform, with high harmonic content, similar to that found in six-step modulation.

This situation is resolved by filtering the vector to produce a value which represents the applied vector over more than one switching period, resulting in different fundamental vectors being applied within each 60° period. The result is that the voltage more smoothly tracks the demand.

The vector is expressed in polar form at the filter input, and its modulation index and angle are filtered separately. The same structure is used for both components. The following description describes the filtering of the angle, but the filtering of the MI is implemented in the same way.

The conventional finite impulse response (FIR) filter structure, considered the best choice due to its simple digital implementation and finite response, has been adapted to allow predicted outputs for each fundamental state to be produced with minimal computation. The new structure is shown in Figure 3.3 for an nth order filter. An nth order filter has $n+1$ filter coefficients, which determine the frequency response of the filter for a given sample rate. These coefficients sum to 1, ensuring the filter has no net effect on the amplitude of the signal.

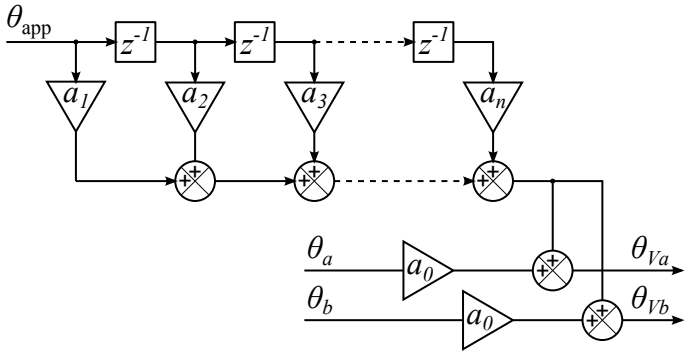

Figure 3.3: Structure of nth order angle filter for single-state modulation

The angle applied at the beginning of the present PWM period θ_{app} is fed into the delay line where it and its $n-1$ previous values are convolved with the filter coefficients a_1-a_n. The first filter coefficient a_0 is multiplied by the angle of each

possible next state θ_a and θ_b and added to the convolution result to form the two effective angles θ_{Va} and θ_{Vb}.

In the case of the MI, the filter input and possible next values will always be either 0 or 100%, depending on whether each is a zero or non-zero vector.

It is possible to design a range of different filters with different frequency and phase responses. FIR filters are designed according to well established principles which can be automated. The order of the filter (the number stages in the delay line) can be varied, as well as the cut-off frequency and windowing method. It was found experimentally that a 7th order moving average filter gives the best response. A moving average filter is a special case of the FIR filter where each of the elements is equally weighted. This gives more weight to the future values than in a classic FIR where the oldest and newest (in this case future) values are multiplied by the lowest coefficient.

3.1.3.2 Reference Frame Phase-shift

The previous section outlines the need to filter the angle of the applied vector to achieve a smooth response. Linear values like the MI do not present a problem, but filtering or averaging angles is more complicated due to their circular nature. This is a known problem in general statistical theory, where it is referred to as the directional or circular mean problem [70, 71].

The voltage vector rotates when expressed relative to the stator, and hence the angle changes. However, the voltage vector produced by the PI controllers is expressed in the e-frame, and is theoretically stationary in the steady-state. Keeping the vector in this reference frame simplifies the implementation of the filter.

For the filter implementation to be successful, the sampled values of the angle must lie within a range of π rad. This is because the definition of the average of two angles greater than π rad apart is unclear, since the closest angle which bisects them lies in the opposite semi-circle. Making the practical assumption that induction machines are operated over a small range of slip values, this does not pose a problem.

The other criterion is that the angle does not cross the $0 \leftrightarrow 2\pi$ rad boundary. The likelihood of this happening is reduced by moving the boundary away from the active range by rotating the reference frame, but still keeping it static in relation to the electric field. In the proposed scheme the e-frame is phase-shifted by a constant π rad, but rotates at the same angular velocity. The e-frame will subsequently be referred to as the e'-frame to indicate this offset. The relationship between the phase-shifted e'-frame, the unmodified e-frame and the s-frame is shown in Figure 3.4, including an arbitrary demand vector.

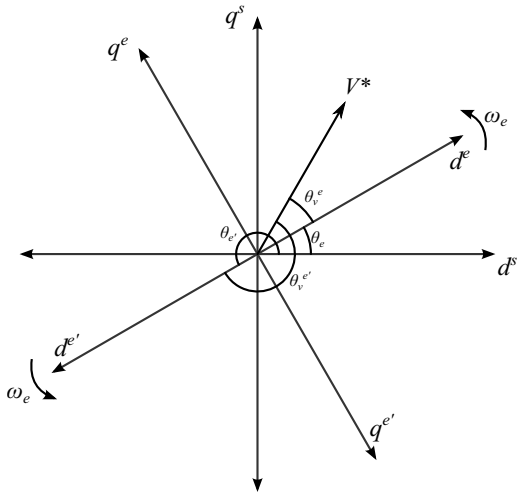

Figure 3.4: Illustration of the phase-shifting of the excitation reference frame

The overall structure of the FOC scheme when the e-frame is phase-shifted is illustrated by Figure 3.5. The simulation results in Section 3.2.3.1 also show how the reference frame is phase-shifted.

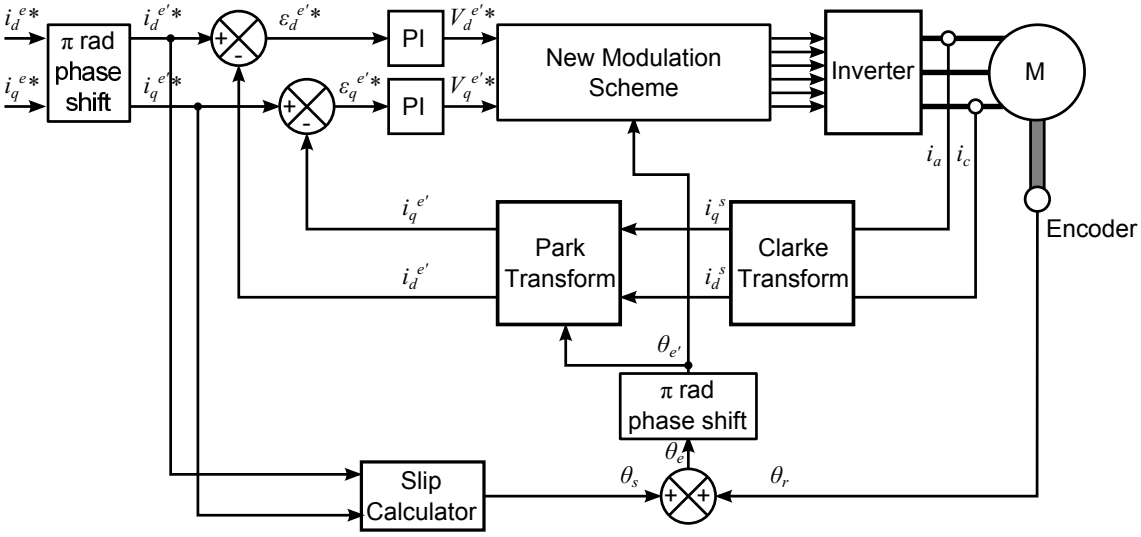

Figure 3.5: The modified FOC scheme with reference frame phase-shifting

3.1.3.3 Cost Function

As previously implied, the control scheme chooses the optimal vector or pattern of vectors based on both the need to recreate the demand accurately and the minimisation of the losses. A cost function is used to allow the relative priorities of these objectives to be controlled. The cost function is evaluated for each possible output and the one that minimises it is chosen.

The general form of the cost function is:

$$g = A \cdot \varepsilon + P_{\text{loss}}. \tag{3.1}$$

The weighting factor A determines the relative priority of the vector error ε and the calculated loss P_{loss}. Increasing A gives more importance to the accuracy of the vector at the expense of the minimisation of loss, and decreasing vice versa.

Both angle and magnitude errors are included in the cost function in single-state modulation, giving:

$$g = A(\varepsilon_m + \varepsilon_\theta) + P_{\text{loss}}. \tag{3.2}$$

A radian is defined as the angle subtended by two radii joined by an arc of the same length. Short arcs are approximately the same length as the corresponding chord. Therefore, a vector which differs in magnitude from the demand vector by a small amount, but has the same angle, is a similar distance from the demand vector to a vector with an angle which differs by the same amount in radians, but has the same magnitude as the demand vector. An illustration of this principle can be found in Figure 3.6, which compares vectors differing by magnitude 0.1 and angle 0.1 rad from an arbitrary vector. The distance to the error vector with a different angle is shown by the cosine rule to be approximately equal to the distance to the error vector with a different magnitude by (3.3).

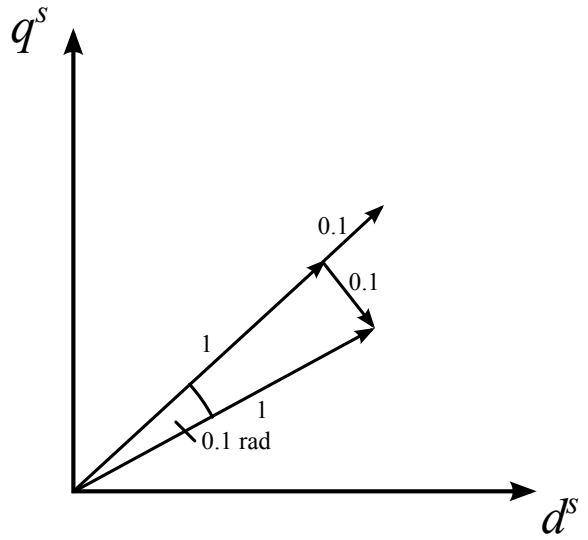

Figure 3.6: Comparison of angle and magnitude errors of vectors

$$\cos\theta = \frac{1^2 + 1^2 - 0.1^2}{2 \times 1 \times 1} \tag{3.3}$$

$$\theta = \cos^{-1} 0.995 = 0.10004 \tag{3.4}$$

This approximation is true when the magnitude of the demand vector is one. For demand vectors with magnitude less than one, the angle error should be weighted, since the same angle error is equivalent to a smaller magnitude error. This can be achieved by multiplying the angle error by the magnitude of the demand vector. This can be proven by considering that the magnitude of the error ε corresponding to a vector with an angle error θ but the same magnitude as the demand m can be expressed using the cosine rule as:

$$\varepsilon = \sqrt{2m^2(1 - \cos\theta)}. \tag{3.5}$$

It can be demonstrated that for $0 \leqslant m \leqslant 1$ and $-0.1 \leqslant \theta \leqslant 0.1$, which covers the range of values practically encountered, that

$$\varepsilon \approx m\theta. \tag{3.6}$$

The errors ε_m and ε_θ are evaluated for each adjacent vector by subtracting the predicted values from the demand:

$$\begin{aligned} \varepsilon_{\theta a} &= \theta_v* - \theta_a \\ \varepsilon_{\theta b} &= \theta_v* - \theta_b \\ \varepsilon_{ma} &= m_v* - m_a \\ \varepsilon_{mb} &= m_v* - m_b. \end{aligned} \tag{3.7}$$

The basic structure of the single-state modulation method is illustrated in Figure 3.7. The point where the calculated losses for each possible state would be introduced is shown. In the first instance, this and the other proposed methods will be considered in isolation from the loss model to assess their performance separately.

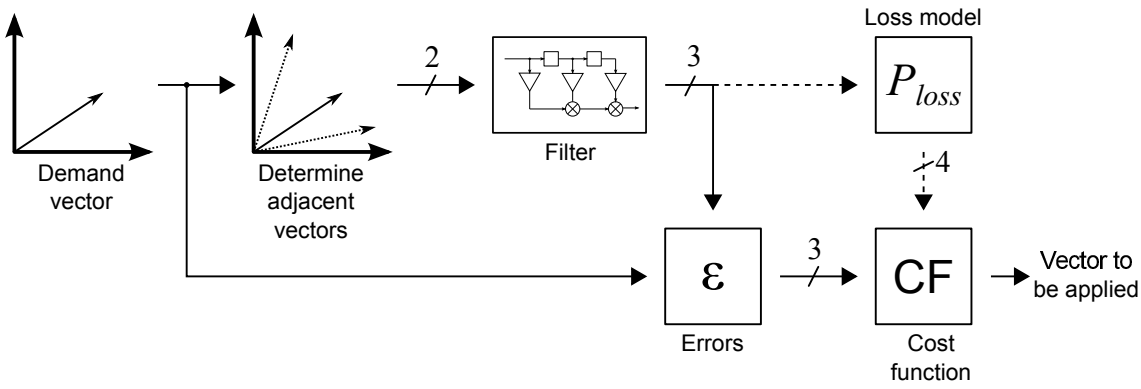

Figure 3.7: Block diagram showing single-state modulation scheme

An example of the generation of the switching signals using this method is shown in Figure 3.8. This shows a modulator being used with saturated reference values. In practice, a modulator is not necessary and the selected switching state could be applied directly, but this demonstration allows for easier comparison with the other proposed methods.

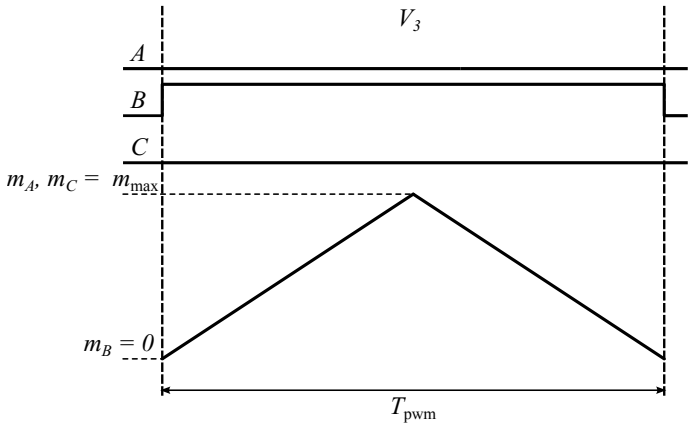

Figure 3.8: Switching waveform derivation in a single-state modulation scheme

This example and the other examples in this section show how the signals might be derived for the same example demand vector as in Section 2.1.1.2 to allow comparison. Note that this is not a singular solution: in this case, the action of the filter may determine that either of the non-zero vectors or the vector V_4 may be the optimal solution.

There is only one degree of freedom in this modulation method - the choice of vector from one of the four available: two adjacent active vectors and the two zero vectors. The zero states are vectorially equivalent, and so the errors between each and the demand vector are equal and are only calculated once.

When the loss model is included, the two zero vectors may give different results and so the system will choose between them on this basis. However, as implemented in this chapter, without the loss model, the two zero vectors always give identical inputs to the cost function, and so arbitration is needed. In this chapter, V_0 is always chosen as the zero vector. This is not ideal as it will result in unequal current sharing between devices on the top and bottom of each leg [12], but the principle of operation is adequately demonstrated.

One possibility for zero vector choice would be to calculate the number of commutations required for each state and choose the minimum. This could be implemented separately as an arbiter when a zero state minimises the cost function, or fed into the cost function. Including the number of commutations in the cost function, where it would also influence the choice of non-zero vectors, would result in a reduction in switching loss.

Even assuming identical devices and constant current during the switching period, the converter losses are dependent on the state because each phase current may be carried by either an IGBT or a diode, leading to different conduction losses. The switching loss is also dependent on the state applied in the previous period.

3.1.4 Cascaded Polar Modulation

It is desirable to minimise the number of states which are evaluated each switching period. However, if each possible state is evaluated only once per period and the optimising state applied for the whole period, the resolution at which the demand vector can be reproduced is poor, as shown by the results of the single-state modulation simulation in Section 3.2.3.

In classic SVM the MI determines the proportion of the PWM switching period occupied by non-zero vectors. It is trivial to determine the MI from the amplitude of the demand vector and use this to calculate the ideal proportion of time to apply the zero vectors, which has no effect on the angle of the applied vector. The angular error and loss model then only need be evaluated for the non-zero vectors, of which only the two adjacent the demand need be considered, and the one that produces the optimal solution is applied in the remaining time. This gives control over the MI to a much higher resolution.

This concept is illustrated in Figure 3.9 with the switching times being calculated as in (3.8). Reference values for use in the modulator are calculated in (3.9) and applied according to the chosen active vector. Phases for which the top switch is on for the applied active vector are given the reference value T_1, and those where the top switch is off are given the reference value T_0. For example, the top switches have states $\{1, 1, 0\}$ for active vector V_2 and so T_1 is applied to phases A and B, and T_0 to phase C.

$$\begin{aligned} T_0 &= m \cdot T_{\text{pwm}} \\ T_1 &= (1 - m) T_{\text{pwm}} \end{aligned} \tag{3.8}$$

$$\begin{aligned} m_0 &= \frac{m_{\max} + m}{2} \\ m_1 &= \frac{m_{\max} - m}{2} \end{aligned} \tag{3.9}$$

A block diagram outlining the whole process is shown in Figure 3.10.

The cascaded polar modulation scheme also uses a filter to improve the response of the system. Since the applied MI is pre-calculated, only the angle is filtered. Some processing of the angles is required before they are input to the filter. The

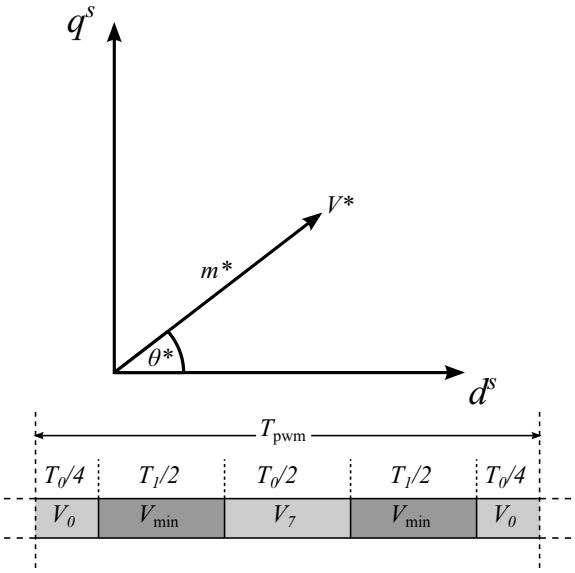

Figure 3.9: Determination of modulation timings under a cascaded polar modulation scheme

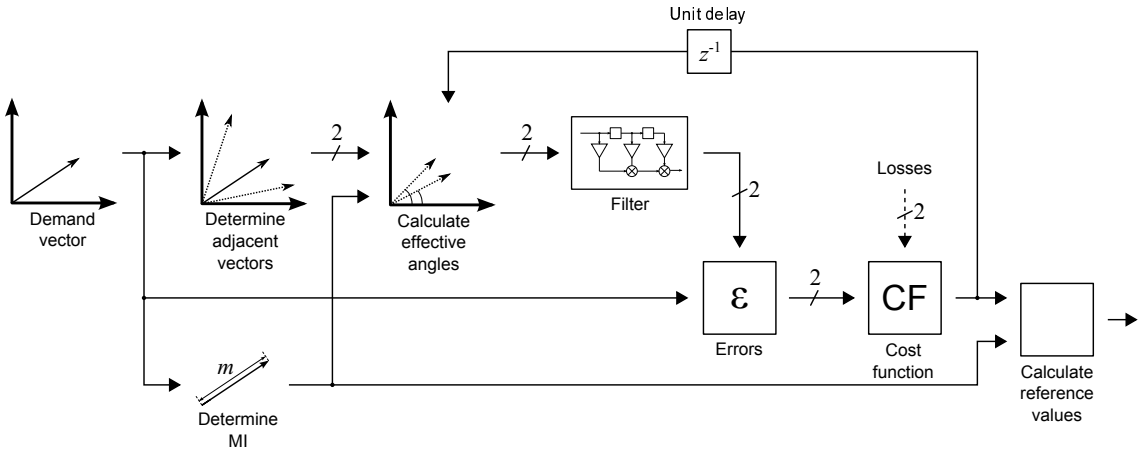

Figure 3.10: Block diagram showing cascaded polar modulation scheme

effective angle of each possible future vector is determined not only by the angle of the non-zero vector, but also the MI, since this determines the period for which the non-zero vector is applied. Thus the effective angles θ'_a and θ'_b are calculated using a mean of the presently applied angle θ_{app} and the possible future angles θ_a and θ_b, weighted according to the MI. Using presently applied angle for the zero vector period T_0 neutralises the effect of the MI on the angle. This calculation is shown in (3.10).

$$\begin{aligned} \theta'_a &= \frac{T_1}{T_{\text{PWM}}} \cdot \theta_a + \frac{T_0}{T_{\text{PWM}}} \cdot \theta_{\text{app}} \\ \theta'_b &= \frac{T_1}{T_{\text{PWM}}} \cdot \theta_b + \frac{T_0}{T_{\text{PWM}}} \cdot \theta_{\text{app}} \end{aligned} \tag{3.10}$$

This modulation scheme also uses the phase-shifted reference frame described in Section 3.1.3.2.

In cascaded polar modulation, only the angle error is included in the cost function, i.e.

$$g = A \cdot \varepsilon_\theta + P_{\text{loss}}. \tag{3.11}$$

There is only a single degree of freedom in this modulation scheme: the choice of active vector from the two adjacent. This will have an effect on conduction loss, since the stator currents may be carried by different numbers of IGBTs and diodes, and switching loss, as different numbers and types of commutation will be required for the two vectors.

3.1.5 Discretised Modulation

This method may be considered a variation on a standard SVM pattern, such as the symmetrical scheme illustrated in Figure 2.7. The modulation values m_a, m_b and m_0 are presented as continuous, but the range of values they can take will usually be limited when implemented digitally. With the digital signal controller used for experimental testing in this project, for example, the range of carrier signal values is limited to 0-7500 for a carrier frequency of 10 kHz. The comparison is computed with integers, and so 7501 is the effective number of different possible lengths of time each state may be applied for.

This enables a finite number of possible switching patterns, albeit a very large number, since it must be raised to the power of the number of phases: $7501^3 = 4.22 \times 10^{11}$. These cannot all be evaluated in a practical system. A more realistic control scheme may be achieved by artificially reducing the resolution of the modulator. The chosen resolution will be a trade-off between control accuracy and computation time.

The number of permutations, and hence computation time, can be reduced by considering the innate constraint shown in (3.12).

$$m_0 \leq m_a \leq m_b \tag{3.12}$$

This reduces the number of permutations to the tetrahedral number of the resolution r (the sum of the first n triangular numbers), for which the number of permutations grows less quickly, as shown in Figure 3.11 and given by (3.13) [72].

$$P_{\text{constrained}} = \frac{1}{6} r(r+1)(r+2) \tag{3.13}$$

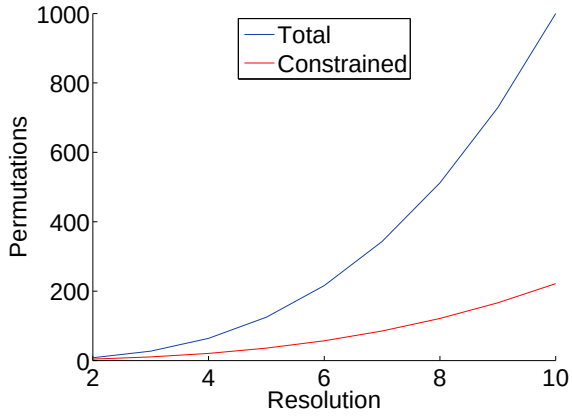

Figure 3.11: Relation of resolution and number of permutations and effects of constraint

The number of permutations may be reduced much further by imposing artificial constraints. By comparing a few possible vectors which belong to the set described by the condition of reduced resolution and which are close to the demand vector, only the options which best reproduce this are considered. The three reference values, one for each phase, are calculated as in the conventional SVM scheme, and discretised according to the chosen resolution. A number of these, this number being referred to as the deviation, which lie either side of the original value are chosen to be evaluated fully.

This results in a number of permutations which is independent of the resolution and dependent only on the deviation. The number of permutations is given by

$$P_{\text{deviation}} = (2D)^3 \tag{3.14}$$

where D is the deviation from the demand value. Whilst this appears similar to the number of permutations in the unconstrained condition, the deviation can be very low and accurate control achieved since the resolution does not affect the number of permutations.

Under these conditions, the choice of resolution becomes a trade-off between limiting the maximum voltage error which will be considered for a given deviation, and maximising the potential reduction in loss by considering operation further from the demand. The choice of deviation is simply a trade-off between potential gains and processing power and is practically limited to very low values—1 or 2—by the cubic relationship with required calculations.

The following is an example of the generation of a set of possible reference values for one phase with a resolution $r = 32$, starting with an arbitrary demand expressed with a resolution of 7500, and a deviation of $D = 2$. The notation $\lfloor x \rfloor$ represents the

floor function, giving the nearest integer less than x; and $\lceil x \rceil$ the ceiling function, the nearest integer greater than x.

$$\begin{aligned} m*_a &= \frac{2578}{7500} \\ m*'_a &= \frac{2578}{7500}(r-1) = 10.66 \\ m'_a &= \{\lfloor m*'_a \rfloor - (D-1), \lfloor m*'_a \rfloor, \lceil m*'_a \rceil, \lceil m*'_a \rceil + (D-1)\} \\ m'_a &= \{9, 10, 11, 12\} \end{aligned} \tag{3.15}$$

Figure 3.12 shows a block diagram of the main processes in the discretised modulation scheme. The signal widths are illustrated on the diagram. The signals between the discretisation stage and the cost function are shown with width of 8, which is valid for a deviation of $D = 1$. For other values of D, the signal will have a width given by (3.14).

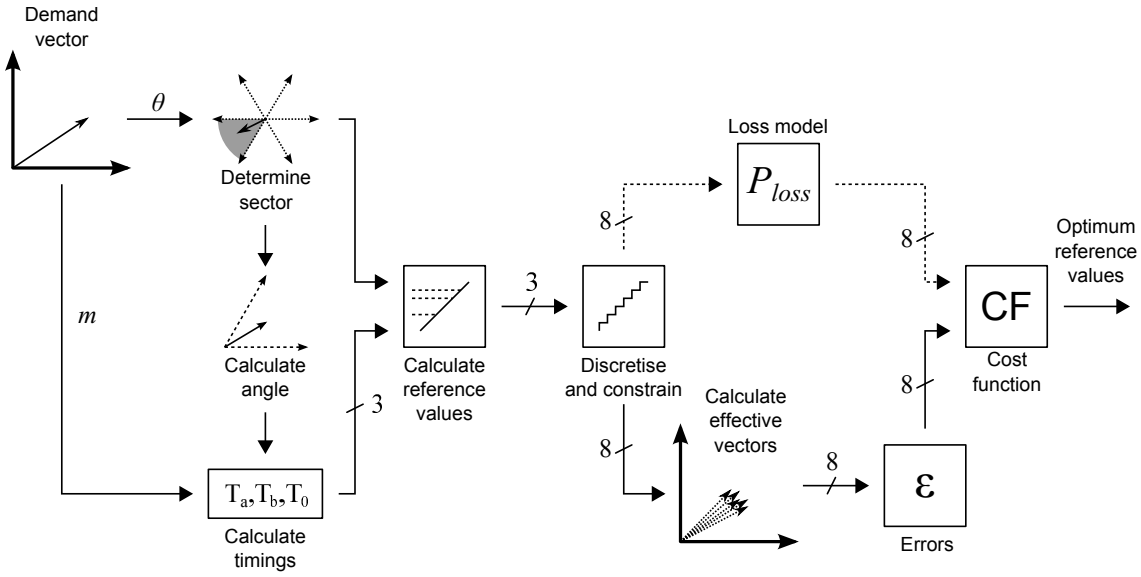

Figure 3.12: Block diagram showing discretised modulation scheme

An example of switching signal generation for a discretised modulation approach with a resolution of 5 is shown in Figure 3.13 to illustrate the principle. The example shows the inherent constraint given in (3.12), but does not show a consideration of constrained deviation.

For the discretised modulation scheme, the equivalent vector must be calculated for each permutation before the errors are calculated for input to the cost function:

$$\begin{aligned} m_{\text{equiv}} &= \frac{T_a + T_b}{T_{\text{pwm}}} \\ \theta_{\text{equiv}} &= \frac{T_a}{T_{\text{pwm}}}\theta_a + \frac{T_b}{T_{\text{pwm}}}\theta_b. \end{aligned} \tag{3.16}$$

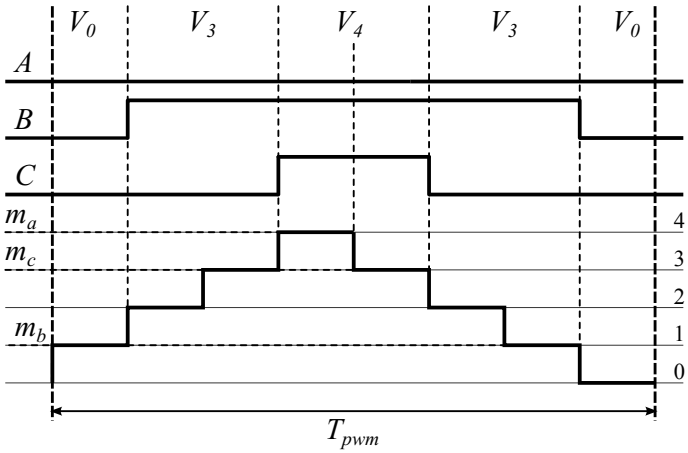

Figure 3.13: Switching waveform derivation in a discretised modulation scheme with resolution of 5

Once the resolution is set by the designer, the controller has one degree of freedom, the reference values, which have been limited to a finite range and constrained as discussed above. The converter losses will be affected accordingly, varying due to the flexibility to pass phase currents through diodes or IGBTs for differing periods. For very low resolutions, the large reductions in switching loss are possible due to the increased potential for discretised reference values equal to each other, the maximum value or zero, and hence vectors to be eliminated from a period completely, reducing switching loss. The example shown above illustrates a situation where the vector V_7 has been omitted in this manner.

For the purpose of comparison, the single-state modulation method described previously could be considered to be equivalent to discretised modulation with a resolution of two; each modulation value has two possible values corresponding to on or off in the corresponding phase, which depends on the sector. Discrete period division with a greater resolution will result in increased control accuracy. Furthermore, the adoption of the symmetrical pattern will reduce harmonic distortion [19].

3.2 Simulation

3.2.1 Implementation

The functionality of the new control scheme is verified by simulating using the MATLAB-Simulink package. This environment was chosen as it allows the mechanical, electrical and control elements of the system to be simulated as one system.

The physical drive system is simulated using a variable step solver to reduce simulation time. The control system is discretised within this and effectively operates as a fixed step system to reflect the nature of its eventual implementation on a digital signal processor (DSP).

The inverter and battery are represented by a simple model which switches terminal voltages between $-0.5V_{dc}$ and $+0.5V_{dc}$. This is a simplification of the real system that assumes zero output impedance of the battery, infinite DC link capacitance, zero voltage drops across the IGBTs and diodes and ideal switching with no deadtime.

The switching frequency is a trade-off between stability of the system and switching loss, both of which increase with frequency [37]. The theoretical lower limit for switching frequency is the Nyquist limit; twice the maximum fundamental frequency. This is determined from the maximum speed and pole number of the machine. However, the practical limit is around 10 times the maximum fundamental frequency. Higher switching frequencies mean more switching events in the same time period, and hence higher switching losses, in direct proportionality. A switching frequency of 10 kHz is chosen for all simulation and practical experimentation as this is a typical value and it is not intended to investigate the effect of switching frequency here.

The IM model is based on the well-known system of equations linking flux, voltage, torque and speed. The terminal voltage is decomposed into an electrical component representing current flowing in the resistive stator winding and a electromagnetic component due to the time-varying magnetic flux. The current can be deduced from the flux and physical constants of inductance and resistance. The principle of conservation of energy is used to calculate power, and hence torque and speed in conjunction with the mechanical equation representing the effects of load and inertia.

The model is described in the s-frame by the following equations:

$$V_{ds} = R_s i_{ds} + \frac{d}{dt}\lambda_{ds} - \omega\lambda_{qs} \tag{3.17}$$

$$V_{qs} = R_s i_{qs} + \frac{d}{dt}\lambda_{qs} + \omega\lambda_{ds} \tag{3.18}$$

$$V'_{dr} = R'_r i'_{dr} + \frac{d}{dt}\lambda'_{dr} - (\omega - \omega_r)\lambda'_{qr} \tag{3.19}$$

$$V'_{qr} = R'_r i'_{qr} + \frac{d}{dt}\lambda'_{qr} + (\omega - \omega_r)\lambda'_{dr} \tag{3.20}$$

$$\lambda_{ds} = (L_{ls} + L_m)i_{ds} + L_m i'_{dr} \tag{3.21}$$

$$\lambda_{qs} = (L_{ls} + L_m)i_{qs} + L_m i'_{qr} \tag{3.22}$$

$$\lambda_{dr} = (L'_{lr} + L_m)i_{dr} + L_m i_{ds} \tag{3.23}$$

$$\lambda_{qr} = (L'_{lr} + L_m)i_{qr} + L_m i_{qs} \tag{3.24}$$

$$T_e = \frac{3}{2}p(\lambda_{ds} i_{qs} - \lambda_{qs} i_{ds}) \tag{3.25}$$

$$T_e = J\frac{2}{p}\frac{d\omega}{dt} + T_L \tag{3.26}$$

The values of the machine parameters used are those derived from the locked rotor and no load tests, as described in Section 4.3.7 and listed in Table 4.1.

The machine is operated at 1500 rpm, giving an electrical frequency of approximately 50 Hz, depending on slip.

The load torque is set to an arbitrary constant value of 10 Nm. A PI controller is used to control the q-axis current demand so that the machine produces torque to match this. The load torque and the actual torque produced by the machine are shown in the traces in each set of results. The torque ripple is then calculated as the difference between maximum and minimum torques produced in each sample, expressed as a percentage of the mean torque.

The spectra are produced by taking a 1 s record sampled at 10 kHz and applying a 16384 point fast Fourier transform (FFT). This results in a resolution bandwidth of approximately 0.6 Hz. The amplitude is plotted on a logarithmic scale.

The total harmonic distortion (THD) of the first 20 harmonics, including the fundamental, of the stator current is calculated using (3.27) [73].

$$\text{THD} = \frac{\sqrt{I_2^2 + I_3^2 + I_4^2 ...}}{I_1} \tag{3.27}$$

The torque ripple and THD calculations for each modulation scheme are shown in Section 3.2.6.

3.2.2 Conventional Modulation

The drive system is simulated with SVM as the conventional modulation system in order to provide a point of comparison. The results are shown in Figure 3.14. The spectrum of the stator current is shown in Figure 3.15.

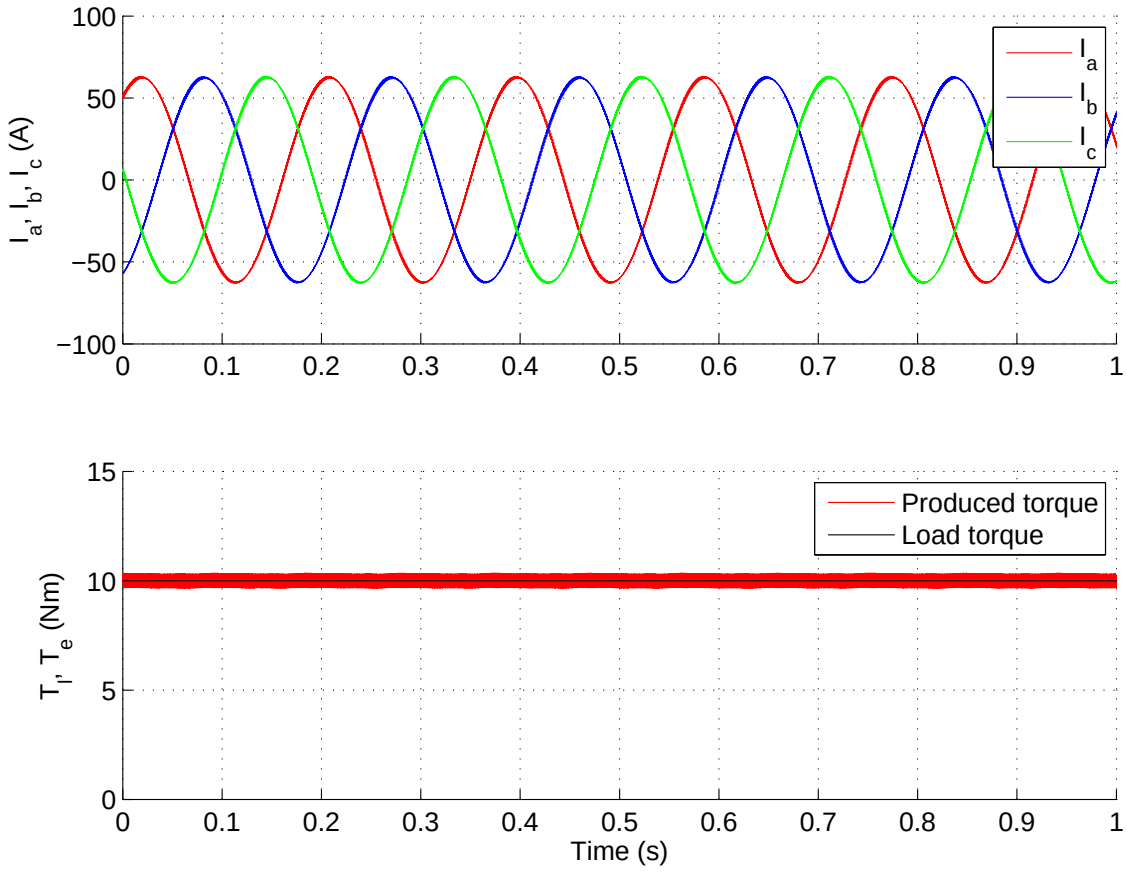

Figure 3.14: Stator currents and torque simulated with conventional modulation

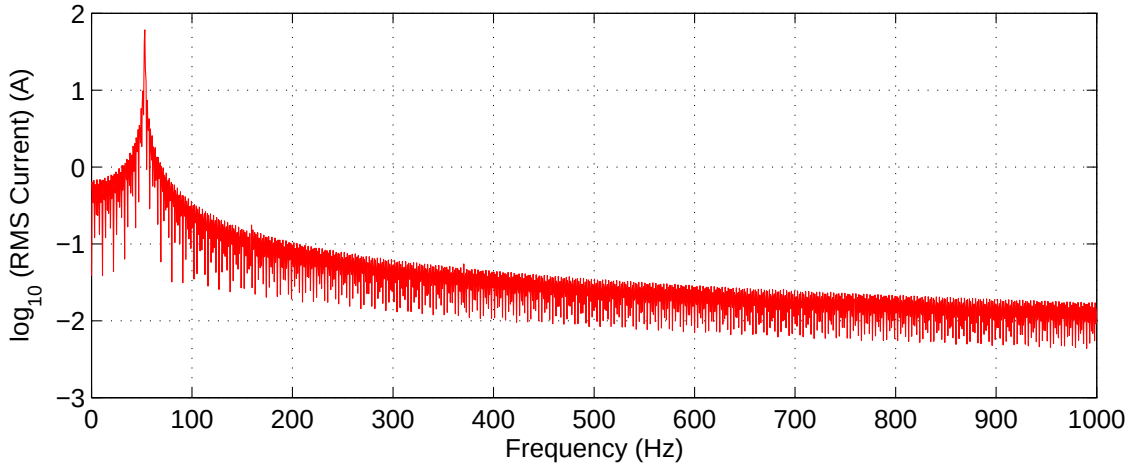

Figure 3.15: Spectrum of stator current in conventional modulation simulation

3.2.3 Single-State Modulation

3.2.3.1 Phase-shifted Reference Frame

Figure A.2 in Appendix A shows the outline of the FOC scheme as simulated, with the reference frame phase-shift as described. The modulation scheme is represented as a subsystem within this, the internal structure of which is shown in Figure A.1. Variables in `monospace` font throughout this section refer to signals shown in these and other figures in Appendix A.

The single-state and cascaded polar modulation schemes proposed use a different reference frame to conventional SVM, as described in Section 3.1.3.2. The relationship between this and other reference frames is demonstrated by the following simulation results.

The various reference frames used are demonstrated by showing the stator current vectors at an arbitrary point during steady-state operation. Figure 3.16a shows the measured stator current in the s-frame. This is calculated from the measured phase currents using (2.8).

The demand current is specified in the e-frame and is shown in Figure 3.16b, with the s-frame axes for reference. In this example, constant demands of 50 A are used for both d- and q-axis current.

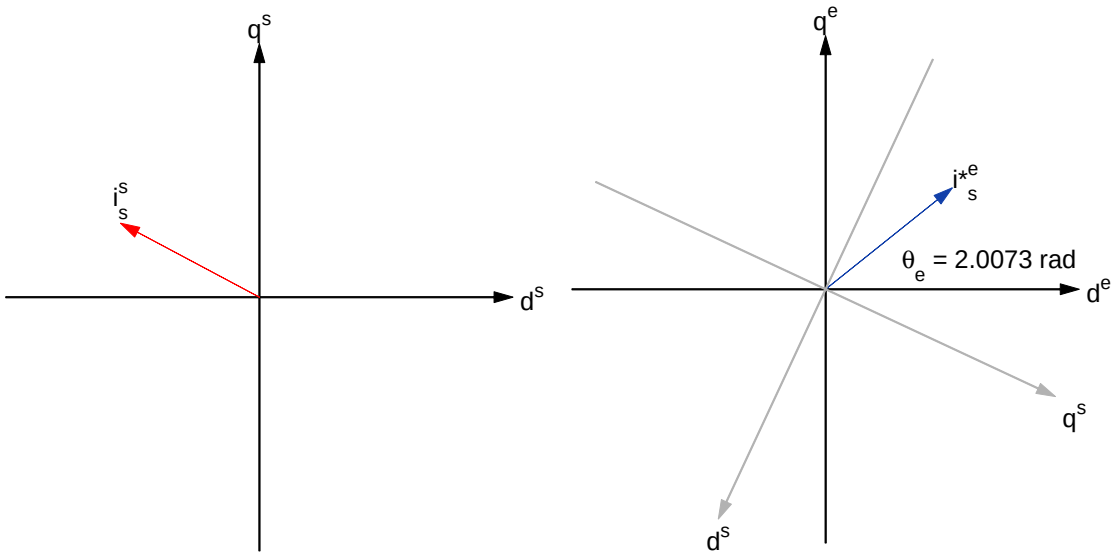

(a) Measured stator current in the s-frame

(b) Demand current in the e-frame

Figure 3.16: Reference frame transformation of current vectors

The demand and measured currents will be compared in the e'-frame. Figure 3.16c shows the two vectors in the e'-frame; again, the s-frame axes are shown for reference.

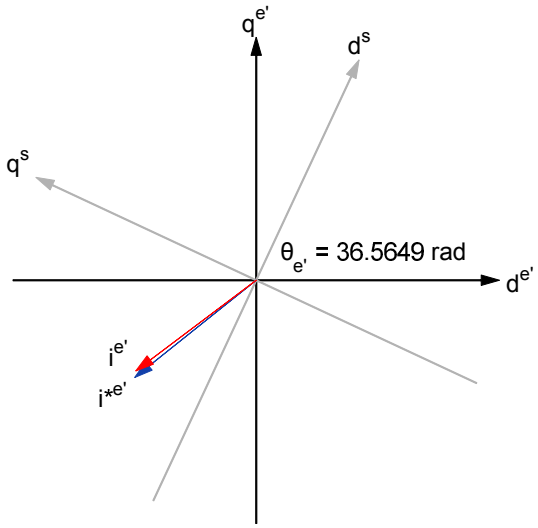

(c) Measured (red) and demand (blue) currents in the e'-frame

Figure 3.16: Reference frame transformation of current vectors (continued)

The other proposed modulation schemes also use the structure seen in Figure A.2 but with different structures within the 'Modulation' subsystem. The discretised modulation scheme (Section 3.2.5) is slightly different because the reference frame is not phase-shifted, but Figure A.2 illustrates its position in the control hierarchy.

3.2.3.2 Fundamental Vectors and Adjacent Vector Selection

The angles of the fundamental vectors in the s-frame are specified as constants, and the vectors are then rotated into the e'-frame by adding `theta_e'`. The adjacent vectors are selected by subtracting each of the fundamental vector angles `theta(1-6)_e'` from the angle of the voltage demand vector `theta_v*_e'` and finding the smallest positive value. This becomes the angle of adjacent vector A `theta_a_e'` and the next vector clockwise becomes the adjacent vector B with angle `theta_b_e'`. The fundamental vectors are shown in the e'-frame for an arbitrary value of `theta_e'` in Figure 3.17 with adjacent vectors A in red and B in blue, and demand vector `V*` in green.

3.2.3.3 Filtering

The vectors are filtered as described in Section 3.1.3.1. The filtering of the MI is simplified since the MI of the fundamental vectors is always unity. For the same reason, the filtered output for the possible next states A and B are the same, and so there are two outputs representing the choice of a zero or non-zero vector in the

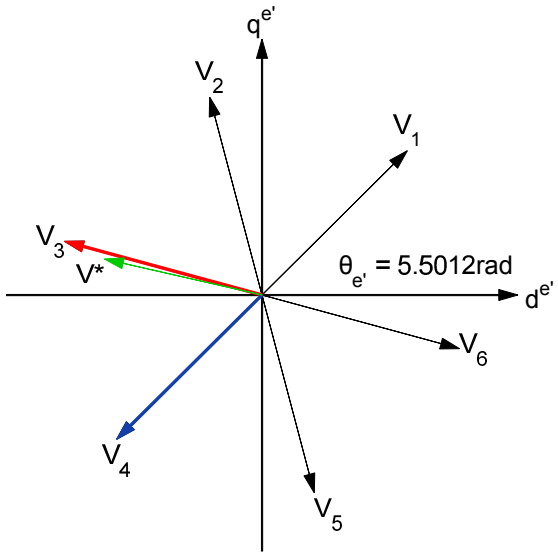

Figure 3.17: Fundamental vectors in the e'-frame for an arbitrary value of `theta_e'` with adjacent vectors highlighted

next switching period. The input and output of the MI filter during a simulation is shown in Figure 3.18. The MI of the vector applied in the previous switching period `m_app` is shown in grey, and takes a value of either 0 or 100%. The filtered MI for a non-zero and zero next state are shown in dark red and green respectively.

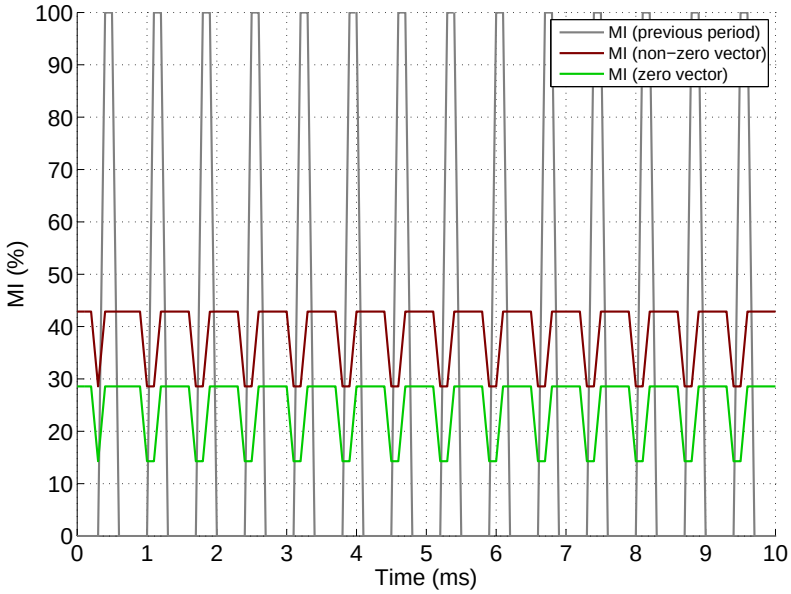

Figure 3.18: MI filter input and output signals in single-state modulation scheme simulation

Corresponding signals for the simulation of the angle filter are shown in Figure 3.19. The angle of the applied vector in the previous period `theta_app_e'`, equal to either one of the fundamental vector angles `theta_a_e'` or `theta_b_e'`, or the demand vector angle `theta_v*_e'` (in the case of a zero vector), is shown in grey. The red and blue traces show the input angles `theta_a_e'` and `theta_b_e'` respectively which

appear as reverse sawtooth waveforms with a period of $\frac{\pi}{3}$ rad due to the rotation of the e'-frame and movement of the demand vector through each of the six sectors. The dark red, dark blue and green lines show the predictive filtered output for next states A, B and 0 respectively.

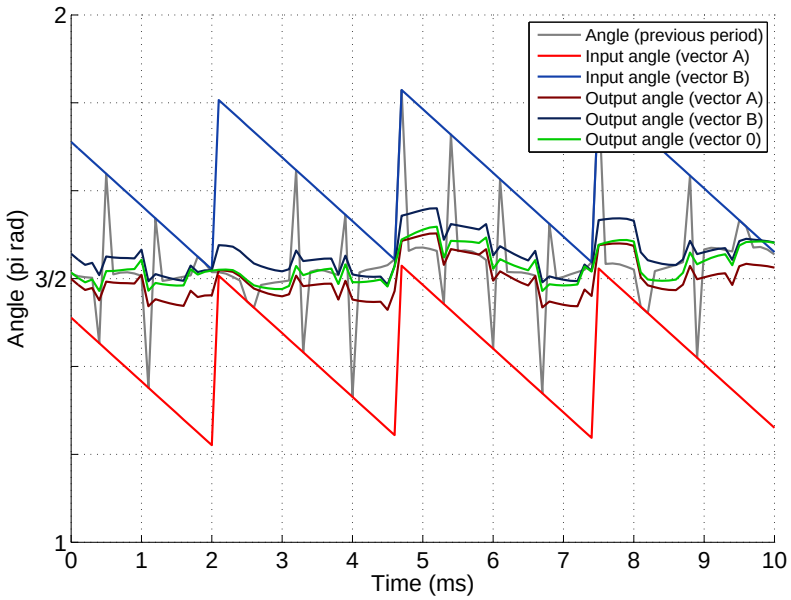

Figure 3.19: Angle filter input and output signals in single-state modulation scheme simulation

3.2.3.4 Error Calculation and Cost Function

The MI errors are calculated as the absolute values of the difference between the demand MI and the filtered MI for each possible next state A, B and 0, `e_0ab` (a vector signal with width 3). The angle errors are calculated in a similar way. The MI and angle errors are summed to give the error for each possible vector.

These values are passed to the cost function where the lowest error value determines the vector to be applied.

3.2.3.5 Results

The single-state modulation scheme is simulated with the conditions described above and 7th order moving average filters.

The sinusoidal current waveforms in the upper part of Figure 3.20 demonstrate that this modulation scheme is capable of producing a rotating electrical field. However, distortion is clearly visible on the waveforms. The effect of this on the torque produced is shown in the red trace in the lower part of the figure. The black trace shows the load torque.

The corresponding spectrum of the phase current is shown in Figure 3.21.

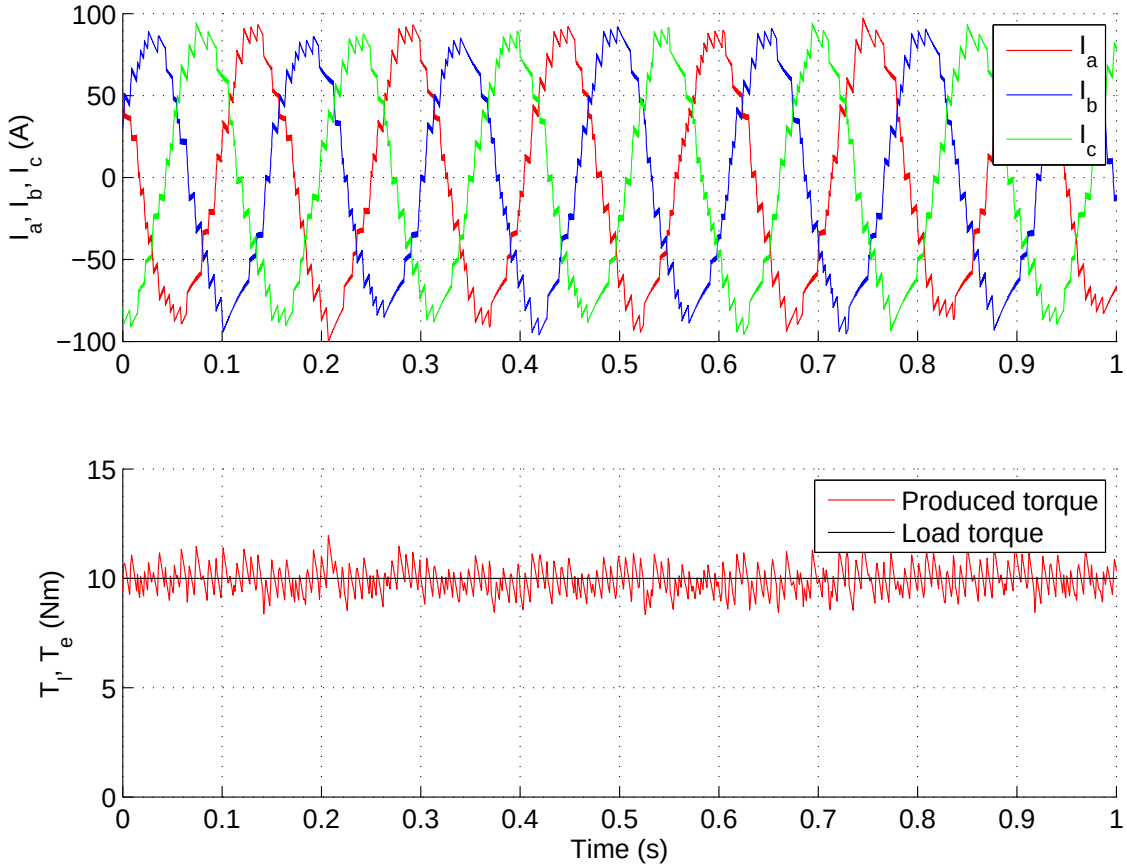

Figure 3.20: Stator currents and torque simulated with single-state modulation

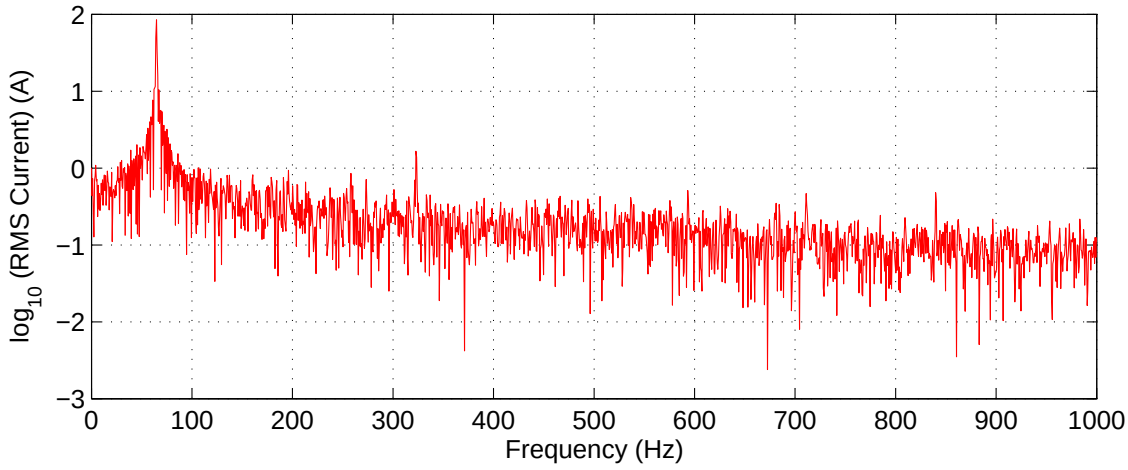

Figure 3.21: Spectrum of stator current in single-state modulation simulation

3.2.4 Cascaded Polar Modulation

The cascaded polar modulation scheme is also implemented with a phase-shifted reference frame, as described in Section 3.2.3.1 and so the FOC structure is the same as shown in Figure A.2.

An overview of this modulation scheme is shown in Figure A.3.

3.2.4.1 Angle Filtering

The values `theta_a_e'` and `theta_b_e'` are passed to the input of the filter. The effective angle of each taking into account the demanded MI is calculated, as in (3.8). The effective angle of the vector applied in the previous PWM period `theta_app` is calculated in the same way, using unit delays, as can be seen in Figure A.3 (blocks 'Select Applied Vector' and 'Calculate Applied Angle').

The filter is then implemented as described in Section 3.1.3.1 and shown in Figure 3.22 (cf. Figure 3.3). Again, a 7th order moving average filter is used.

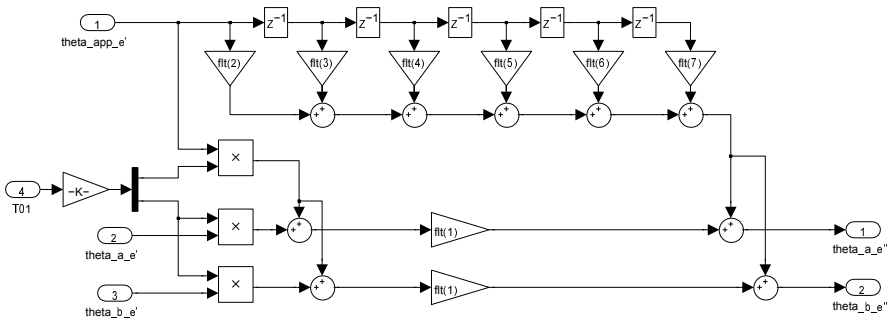

Figure 3.22: Filter structure in Simulink environment

The action of the filter block can be seen in Figure 3.23. The blue and red traces show the values of `theta_a_e'` and `theta_b_e'`, the inputs to the filter, respectively The grey trace shows the effective angle applied in the previous period `theta_app_e'`, and the dark blue and dark red traces the filtered values of the two angles `theta_a_e''` and `theta_b_e''`, the minimum of which will be applied in the next period.

3.2.4.2 Reference Value Formulation

The minimum of `ea` and `eb` is found and this determines which of the vectors A and B is applied in the next period. The minimising vector and the calculated MI are used to produce reference values for the modulator.

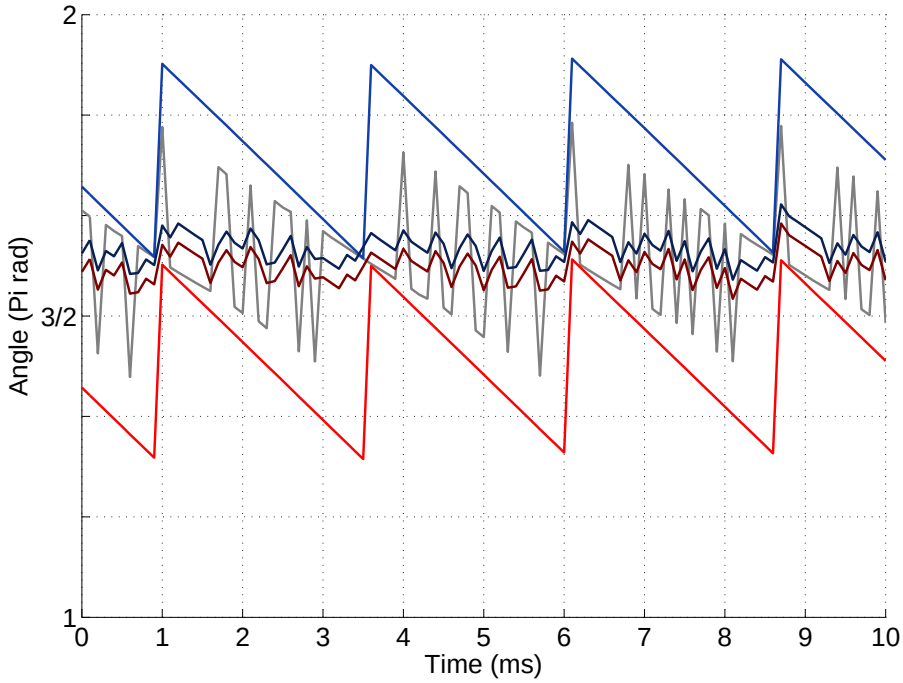

Figure 3.23: Angle filter input and output signals in cascaded modulation scheme simulation

3.2.4.3 Results

The resulting stator currents, and the torque demanded and produced are shown in Figure 3.24. Figure 3.25 shows the spectrum of the stator current.

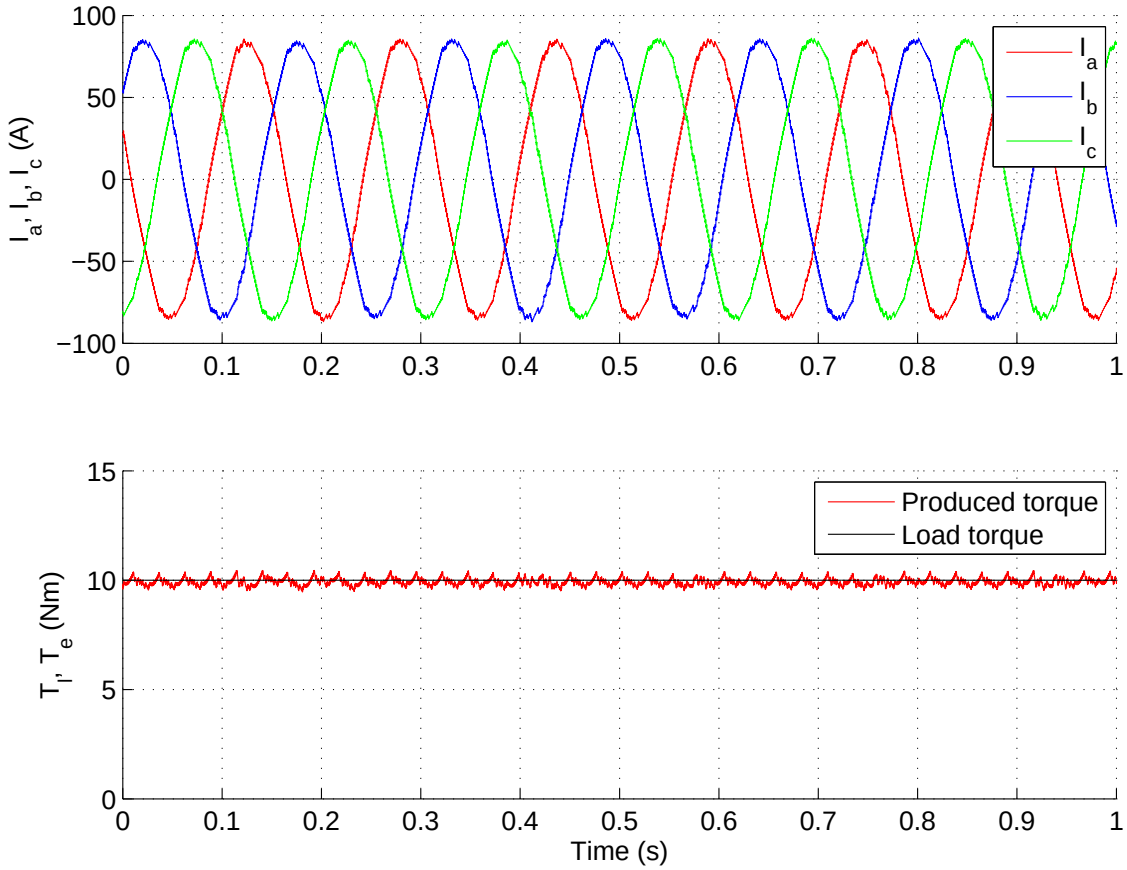

Figure 3.24: Stator currents and torque simulated with cascaded polar modulation

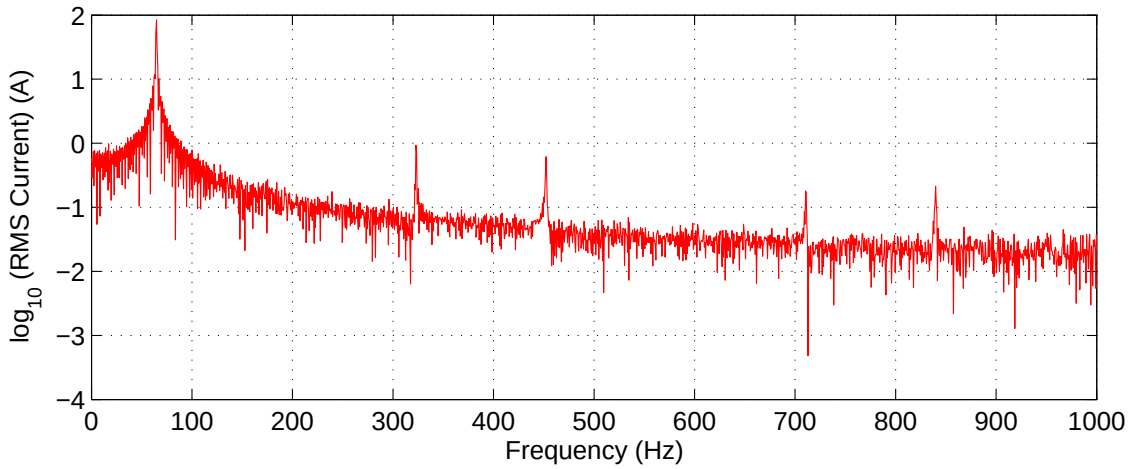

Figure 3.25: Spectrum of stator current in cascaded polar modulation simulation

3.2.5 Discretised Modulation

The discretised modulation scheme simulation diagram is shown in Figure A.4. This is unlike the other two proposed modulation schemes since it operates in the conventional e-frame. In the first part of the block diagram, the modulator reference values are calculated from the demand vector in the same way as in the conventional scheme. After the discretisation and peturbation of these values, the inverse calculations are performed to find the equivalent times and hence vectors, as described in Section 3.1.5.

The equivalent time calculations may give negative results in some circumstances, which is a natural result of peturbing the operating point when these values are close to zero, but is obviously not possible to physically implement. These permutations are removed from the possible vectors by setting the `sat` flag for that permutation in the simulation, as seen in Figure A.4. This causes the corresponding error value to saturate before it is input to the cost function and ensures undefined behaviour does not occur.

3.2.5.1 Results

The discretised modulation method is simulated with high, medium and low resolutions of 7500, 750 and 75 respectively. The high resolution test represents the same effective resolution as the conventional control strategy implemented on the test platform described with a switching frequency of 10 kHz due to the digital nature of the system.

The stator currents and torque for the high, medium and low resolution simulations are shown in Figure 3.26, Figure 3.28 and Figure 3.30 respectively. The stator current spectra for the different simulations are shown in Figure 3.27, Figure 3.28 and Figure 3.31. Stator currents and torque are shown for the discretised modulation scheme low resolution with the same speed but negative torque, representing a regenerative condition where power flows in the opposite direction.

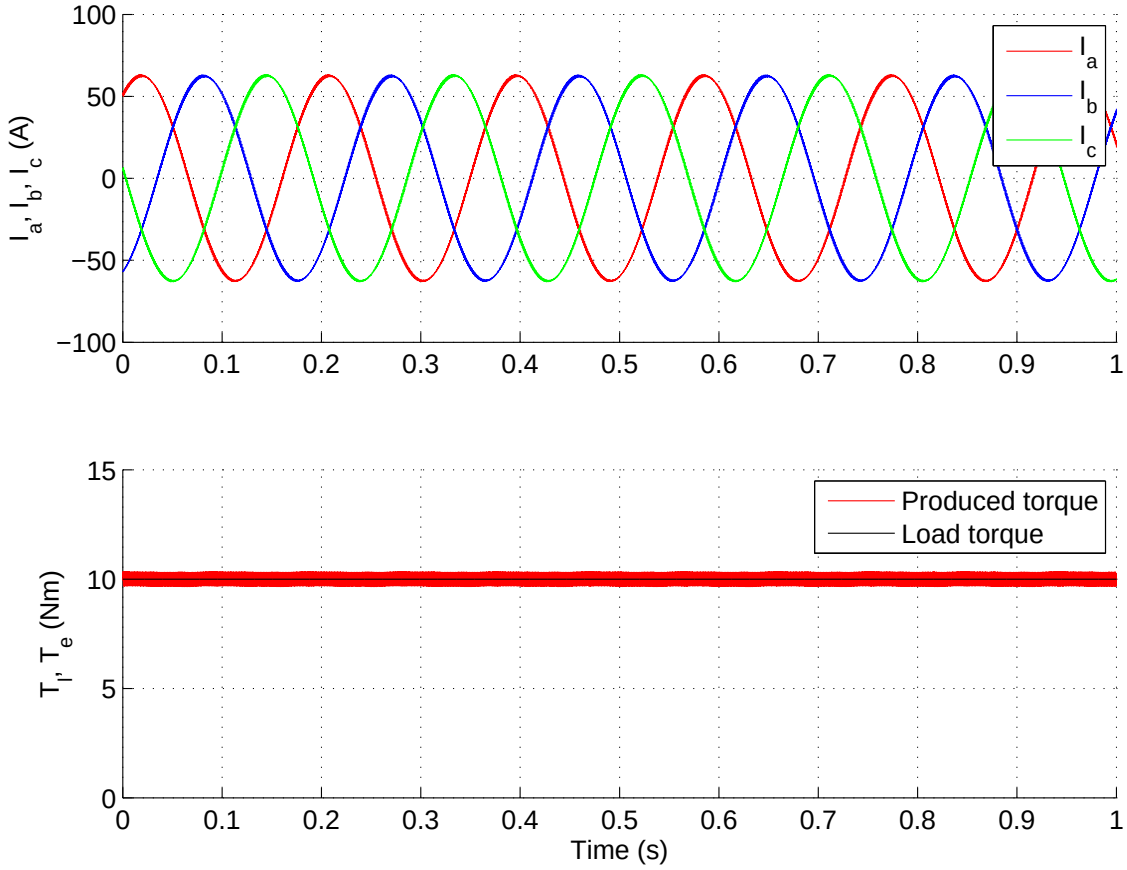

Figure 3.26: Stator currents and torque simulated with high resolution (7500) discretised modulation

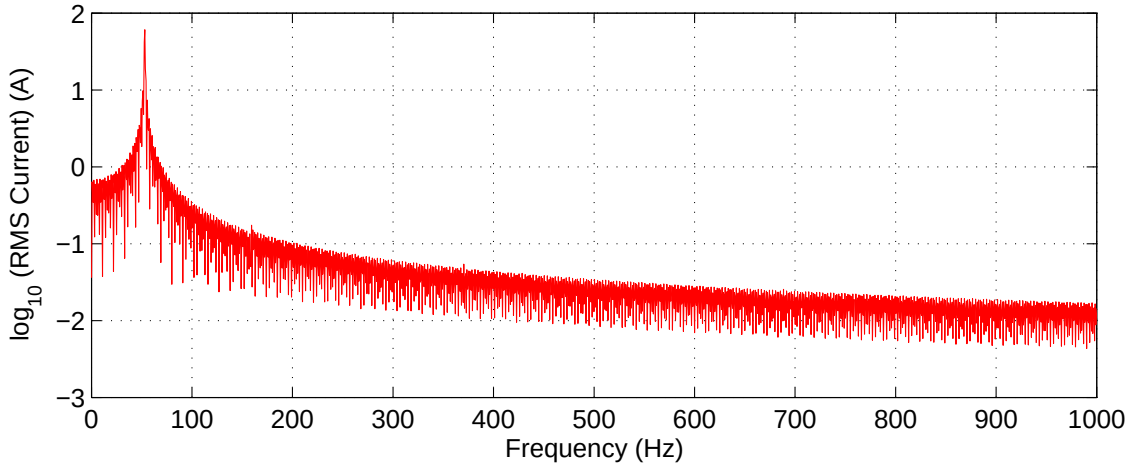

Figure 3.27: Spectrum of stator current in high resolution (7500) discretised modulation simulation

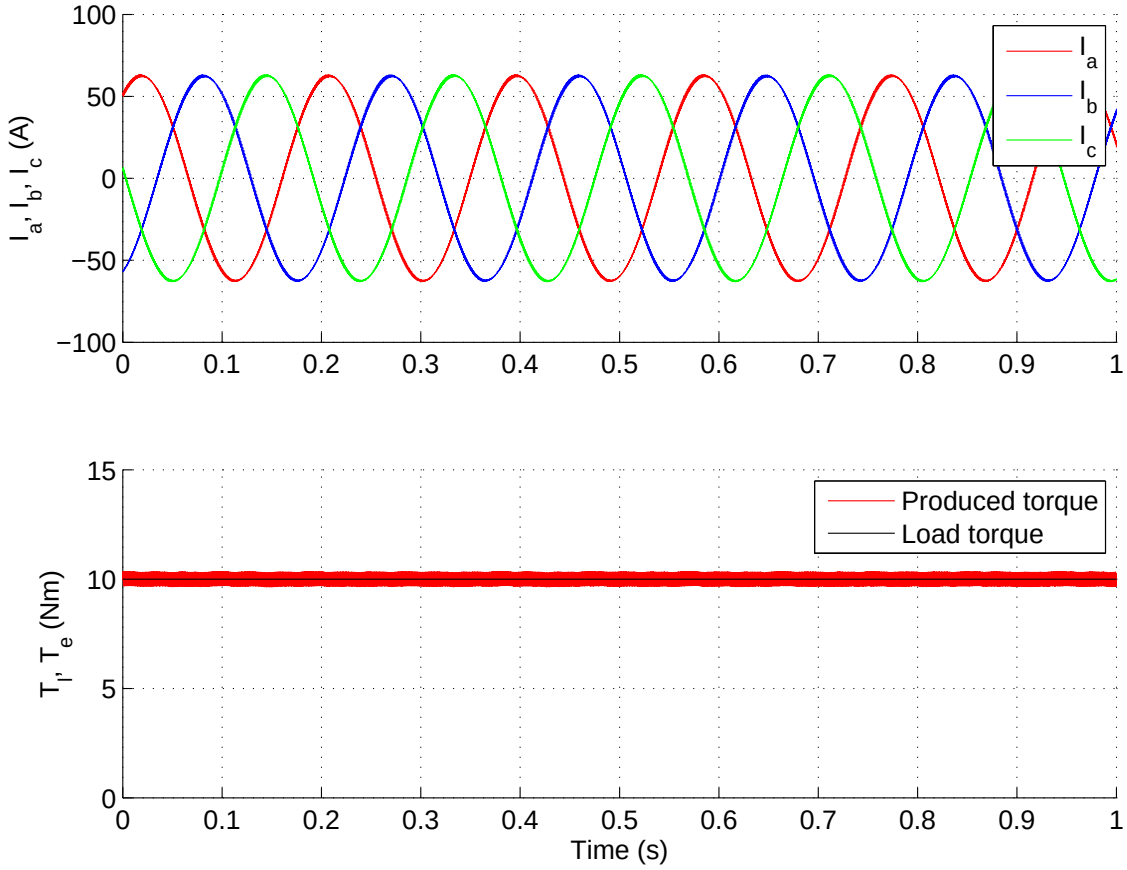

Figure 3.28: Stator currents and torque simulated with medium resolution (750) discretised modulation

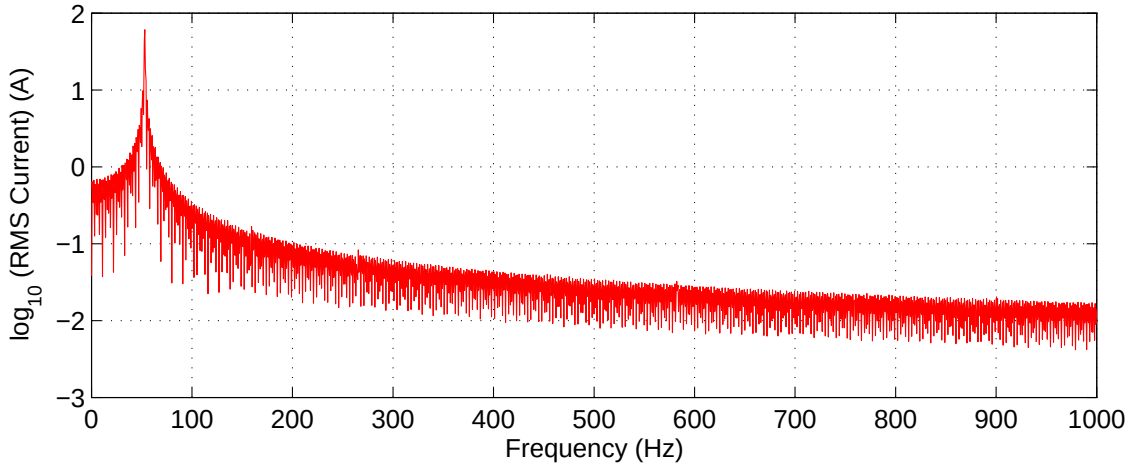

Figure 3.29: Spectrum of stator current in medium resolution (750) discretised modulation simulation

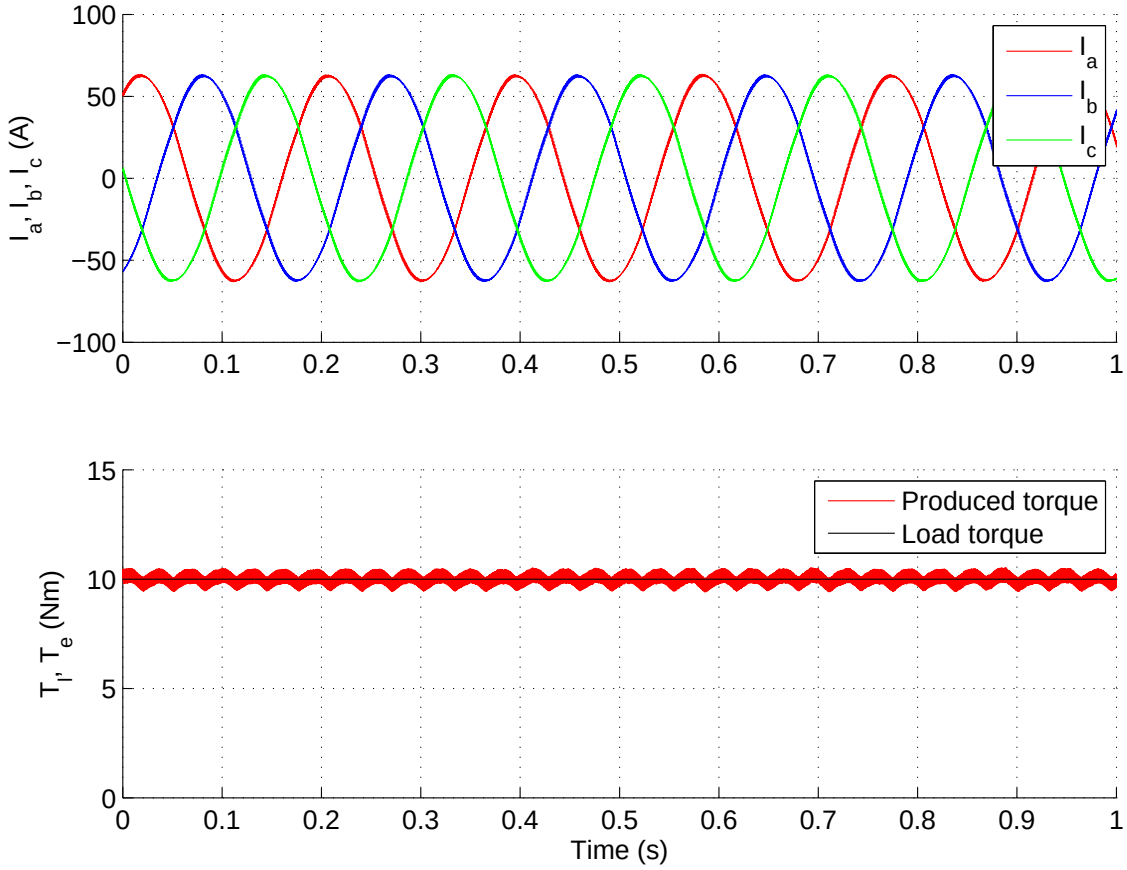

Figure 3.30: Stator currents and torque simulated with low resolution (75) discretised modulation

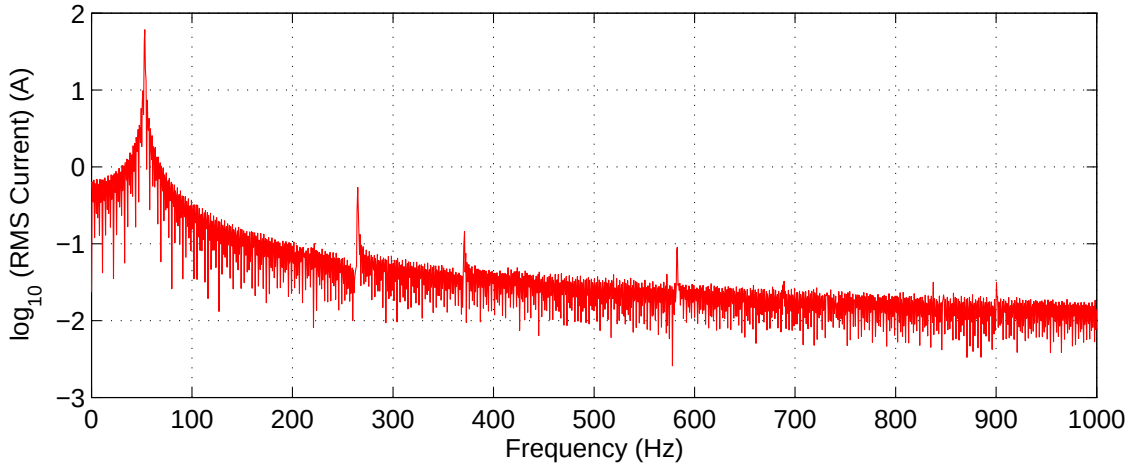

Figure 3.31: Spectrum of stator current in low resolution (75) discretised modulation simulation

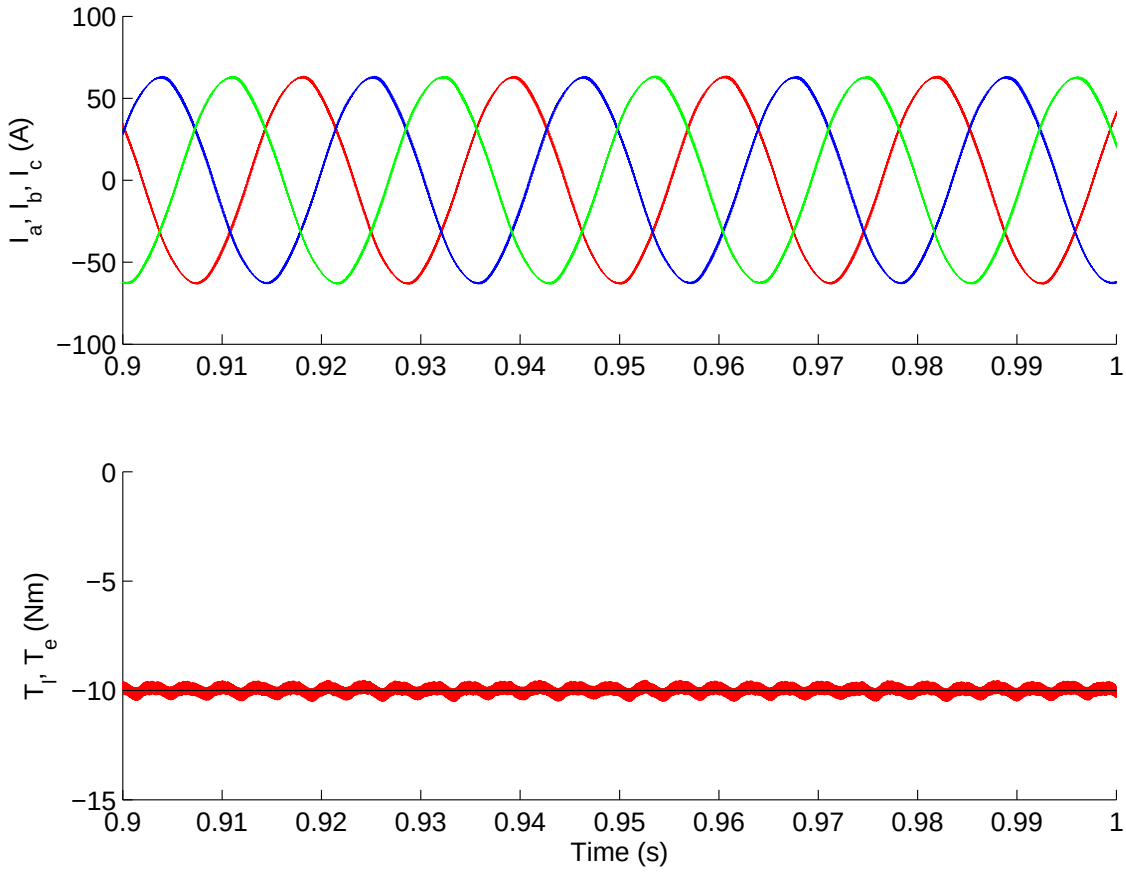

Figure 3.32: Stator currents and torque simulated with low resolution (75) discretised modulation under regeneration

3.2.6 Comparison

The phase current produced by the single-state modulation scheme is heavily distorted, and this is reflected in both the higher torque ripple and higher THD, which are shown in Table 3.1. The torque ripple increases noise, vibration and harshness which would compromise the comfort of a vehicle [74].

The torque ripple produced by the higher frequency harmonics is small and will be damped by the drivetrain before it reaches the wheels, but these harmonics result in eddy currents which increase iron loss. If this effect is too large, it will cancel out any gains in efficiency achieved by the loss model control scheme.

It is clear from the current and torque waveforms that the cascaded polar modulation scheme improves on the single-state method. However, a comparison of the calculated metrics reveals that this scheme still creates greater harmonic distortion than conventional SVM.

It is possible that improvements may be possible in the filters in both these schemes that would improve their performance such that they become viable, but this would require a much more detailed analysis of the control system response than has been made here.

Discretised modulation appears to offer the greatest performance of the all the novel schemes being considered here. Even with a low resolution, both torque ripple and THD are improved compared with the cascaded polar method, though the difference is small. The high and medium resolution results are within two significant figures of those for the conventional scheme, with the high resolution variant even appearing to improve on the conventional.

It is important to note that the aim of the work in this chapter is not to improve on the conventional modulation scheme, merely to change the structure of it to allow the loss model to be integrated, but without degrading the performance.

Modulation scheme	Torque ripple (%)	THD (%)
Conventional	6.9343	0.1110
Single-state	41.8232	3.6795
Cascaded Polar	11.0832	1.5835
Discretised, high resolution	6.9089	0.1098
Discretised, medium resolution	6.9548	0.1155
Discretised, low resolution	11.0243	0.9717

Table 3.1: Calculated torque ripple and THD in simulation

Chapter 4

Test Platform, Benchmarking and Verification of New Modulation Techniques

This chapter describes the implementation of a platform for practical testing of the control techniques proposed in this work and the "benchmarking" of this system for comparison with the novel scheme.

A novel theory of control is the principal aim of this work, but this is worthless without results from a thorough testing process which demonstrate its feasibility and efficacy in practice. This testing process should reflect as closely as possible the intended application, and the measures taken to achieve this—in this case the attempted replication of an EV system in a laboratory setting—are detailed in the relevant sections of this chapter.

In order to provide meaning to the results of tests on the novel control scheme, comparable test results must be available for a conventional scheme to demonstrate any differences in efficiency or other metrics. A description of the implementation of the conventional control scheme is presented along with procedures taken to undertake efficiency testing to a sufficiently high level of accuracy within the practical constraints discussed. The results of this testing are presented, ultimately resulting in an efficiency map of the conventional system covering a defined operating area, which can be compared with results in the same format for the novel scheme.

4.1 Rationale for Practical Testing

Simulation is a valuable tool in many aspects of engineering design. However, the nature of the theory being developed here limits the value of simulation. The loss model developed is essentially a series of calculations giving the total power loss of the system under given operating conditions. To verify the accuracy of this, it is necessary to compare the results to accurate, reliable results obtained under the same conditions. If these reference results are obtained by simulation, the comparison is merely between two different methods of calculation. Unless the reference method is extremely accurate, the validity of the comparison becomes questionable, particularly as the accuracy of the proposed system approaches that of the reference simulation.

A suitably accurate reference simulation would require detailed magnetic and thermal modelling. An alternative method is to compare the proposed loss model to measurements from a real system. The accuracy of the measurement equipment is also a limiting factor for the validity of this comparison. The major advantage of this approach is its inherent and evident real-world accuracy, and so despite the limitations identified, it is felt that an emphasis on practical results is the most robust approach.

4.2 Rationale for Full-Scale Testing

A common method of testing a control scheme such as the one proposed in this work would be to test it using smaller scale devices than would usually be found in the real application. In this case, instead of testing using a 40 kW nominal machine, a much less powerful machine—of say 1 kW—would be used. The drive and other components of the test setup would be rated appropriately to this, including a smaller dynamometer with a less powerful drive and DC supply with lower voltage and current ratings. Auxiliary systems such as cooling also have much lower requirements. This has obvious practical advantages in setup time and cost. The implication is that the system could be scaled up to suit the real application and the effects demonstrated would be similar and proportionate.

Since the overall goal of this work is to improve the efficiency of induction motor drives, accurate power measurements are a key requirement of testing. Unless a calorimeter is used, efficiency must be determined by measuring both the input and output power. In a very efficient system these are both large with respect to the calculated power loss, which exaggerates the inaccuracy of the measurement. Therefore the input and output power must be measured very accurately. This is easier if these have large absolute values.

4.3 Test Platform

A platform for testing the proposed novel control scheme was implemented according to the schematic in Figure 4.1.

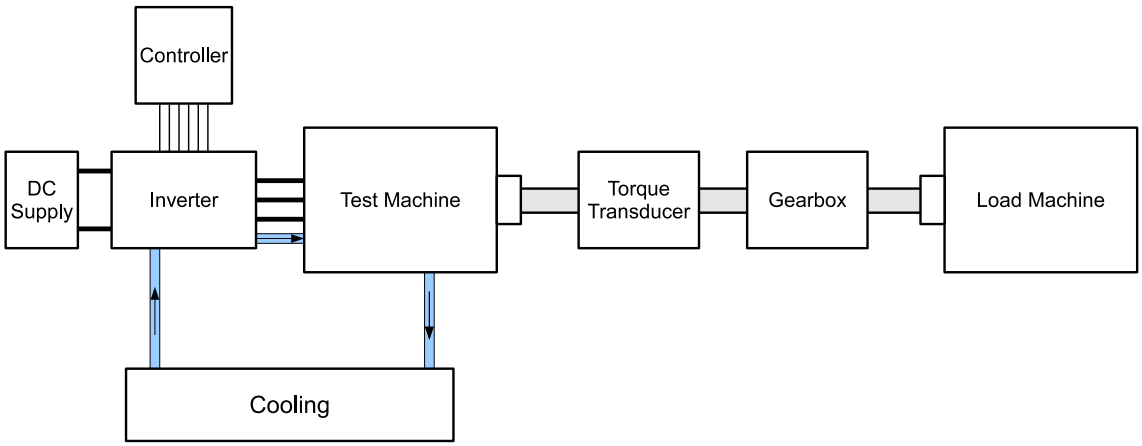

Figure 4.1: Schematic diagram of the test platform

4.3.1 Motor

An induction machine was chosen for reasons previously discussed. The model selected was a 200-330W model was from Swiss manufacturer MES-DEA. This is a 3-phase 4-pole 160 V star-connected induction machine with a nominal power output of 40 kW, making it suitable for a small passenger BEV or a larger mid-strength HEV. The machine is designed for an EV application and so is water-cooled, as it would be in an automotive environment.

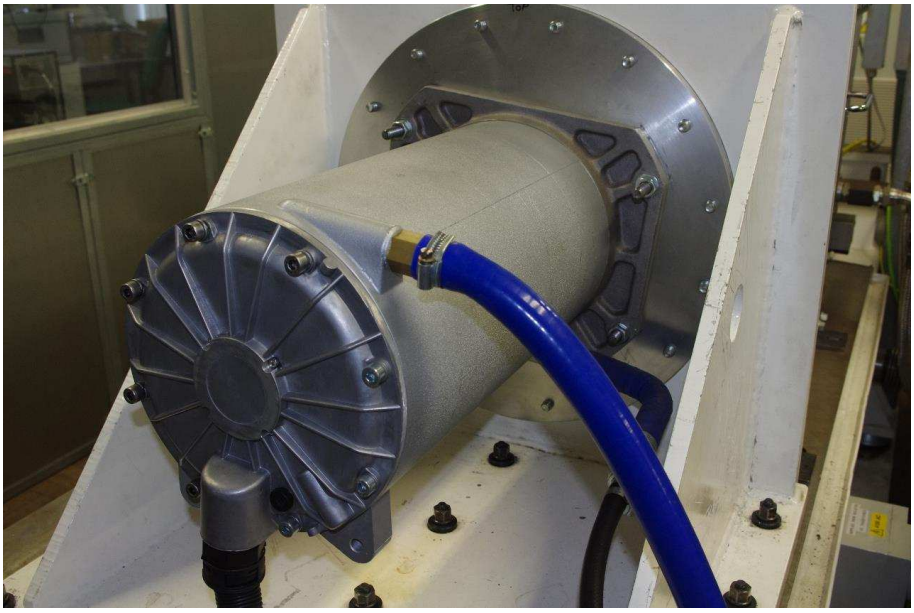

Figure 4.2: Photograph showing induction motor used for practical testing

4.3.2 Inverter

A Semikron SKAI module was chosen as the basis for the drive. This contains six 1 200 V IGBTs arranged in a standard 6-switch inverter topology, with antiparallel diodes, as shown in the schematic in Figure 2.1. Gate drives, a 1 mF film DC link capacitor, current sensors on two output phases, a temperature sensor on the heatsink and a voltage sensor on the DC bus are also included in the package, which is liquid-cooled. An onboard controller provides analogue outputs from these sensors, digital error signals and an interface for each gate drive.

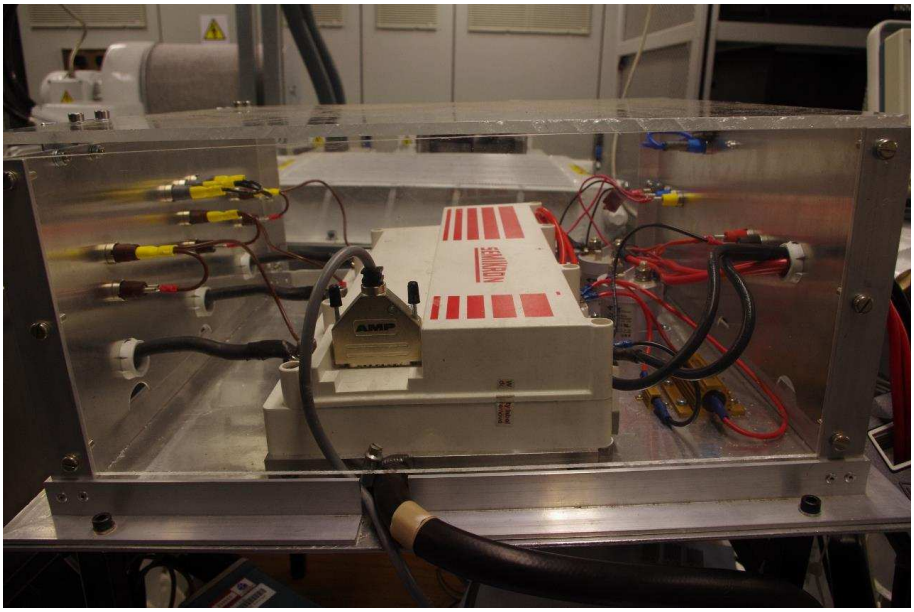

Figure 4.3: Photograph showing inverter used for practical testing

Some protection is also provided by this controller. In the event of overvoltage on the DC bus, overcurrent on any output phase, overtemperature of the heatsink or undervoltage of the supply for the controller itself, all the IGBTs are turned off regardless of control signals and error signals are generated.

The inverter switching frequency used is 10 kHz, for reasons discussed in Section 3.2.1.

4.3.3 Controller

The control scheme runs on a DSP which ultimately generates signals to the inverter module to control switching of the IGBTs. This part of the system is based on a Texas Instruments TMS320 F28335 DSP. This is a 16-bit floating point chip which runs at a clock speed of 150 MHz. It has built-in hardware for PWM generation,

analogue-to-digital conversion (ADC), quadrature encoder interfacing and a serial communications interface.

The DSP is mounted on a commercial interface board from Spectrum Digital, which provides a JTAG controller for programming, RS232 line drivers for the DSPs serial communications interface, additional memory and a 3.3 V power supply. Various input/output (I/O) pins are brought out to connectors on this board.

A further interface board improves access to the features on the DSP. Each PWM channel is brought out through MOSFET drivers which enable these to drive greater loads and protect the DSP, with signals transmitted at 15 V for greater noise immunity.

The controller boards, interface board and logic power supply are shown in Figure 4.4. Note that this system was designed as an experimental platform, and that many redundant components could be removed and the layout condensed for a implementation in a real production setting.

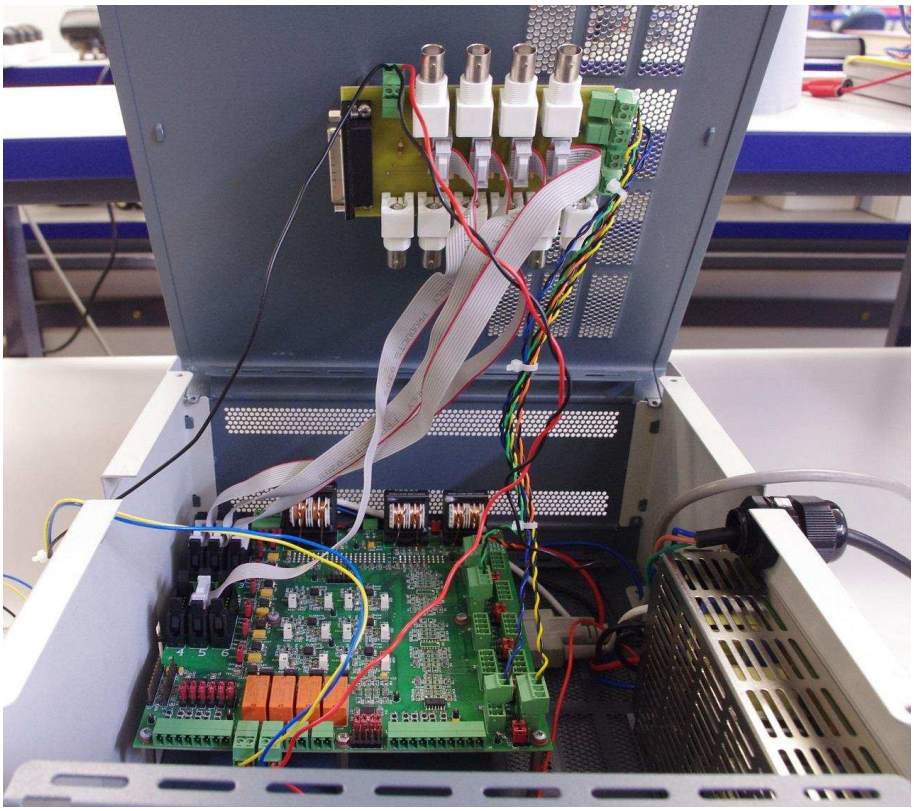

Figure 4.4: Photograph showing control hardware in test platform

The ADC channels are brought out to their own connectors, each having two clamping diodes to prevent the DSP pins being driven below -0.3 V or above 3.3 V. The two channels which read the phase current sensors from the inverter have been fitted with offset shift circuits to allow conversion of the bipolar signals, since current flowing into a phase is indicated by a negative voltage from that sensor. These circuits

convert signals from 0–3 V to ± 1.5 V, allowing a range of 300 A (± 150 A) to be read at 10 mV/A.

Some of the general purpose I/O pins are connected to 5 A rated relays on the interface board. Two of these are utilised to switch the DC link contactors, which require up to 3.8 A to pull them in. This interface board also provides 15 and 5 V power supplies, each with a choke to reduce common mode noise. The board has been designed to reduce interference to an absolute minimum, with an internal ground plane, separate power tracks for the worst sources of noise, and decoupling capacitors where appropriate, amongst other measures.

An additional printed circuit board (PCB) was designed and built to provide equipment-specific interfaces. This provides appropriate connections between the DSP interface board and the inverter module for the gate drive, sensor, encoder and error signals. The power supply for the inverter's internal control logic and the inverter error signals are also routed through this board. BNC connections are provided to allow monitoring of the gate drive signals on an oscilloscope for testing purposes.

The operation of the controller—start/stop and current demands—during testing is controlled by a software interface implemented in the National Instruments LabVIEW package, which is shown in Figure 4.5. This communicates with the controller over an RS232 serial channel.

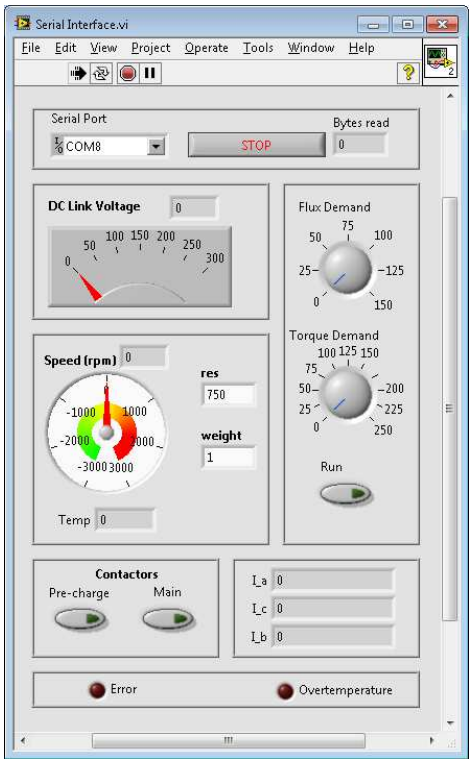

Figure 4.5: Software control panel in NI LabVIEW

4.3.4 Power Source

A DC power supply represents the EV battery in the test system. Two 600 V Regatron TopCon supply units are used together in a master-slave configuration. Each supply has a current rating of 66 A and a maximum power of 32 kW, which are doubled when the two devices are connected in parallel. These are supplied from a 415 V 3-phase mains connection.

A constant DC source is not a completely accurate representation of an EV battery, as the voltage of a battery will vary with state-of-charge, temperature and load, due to non-zero internal resistance. However, this work is concerned with the improvement of the drive system only and so battery-related effects are not desirable.

The rated line-to-line voltage is 160 V. The DC voltage needed to achieve this as a maximum is given by

$$V_{\mathrm{dc}} = 2\frac{\sqrt{2}}{\sqrt{3}}V_{l-l} = 261.3 \text{ V}. \tag{4.1}$$

Thus it seems reasonable to adopt the manufacturers recommended DC link voltage of 240 V.

A contactor is used to allow controlled connection and isolation of the DC supply and the DC bus of the inverter. In order to avoid a large in-rush current when the DC supply is switched across the DC link capacitor, a second contactor switches in a pre-charge resistor, through which the capacitor charges. When the capacitor is fully charged, the main contactor can be closed, short-circuiting the pre-charge resistor.

To give a charging time of 5 s, a 1 kΩ resistor is required, since the charging is approximately equal to five time constants of the RC circuit:

$$T = 5RC. \tag{4.2}$$

Although the power dissipated reduces as the capacitor charges, when the full voltage is applied, the instantaneous power dissipated by the resistor will be:

$$P = \frac{V^2}{R} = 57.6 \text{ W} \tag{4.3}$$

Thus a 100 W, 1 kΩ wirewound resistor was used.

Since the contactors will potentially need to break large DC currents they must be appropriately rated, since AC rated contactors rely on zero-crossing points at which less force is required to open the contacts.

4.3.5 Dynamometer

The test platform utilises part of an existing test facility at Newcastle University. Part of this facility is a testbed for machines up to 100 kW and 30 000 rpm. The dynamometer is a 100 kW IM (seen in the top of Figure 4.6) with a maximum speed of 6 000 rpm, with a 1:5 speed-increasing gearbox on the output shaft (centre of Figure 4.6). A torquemeter (shown at the bottom of Figure 4.6) is located on the gearbox output shaft to enable measurement of torque and speed.

The load machine is driven from a 100 kW variable speed regenerative drive with an active front end, seen in the background of Figure 4.6.

This drive system forms a load for the system under test, allowing different load conditions to be represented.

The test machine was installed on the testbed with the aid of a custom-made mounting plate and coupled to the torquemeter using flexible couplings and a shaft designed to meet the torque requirements of this project.

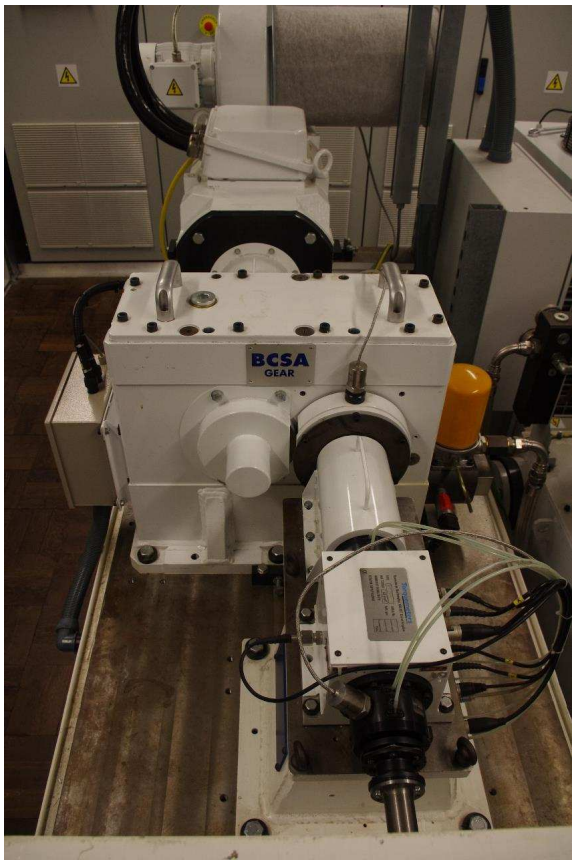

Figure 4.6: Photograph showing dynamometer, torque transducer and gearbox

4.3.6 Auxiliary Systems and Issues

4.3.6.1 Safety and Protection

A thorough risk assessment was carried out on the test platform, and a range of safety measures implemented.

The most obvious precaution to take is to ensure all conductors and connectors are appropriately insulated and covered to ensure they cannot be touched accidentally whilst live. To this end, a cover was built to enclose the inverter, on which the DC link and output phase connections are exposed. The cover has an interlock mechanism in the form of a switch which is connected to the emergency stop input on the DC supply, shutting it down in the event that the cover is opened.

To ensure that the DC link capacitor does not remain charged after the DC link voltage is removed, a bleed resistor is connected across the DC rails. A 10 kΩ resistor was chosen, which with a 240 V DC link will dissipate a negligible amount compared to the power being supplied to the inverter even under a light load, but discharge the capacitor relatively quickly:

$$P_{\text{bleed}} = 5.76 \text{ W} \tag{4.4}$$

$$T_{\text{bleed}} = 5RC = 50 \text{ s} . \tag{4.5}$$

A 15 W wirewound resistor was used for this part.

Rotating machinery presents a risk of entanglement and injury from debris should a mechanical failure occur. An appropriately-rated steel guard placed over the exposed rotating components minimises both these risks.

The testbed is located in its own enclosed cell, which is temperature-controlled and soundproofed. Safety is further improved if the equipment within can be operated from outside of this room, and there is a bench and control panel for the load drive outside to enable this. The test drive system is operated from a computer which is controlled remotely from this bench, as shown in Figure 4.7.

It is important to be able to shut the test down from this bench in the event of unexpected operation. The control panel features an emergency stop button which shuts off power to the load drive when pressed. A separate emergency stop switch was installed which shuts down the DC supply, and thus the entire system can be shut down without entering the test cell.

Figure 4.7: Photograph showing control bench outside test cell

4.3.6.2 Cooling

A new water-based cooling system was commissioned to meet the cooling needs of the test motor and inverter. A 16 kW industrial cooler with a nominal flow rate of 45 l/min is used to pump coolant through both the motor and inverter, which are connected in series.

4.3.6.3 Instrumentation

A 64-line quadrature encoder is built into the test motor's output shaft bearing.

The signals indicate the rotor position and are used in the FOC algorithm. This information is also used by the DSP to calculate the motor speed which is transmitted to the control PC within the test cell via a serial connection. Here the LabVIEW 'virtual instrument' panel, seen in Figure 4.5, is used to display this and other data transmitted from the DSP.

The currents from the two current sensors, the DC link voltage, the heatsink temperature and the status of the two inverter signals are also sent and displayed on the LabVIEW panel.

4.3.7 Motor Parameterisation Tests

Some parameters of the machine are needed in order to accurately implement vector control and for the loss model. A standard test procedure was followed to determine the stator and rotor resistance, stator and rotor leakage inductance, magnetising inductance and equivalent iron loss resistance [60, 65]. This consists of three tests: a DC resistance test, a locked rotor test, and a no load test.

4.3.7.1 DC Resistance Test

The DC resistance test is a simple measurement of the resistance of the stator windings, performed by applying a voltage to two terminals of the machine. The resulting current is measured and the resistance derived according to Ohm's Law. Since the motor is star connected, this gives the resistance across two phases. All the phases are assumed to be equal, so the per-phase resistance is this figure divided by two.

Four tests were performed, at 60, 80, 100 and 120 A, and the mean of the results taken. The test was performed with a coolant temperature of 20 ℃ ±1.

4.3.7.2 Locked Rotor Test

The principle of the locked rotor test is to enforce a slip of one by physically locking the rotor whilst applying a rotating electric field. A small fraction of the rated voltage is applied to the machine to ensure negligible magnetising current, and also to limit the stator current, preventing heating of the stator windings. This leaves an equivalent circuit which effectively consists of only the stator resistance and leakage reactance, and the rotor resistance and leakage reactance as referred to the stator. Hence the real power input, which is measured, can be considered to be the sum of stator and rotor resistance effects. Calculating the stator copper loss from the stator resistance as measured in the DC resistance test, the rotor resistance referred to the stator can be determined:

$$R'_r = \frac{P_{in}}{3I^2} - Rs. \tag{4.6}$$

Once the real component of the equivalent circuit impedance has been calculated from the real power, as above, the imaginary component can be derived from the total impedance, $Z = U/I$.

$$X = \sqrt{Z^2 - R^2} \tag{4.7}$$

This includes the stator and referred rotor leakage reactances, but it is not possible to know the value of each. It is commonly assumed that they are equal and so each is half the total reactance.

The test is performed at a range of frequencies and the calculated referred rotor resistance is extrapolated to zero to give a DC value.

4.3.7.3 No Load Test

In a similar way to the locked rotor test, the no load test eliminates part of the equivalent circuit, this time the rotor branch, by enforcing a slip of zero. The machine is run at synchronous speed with no load, resulting in an electrical field rotating synchronously with the rotor. Hence the current induced in the rotor is negligible, and the torque produced is effectively zero.

In this case, the test was performed with a 50 Hz sinusoidal supply, giving a synchronous speed of 1 500 rpm.

The machine is magnetised and so iron loss occurs. Copper loss occurs in the stator and this can be calculated from the stator resistance and current. Electrical power is supplied to overcome friction and windage losses. Since the current flowing in the rotor is negligible, rotor copper loss can be ignored. Therefore the real input power is the sum of the stator copper loss and iron loss.

The magnetising voltage can be calculated by subtracting the voltage drop across the stator leakage reactance, known from the locked rotor rest, from the phase voltage. Then the equivalent iron loss resistance can be derived from the measured real input power:

$$R_f = \frac{V_m^2}{P_{\text{in}}} \tag{4.8}$$

The magnetising reactance can be determined from the magnetising voltage and current, using the assumption that the current flowing is entirely magnetising current, which is approximately true:

$$X_m = \frac{V_m}{I}. \tag{4.9}$$

4.3.7.4 Results

The parameters in Table 4.1 were obtained after following the procedure described above.

Parameter	Symbol	Value
DC stator resistance	R_s	10 m Ω
Referred DC rotor resistance	R'_r	15.02 m Ω
Stator leakage inductance	L_{sl}	152.87 µH
Referred rotor leakage inductance	L'_{rl}	152.87 µH
Equivalent iron loss resistance	R_f	9.23 Ω
Magnetising inductance	L_m	2.2 mH

Table 4.1: Measured machine parameters

4.4 Benchmarking

The goal of the benchmarking process is to obtain a set of results from a conventional system against which to assess the effectiveness of the proposed solution by comparison. To this end, a controller using the traditional FOC strategy in conjunction with a SVM scheme is implemented on the test platform. Using the same same hardware (motor, inverter, DSP, etc.) for the benchmark results and the novel scheme proposed in this work allows for the most credible comparison of the control scheme itself.

Since this work concerns a scheme for improved efficiency, most of the effort here is concentrated on determining the efficiency of the conventional system.

4.4.1 Operating Region

Efficiency figures are often quoted for a given piece of equipment, such as a motor or inverter, but often only for a single operating point; usually the point of maximum efficiency. In fact, the efficiency usually varies widely over the operating region. In the case of an electric motor, the efficiency is strongly dependent on speed and load torque.

Therefore, in order to provide a thorough comparison, the efficiency is determined at a number of points across the operating region, and the results compiled into an "efficiency map". The efficiency at each point can be represented by the height of a surface from the speed-torque plane in a three-dimensional plot or by colour in a two-dimensional plot.

The base speed of the test motor is 3 000 rpm, the nominal torque is 130 Nm and the maximum speed is 9 000 rpm. Therefore the operating region has a constant torque region bounded by the nominal torque up to the base speed, and thereafter, by a constant power region bounded by the nominal power of 40 kW, up to the maximum speed.

The area covered by the efficiency tests is limited by practical constraints: namely, the current capability of the dynamometer drive, which limits the torque that can be tested to 63.6 Nm. To allow some margin for error the maximum test torque has been set as 60 Nm.

The maximum test speed has also been set to 3 000 rpm, the base speed of the machine under test. This avoids the need to operate in the field-weakening region, which simplifies the comparison of results.

In electric vehicle applications, the machine is usually operated as a generator when required to provide braking torque and recuperate energy that would otherwise be lost using friction brakes. This mode of operation is not tested in this project due to a restriction on the export of power to the grid in the test facility. In theory, the principles seen are equally applicable to generator operation.

The operating region of the motor and the area covered by the tests are shown in Figure 4.8.

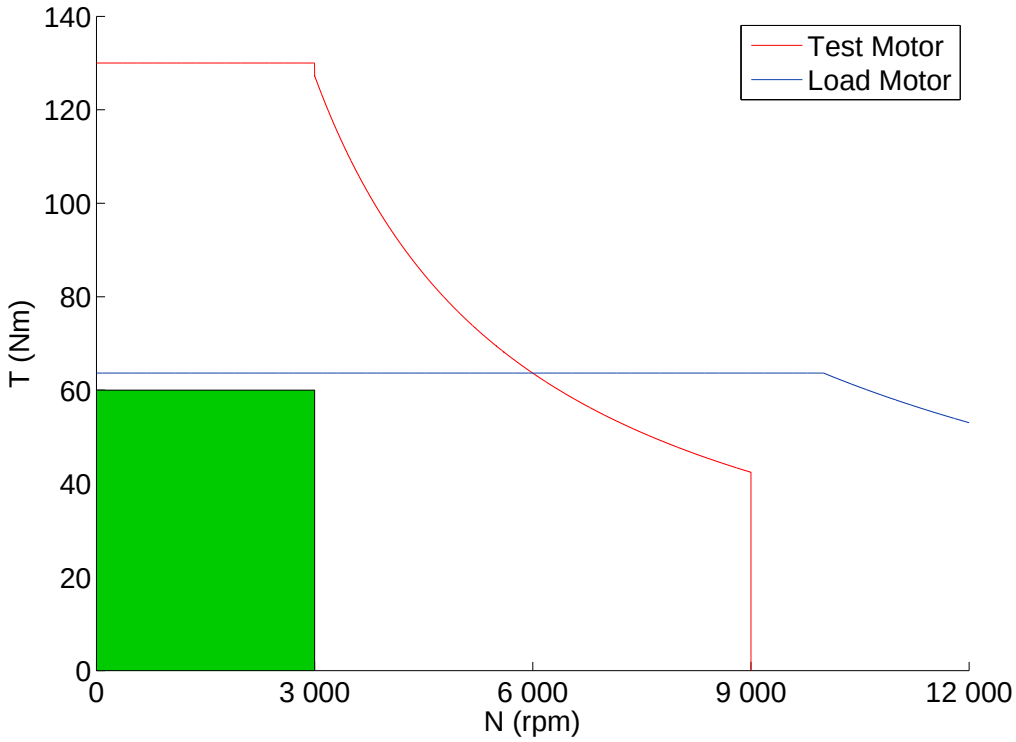

Figure 4.8: Operating region covered by efficiency map and showing limits of test area

4.4.2 Measurement Uncertainty Requirements

Although accurate, determining efficiency using a calorimetric set-up is expensive and time-consuming. Thus it is rarely used in practice, and is usually limited to very small drives lower than 1 kW, although there is anecdotal evidence that some larger companies are beginning to use this method for higher power drives.

Furthermore, whilst using datasheet loss figures for each component in the system can be shown to be an accurate method of efficiency determination [75], it is not considered appropriate for this situation, since some of this information will be used in the loss model.

Thus it is decided that the input-output method of efficiency measurement will be used for benchmarking and subsequent comparison.

It is important when considering the measured efficiency to have a quantified certainty of its accuracy in order to show that differences between results are due to different control schemes, and not simply measurement error. A typical or maximum relative error is usually quoted for measurement instrumentation, such as voltage probes or current clamps, by the device manufacturer. All the errors in the chain of measurement devices and interfaces must be considered in order to give a true calculation of the uncertainty of the final result.

When a measured signal passes through many devices, an error at one stage will be compounded by further stages. Therefore it is important to consider the effect of all sources of error in the measurement chain. Calculations for the typical uncertainty of efficiency results are presented in the following section based on the actual equipment used.

4.4.3 Instrumentation

Power flow in the drive system may be measured at three points: the input (power at the DC link P_{dc}), the inverter output/motor input (three-phase power P_{3ph}) and the motor shaft (mechanical power P_{mech}). Since the machine is operated in motoring mode, positive power is defined as flowing from the DC link, via the inverter and motor, to the output shaft, as shown in Figure 4.9. This allows the efficiencies of the inverter, η_{inv}, and motor, η_{mot}, to be calculated separately, as well as the total drive efficiency, η_{drive}. These are defined as follows:

$$\eta_{\mathrm{inv}} = \frac{P_{\mathrm{3ph}}}{P_{\mathrm{dc}}} \tag{4.10}$$

$$\eta_{\mathrm{mot}} = \frac{P_{\mathrm{mech}}}{P_{\mathrm{3ph}}} \tag{4.11}$$

$$\eta_{\mathrm{drive}} = \frac{P_{\mathrm{mech}}}{P_{\mathrm{dc}}} \tag{4.12}$$

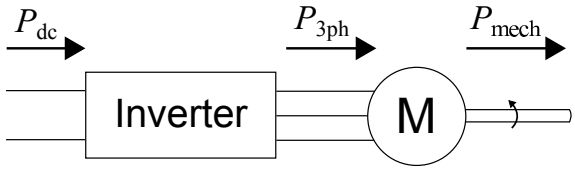

Figure 4.9: Power flow definitions

The DC power is measured by recording the DC link voltage and current using isolated voltage probes and current clamps. The mean power over the measurement period is then calculated simply as the integral of the product of these two measurements at each sampling point, over the measurement period.

The three-phase power is determined from each phase current and the voltage between each line and a neutral reference point.

The mean power can then be computed as the integral of these instantaneous powers over the measurement period. In order to arrive at a true figure, the measurement period is windowed such that a whole number of electrical cycles is covered, since the instantaneous power varies over the cycle. This is achieved by applying a zero-crossing detection algorithm on one of the phase currents to find the first and last positive-going zero crossings in each record. All the power calculations are then performed over this interval for consistency.

The mechanical power is defined as the product of speed and torque and is again integrated over the measurement period to find the mean. The speed is obtained from the quadrature encoder in the motor bearing. The pulse train is recorded and

the number of pulses in the measurement period counted to determine the speed. Signals from the torquemeter are read by a signal processor, which outputs a voltage proportional to the torque.

The sampling frequency needs to be significantly higher than the PWM frequency, in order to detect the edges of the voltage waveforms with sufficient accuracy. Therefore a sampling frequency of 1 MS/s, one hundred times the PWM frequency, is used. This effectively means that the duty cycle of the voltage signals is measurable with a resolution of 100. The sampling window needs to be sufficiently long to ensure a truly representative mean power is measured over several electrical and mechanical cycles. A sampling window of 0.5 s is chosen, giving records of 500 000 samples.

4.4.3.1 Measurement Uncertainty

In this section the total uncertainty of each measured value will be calculated from the stated error of each piece of equipment involved.

Where two or more measured quantities, say x and y, are combined to calculate a resultant quantity, say z—when a power is calculated from a measured voltage and current, for instance—the uncertainty in each measured value must be combined to determine the total uncertainty in the resultant quantity.

When two measured quantities are multiplied or divided, the combined relative uncertainty of the result ε_z is the summation in quadrature of the individual relative uncertainties ε_x and ε_y [76]:

$$z = xy \tag{4.13}$$

$$\varepsilon_z = \sqrt{\varepsilon_x^2 + \varepsilon_y^2}. \tag{4.14}$$

When two quantities are added or subtracted, the combined absolute uncertainty of the result S_z is the summation in quadrature of the individual absolute uncertainties S_x and S_y, usually expressed as a standard deviation S:

$$z = x + y \tag{4.15}$$

$$S_z = \sqrt{S_x^2 + S_y^2}. \tag{4.16}$$

The combined relative uncertainty can be calculated as

$$\varepsilon_z = \frac{S_z}{z} = \frac{1}{z}\sqrt{\varepsilon_x x^2 + \varepsilon_y y^2}. \tag{4.17}$$

Where the "typical" error is given for a piece of equipment by its manufacturer, this is assumed to be the relative standard deviation, whereas if a "maximum" error is given, this is assumed to be the maximum error of a rectangular distribution and so the standard deviation is obtained by dividing by $\sqrt{3}$. If not explicitly stated, quoted error values are assumed to be maximum values.

The sources of uncertainty in the chain of measurement are illustrated in Figure 4.10.

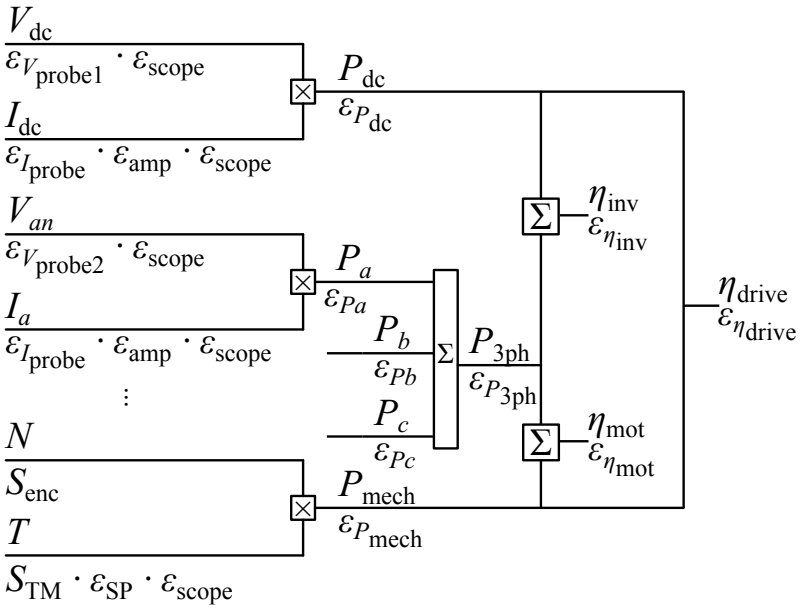

Figure 4.10: Chain of measurement and sources of error

The stated error for each piece of equipment is given in Table 4.2. The speed is determined by counting the edges of the square wave encoder signal and so the uncertainty comes from the restriction of an integral number of edges.

Table 4.2: Stated errors for each device in the measurement chain

Device	Symbol	Error
Pico TA041 differential voltage probe	$\varepsilon_{\text{Vprobe1}}$	2%
Yokogawa 700924 differential voltage probe	$\varepsilon_{\text{Vprobe2}}$	2%
Tektronix TCP404XL current probe	$\varepsilon_{\text{Iprobe}}$	1%
Tektronix TCPA400 current probe amplifier	ε_{amp}	1%
Quadrature encoder [1]	S_{enc}	0.9375 rpm
Torquemeter [2]	S_{TM}	0.32 Nm
Signal processor	ε_{SP}	0.1%
Yokogawa DL850V with input module 701251	$\varepsilon_{\text{scope}}$	0.25%

[1] The maximum error is 1 edge per 1 s sample, representing 1⁄64 of a revolution using a 64 line encoder, equivalent to 0.9375 rpm.

[2] The torquemeter accuracy is given as within 0.01% of full scale (160 Nm).

Table 4.3 shows combined uncertainties for quantities derived from these measurements, and their calculations. The 3-phase power is obtained by addition of each of the 3-phase power measurements, and hence the calculation of its uncertainty involves absolute values. However, as long as the ratio between total 3-phase power and the individual phase powers is constant, the relative uncertainty is constant. Thus a constant ratio of 3:1 is used; i.e. a balanced condition is assumed. The speed and torque errors are absolute values, and so the larger relative uncertainty is obtained at lower measured values. In order to represent the worst case scenario, the values from the lowest test point, 500 rpm and 10 Nm, are used in these calculations.

Table 4.3: Combined uncertainties for derived values

Quantity	Calculation	Uncertainty (%)
V_{dc}	$\varepsilon_{V_{\mathrm{dc}}} = \sqrt{\varepsilon^2_{V_{\mathrm{probe1}}} + \varepsilon^2_{\mathrm{scope}}}$	2.016
I_{dc}	$\varepsilon_{I_{\mathrm{dc}}} = \sqrt{\varepsilon^2_{I_{\mathrm{probe}}} + \varepsilon^2_{\mathrm{amp}} + \varepsilon^2_{\mathrm{scope}}}$	1.414
P_{dc}	$\varepsilon_{P_{\mathrm{dc}}} = \sqrt{\varepsilon^2_{V_{\mathrm{dc}}} + \varepsilon^2_{I_{\mathrm{dc}}}}$	2.462
V_{an}, V_{bn}, V_{cn}	$\varepsilon_{V_{an}}, \varepsilon_{V_{bn}}, \varepsilon_{V_{cn}} = \sqrt{\varepsilon^2_{V_{\mathrm{probe2}}} + \varepsilon^2_{\mathrm{scope}}}$	2.016
I_a, I_b, I_c	$\varepsilon_{I_a}, \varepsilon_{I_b}, \varepsilon_{I_c} = \sqrt{\varepsilon^2_{I_{\mathrm{probe}}} + \varepsilon^2_{\mathrm{amp}} + \varepsilon^2_{\mathrm{scope}}}$	1.436
P_a, P_b, P_c	$\varepsilon_{P_a}, \varepsilon_{P_b}, \varepsilon_{P_c} = \sqrt{\varepsilon^2_{V_{an}} + \varepsilon^2_{I_a}}$	2.475
$P_{3\mathrm{ph}}$	$\varepsilon_{3\mathrm{ph}} = \frac{1}{P_{3\mathrm{ph(nom)}}} \sqrt{\varepsilon_{P_a} P_a^2 + \varepsilon_{P_b} P_b^2 + \varepsilon_{P_c} P_c^2}$	9.083
N	$\varepsilon_N = \frac{S_{\mathrm{enc}}}{N_{\mathrm{nom}}}$	0.188
T	$\varepsilon_T = \sqrt{\frac{S_{\mathrm{TM}}}{T_{\mathrm{nom}}}^2 + \varepsilon^2_{\mathrm{SP}} + \varepsilon^2_{\mathrm{scope}}}$	3.211
P_{mech}	$\varepsilon_{P_{\mathrm{mech}}} = \sqrt{\varepsilon^2_N + \varepsilon^2_T}$	3.217
η_{inv}	$\varepsilon_{\eta_{\mathrm{inv}}} = \sqrt{\varepsilon^2_{P_{3\mathrm{ph}}} + \varepsilon^2_{P_{\mathrm{dc}}}}$	9.411
η_{mot}	$\varepsilon_{\eta_{\mathrm{mot}}} = \sqrt{\varepsilon^2_{P_{\mathrm{mech}}} + \varepsilon^2_{P_{3\mathrm{ph}}}}$	9.636
η_{drive}	$\varepsilon_{\eta_{\mathrm{drive}}} = \sqrt{\varepsilon^2_{P_{\mathrm{mech}}} + \varepsilon^2_{P_{\mathrm{dc}}}}$	4.051

The total uncertainty figures show that a much higher confidence can be placed in the total drive efficiency results than the results for the efficiency of either the inverter or motor in isolation. This is due to the uncertainty in the three-phase power measurement which is included in the individual inverter and motor measurements, but not the total drive measurement. This measurement has a high uncertainty largely due to the higher number of individual measurements—three voltages and three currents compared with two quantities each in the other power measurements—which contribute uncertainty to the total.

4.4.4 Results

The test procedure was carried out and data obtained according to the process described above.

The results are plotted on speed-torque axes, with the efficiency represented by colour. Interpolation is used to better show trends.

At higher power, and hence current, the resistive component of the IGBT voltage drop given by $R_{ce} \cdot i_c$ becomes more significant and the conduction loss in the inverter becomes less dominated by the constant voltage drop V_{ce0}. Therefore, despite the i_c^2 dependence, there is proportionately less conduction loss at higher current, and hence higher efficiency. A similar argument can be made for the diode.

The switching loss in the IGBT is proportional to current, but in the diode it is proportional to $i^{0.6}$, and so the total switching loss reduces relative to the current at higher current, and hence power.

The combined effects of these relationships between inverter losses and current explain the higher efficiency seen at higher torque and speed in Figure 4.11.

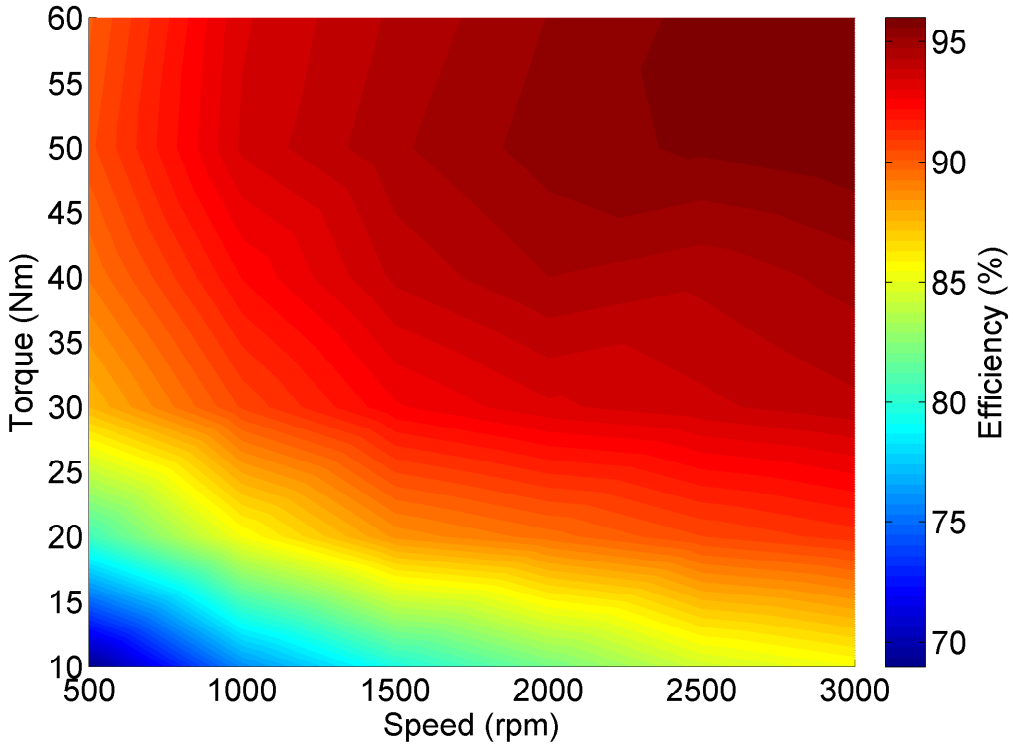

Figure 4.11: Map of inverter efficiency under conventional control

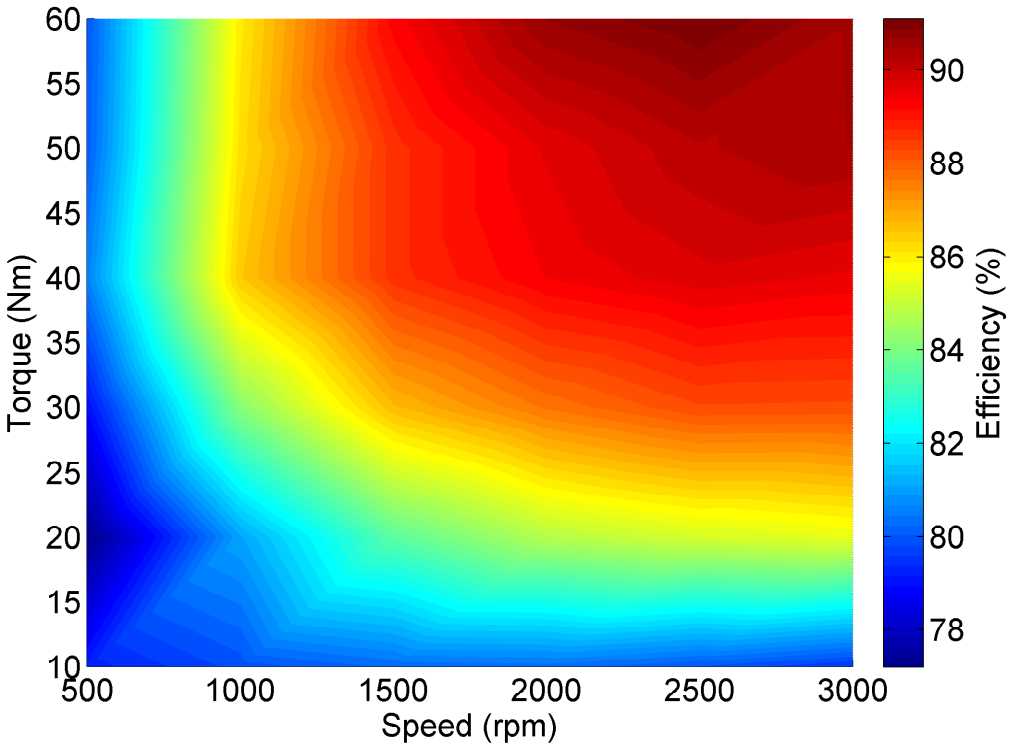

Figure 4.12: Map of motor efficiency under conventional control

The efficiency map of the motor in Figure 4.12 shows an efficiency profile typical of an induction machine. At very low speeds, higher rotor losses result in lower efficiency [77]. The efficiency increases with speed due to reduced copper and iron losses, relative to mechanical output power, though eventually this trend will reverse due to increasing friction and windage losses, in the region off the right-hand side of this figure.

Figure 4.13 shows how the efficiencies of the inverter and motor combine to give the total drive efficiency.

Figure 4.14 shows the torque ripple produced at each operating point with conventional modulation. The torque ripple is defined as the ratio of the range of torque values encountered to the mean torque. Generally the torque ripple is lower at higher load torque, due to the damping effect of the higher load. This varies widely with speed however, possibly indicating some interaction between the motor and the rest of the mechanical system which results in the torque ripple being amplified or attenuated more at certain frequencies than others.

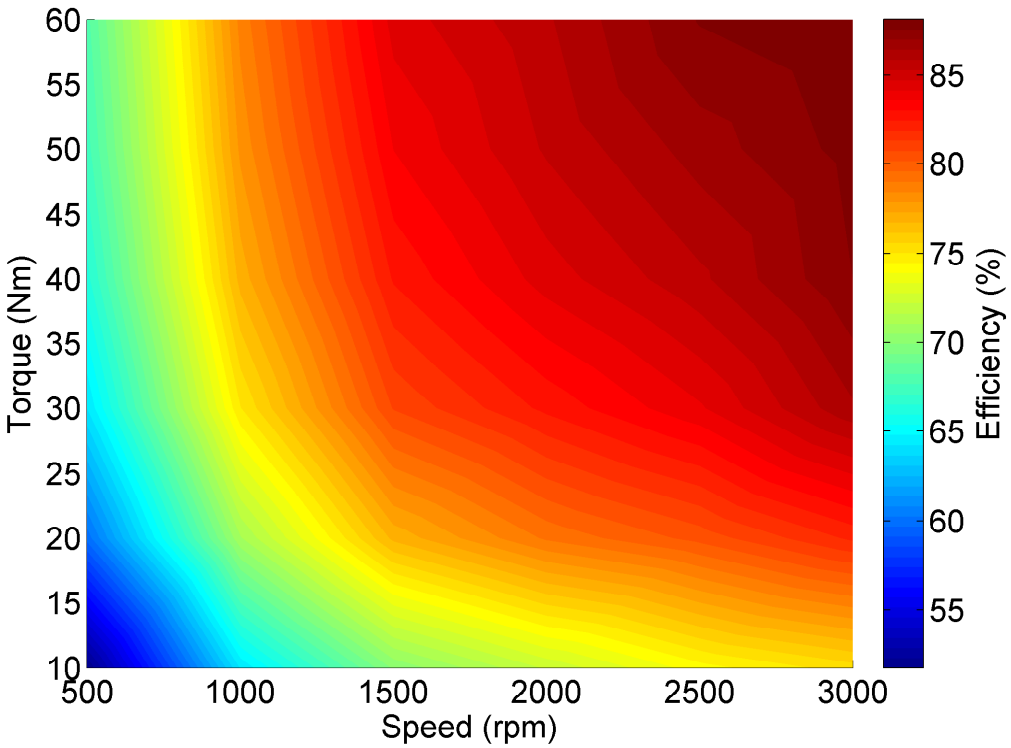

Figure 4.13: Map of total drive efficiency under conventional control

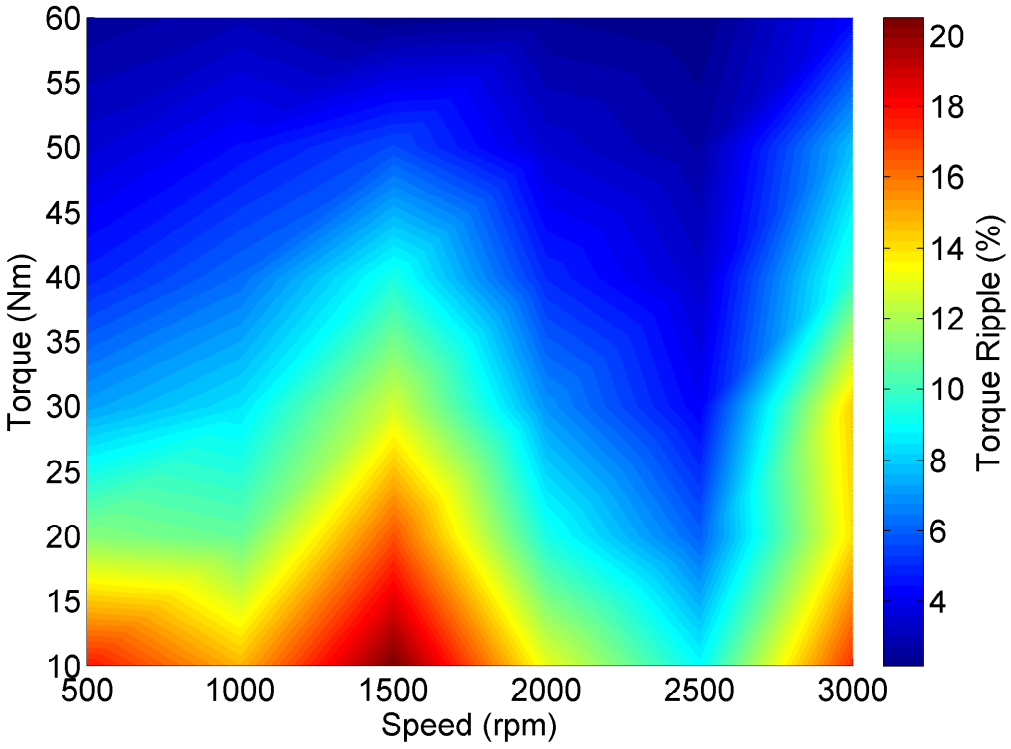

Figure 4.14: Map of torque ripple under conventional control

4.5 Practical Verification of New Controller Structure

As the modulation method with the most promising simulation results of the three proposed, the discretised modulation scheme was implemented on the practical test system described in Section 4.3. High resolution (7500) and low resolution (75) conditions were tested. The resolution is defined in Section 3.1.5. Corresponding results from the conventional FOC scheme are presented for comparison.

All the tests are performed at 1 500 rpm with a 30 Nm load. Note that in this practical situation, the load is determined in the open loop by the torque current demand given to the system under test. The load machine is operated under speed control; its controller regulates the load torque to match that which it sees from the test system in order to maintain the speed set point. This means there is a constant offset from the nominal torque in some of the test results which is very difficult to eliminate in a practical test situation. This is due to an inability to accurately determine the precise torque current demand required to produce exactly the nominal torque in the open loop, rather than an inaccuracy in the controller itself.

4.5.1 Conventional Modulation

Figure 4.15 shows the phase currents and torque from the conventional modulation test. The spectrum of the phase current is shown in Figure 4.16.

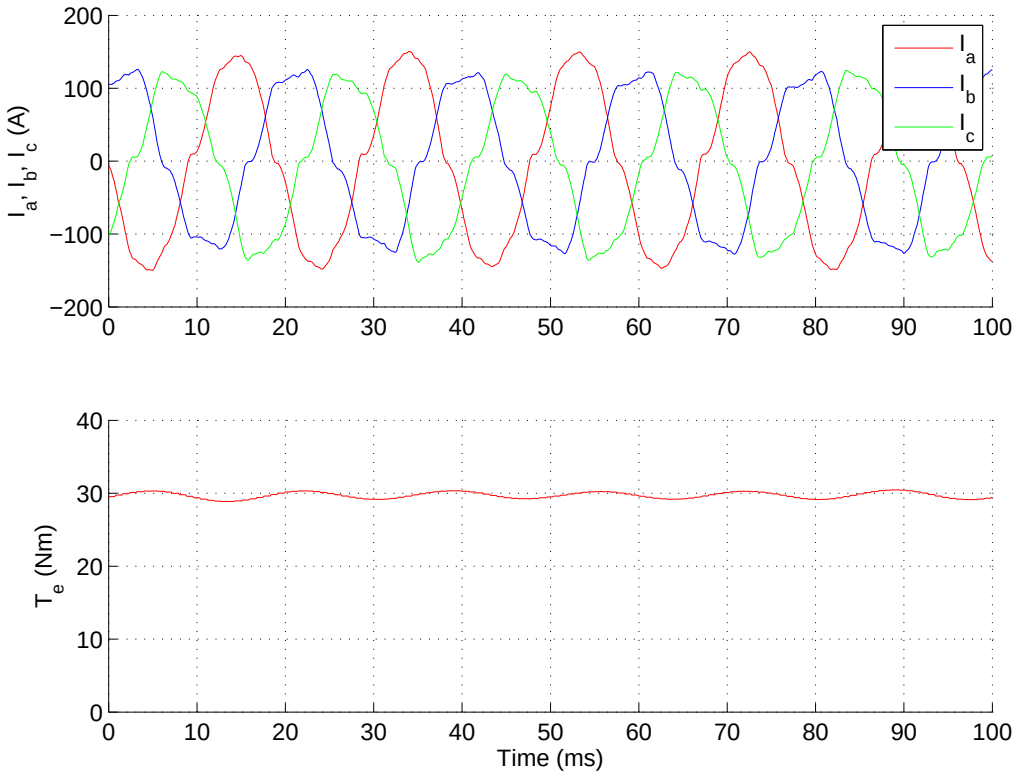

Figure 4.15: Measured phase currents and torque with conventional modulation scheme - 1 500 rpm, 30 Nm nominal

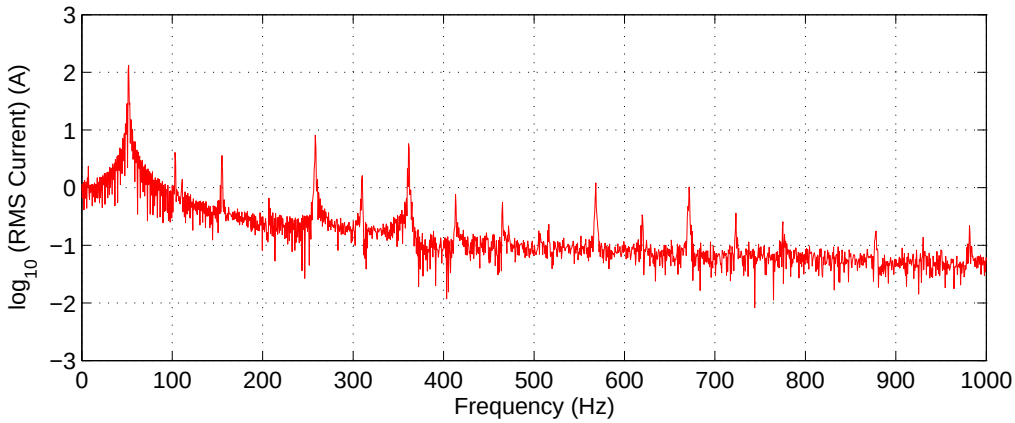

Figure 4.16: Spectrum of measured phase current with conventional modulation scheme

4.5.2 Discretised Modulation

4.5.2.1 High Resolution

Figure 4.17 shows the phase currents and torque from the high resolution (7500) discretised modulation test. The phase current spectrum is shown in Figure 4.18.

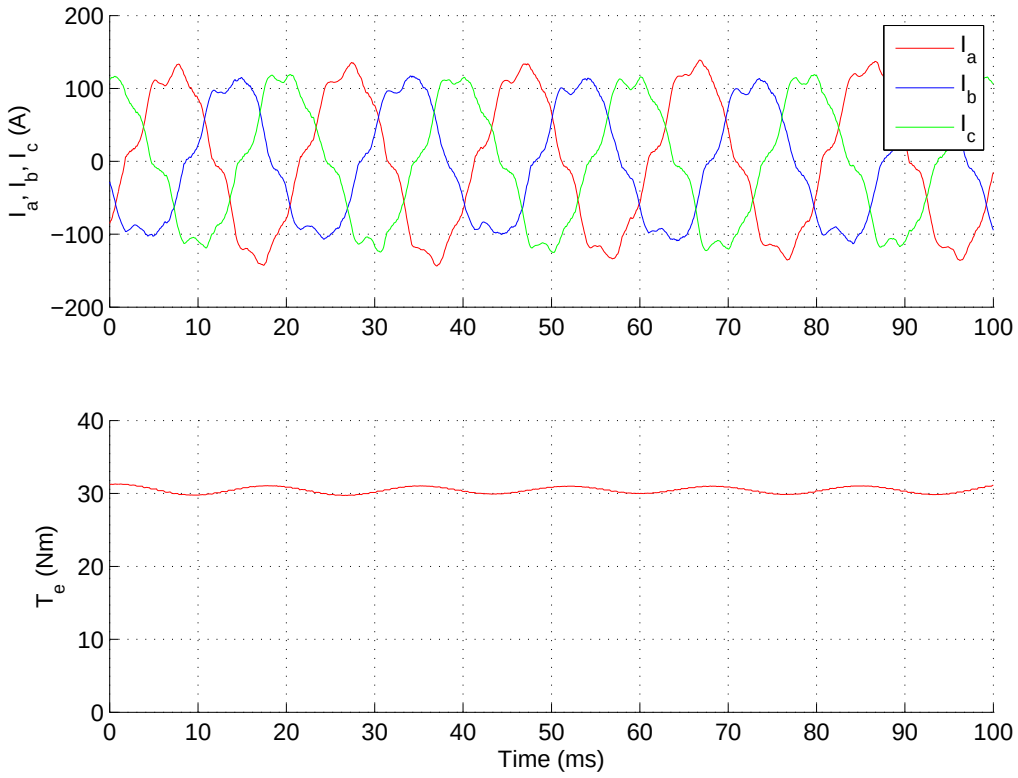

Figure 4.17: Measured phase currents and torque with high resolution (7500) discretised modulation scheme - 1 500 rpm, 30 Nm nominal

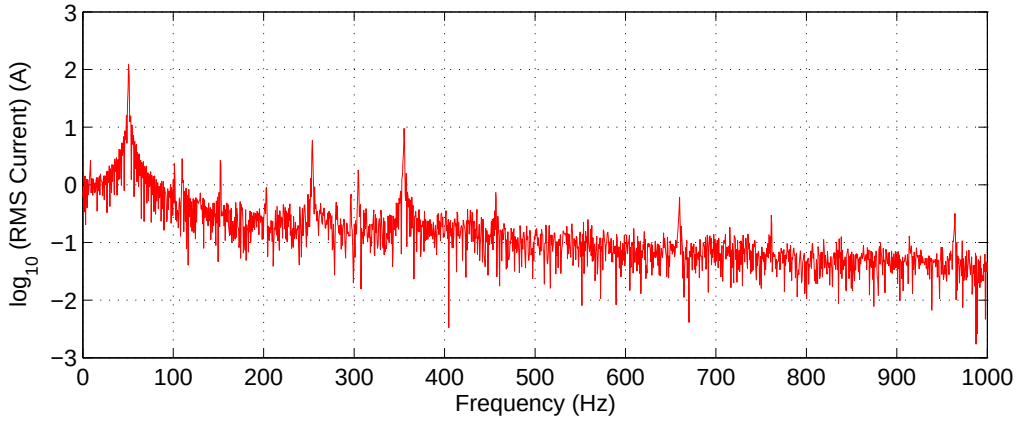

Figure 4.18: Spectrum of measured phase current with high resolution (7500) discretised modulation scheme

4.5.2.2 Low Resolution

Figure 4.19 shows the phase currents and torque from the low resolution (75) discretised modulation test. The phase current spectrum is shown in Figure 4.20.

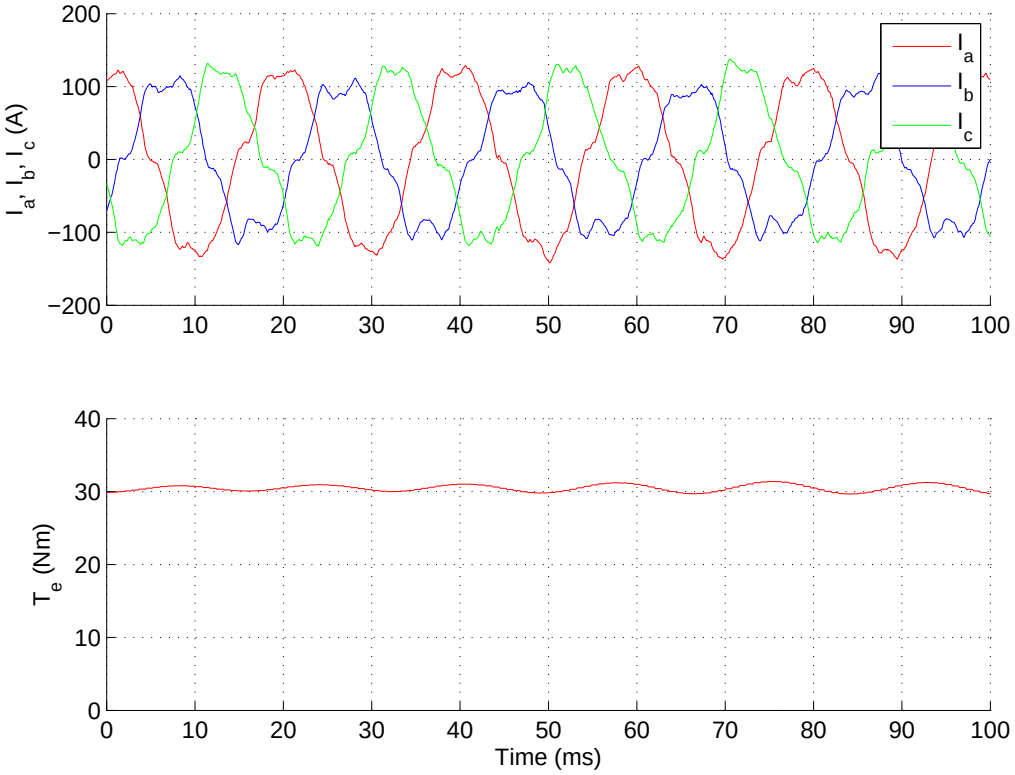

Figure 4.19: Measured phase currents and torque with low resolution (75) discretised modulation scheme - 1 500 rpm, 30 Nm nominal

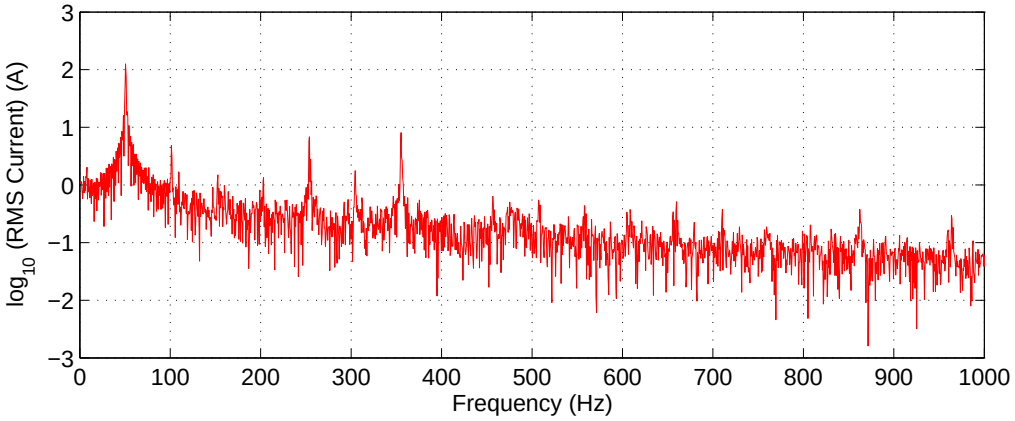

Figure 4.20: Spectrum of measured phase current with low resolution (75) discretised modulation scheme

4.5.3 Comparison

The first thing to note from the practical results is the lower torque ripple than the corresponding simulation results under all schemes, which is likely due to the larger load torque used. The THD in all schemes is considerably higher than in simulation. This is due to non-idealities not represented by the simulation, such as inaccuracies, delay and noise in the current sensor and ADC systems and deadtime in the switching action necessary to avoid shoot-through faults.

Unfortunately the reduction in THD under discretised modulation seen in simulation is not reflected in the practical results, but the increase is not prohibitive.

The torque ripple does appear lower than conventional, however, and as expected, this is more significant with higher resolution.

These results confirm what this section of work set out to prove, which is that it is possible to modify the control structure to allow for inclusion of external variables as control objectives through use of a cost function, without significantly degrading performance. Some figures of merit may even be improved.

Modulation scheme	Torque ripple (%)	THD (%)
Conventional	7.20	8.43
Discretised, high resolution	6.14	9.90
Discretised, low resolution	6.34	9.59

Table 4.4: Measured torque ripple and THD in practical tests at 1 500 rpm, 30 Nm

Chapter 5

Implementation of Loss Model and Loss Minimisation Scheme

5.1 Loss Model Implementation

5.1.1 Inverter Loss Model

The inverter loss model used is based on parameters provided by the device manufacturer, measured values of current and calculated device voltages.

5.1.1.1 Conduction Loss

The manufacturer provides values for the IGBT collector-emitter threshold voltage V_{ce0} and on-state collector-emitter resistance R_{ce}, which are shown in Table 5.1.

Parameter	Value	Units
IGBT collector-emitter threshold voltage V_{ce0}	0.85	V
IGBT on-state collector-emitter resistance R_{ce}	3.1	mΩ
Diode forward threshold voltage V_{f0}	0.8	V
Diode on-state forward resistance R_f	1.87	mΩ

Table 5.1: IGBT and diode characteristic values

This enables the collector-emitter voltage drop v_{ce} to be calculated for a given collector current i_c, which is considered to be equal in magnitude to the output current of the leg in question when that IGBT is conducting. The conduction power loss

P_{condQ} can then be calculated as the product of v_{ce} and i_c, and hence the conduction energy loss as this power over the appropriate conduction time:

$$v_{ce} = V_{ce0} + R_{ce} \cdot i_c \tag{5.1}$$

$$P_{\text{cond}Q} = v_{ce} \cdot i_c = V_{ce0} \cdot i_c + R_{ce} \cdot i_c^2 \tag{5.2}$$

$$E_{\text{cond}Q} = P_{\text{cond}Q} \cdot T \tag{5.3}$$

In a similar way, the diode conduction loss $P_{\text{cond}D}$ can be calculated from the datasheet values for the forward threshold voltage V_{f0} and forward resistance R_f, also found in Table 5.1. Likewise, when a diode is conducting, its forward current i_f is assumed to be the same magnitude as the appropriate phase current.

$$v_f = V_{f0} + R_f \cdot i_f \tag{5.4}$$

$$P_{\text{cond}D} = v_f \cdot i_f = V_{f0} \cdot i_f + R_f \cdot i_f^2 \tag{5.5}$$

$$E_{\text{cond}D} = P_{\text{cond}D} \cdot T \tag{5.6}$$

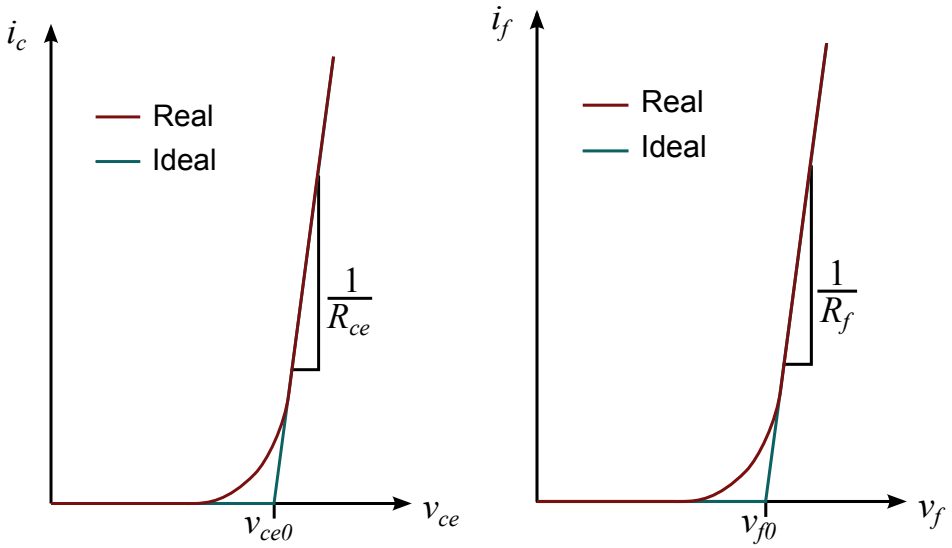

(a) Real and ideal IGBT output characteristics for a given gate-emitter voltage

(b) Real and ideal diode forward characteristics

Figure 5.1: IGBT and diode characteristics

The representation of the conduction losses in the IGBT and diode in this way is a linearisation of their output characteristics. The forward characteristic of a real IGBT for a given gate-emitter voltage V_{ge} is shown in Figure 5.1a alongside its ideal representation, as described by (5.1). The output characteristic of a diode is very similar to the IGBT forward characteristic. The real and ideal shape of this is shown in Figure 5.1b. The ideal characteristic is described by (5.4).

This model is a good approximation, particularly at higher currents. However, the on-state resistances of the devices will vary with temperature and this is not taken into account here.

The conduction loss models given here are very similar to all the models reviewed in Section 2.4. The only difference is that this model considers the voltage drop across the IGBT given by V_{ce0} making it more accurate, particularly at lower currents, than the physical model which only considers the IGBT as a simple resistance in the on-state.

5.1.1.2 Switching Loss

The switching loss is calculated from formulae supplied by the manufacturer. The energy loss in the IGBT is given as

$$E_{\mathrm{sw}Q} = 80.8 \times 10^{-3} \cdot \frac{i}{400} \cdot \left(\frac{V}{600}\right)^{1.4} \cdot (1 + 0.003(T_j - 125)). \tag{5.7}$$

The energy loss in the diode is given as

$$E_{\mathrm{sw}D} = 25.6 \times 10^{-3} \cdot \left(\frac{i}{400}\right)^{0.6} \cdot \left(\frac{V}{600}\right)^{0.6} \cdot (1 + 0.0065(T_j - 125)). \tag{5.8}$$

The dependence on temperature is neglected by assuming a constant temperature of 25 °C to reduce computational demand, and the voltage is assumed to equal the nominal DC link voltage of 240 V leaving the following:

$$E_{\mathrm{sw}Q} = 39.2 \times 10^{-5} i \tag{5.9}$$

$$E_{\mathrm{sw}D} = 5.17 \times 10^{-3} \cdot \left(\frac{i}{400}\right)^{0.6}. \tag{5.10}$$

In comparison to the models identified in the literature, the switching loss model given here is most similar to the analytical and empirical models. The analytical approach may be slightly less accurate than the model seen here, since it fits a curve only containing integral exponents to the characteristic, where as the expression here has non-integral exponents chosen to match the characteristic more closely. The empirical approach is much more complex and may represent the losses more exactly, but is impractical to implement over a wide range of operating conditions.

5.1.1.3 Validation

The inverter loss model presented above was validated against the experimental data from which the efficiency map in Figure 4.12 was produced. The total RMS current at each operating point was determined from the measured currents, and the loss

calculated according to (5.2), (5.5), (5.9) and (5.10) given a switching frequency of 10 kHz. The diodes and IGBTs were each assumed to be conducting for half the time. The corresponding efficiency was then calculated using the measured DC power at each point, and the results plotted in Figure 5.2. The results show very good agreement when compared with the measured efficiency in Figure 4.11, indicating that the model accurately represents the real losses.

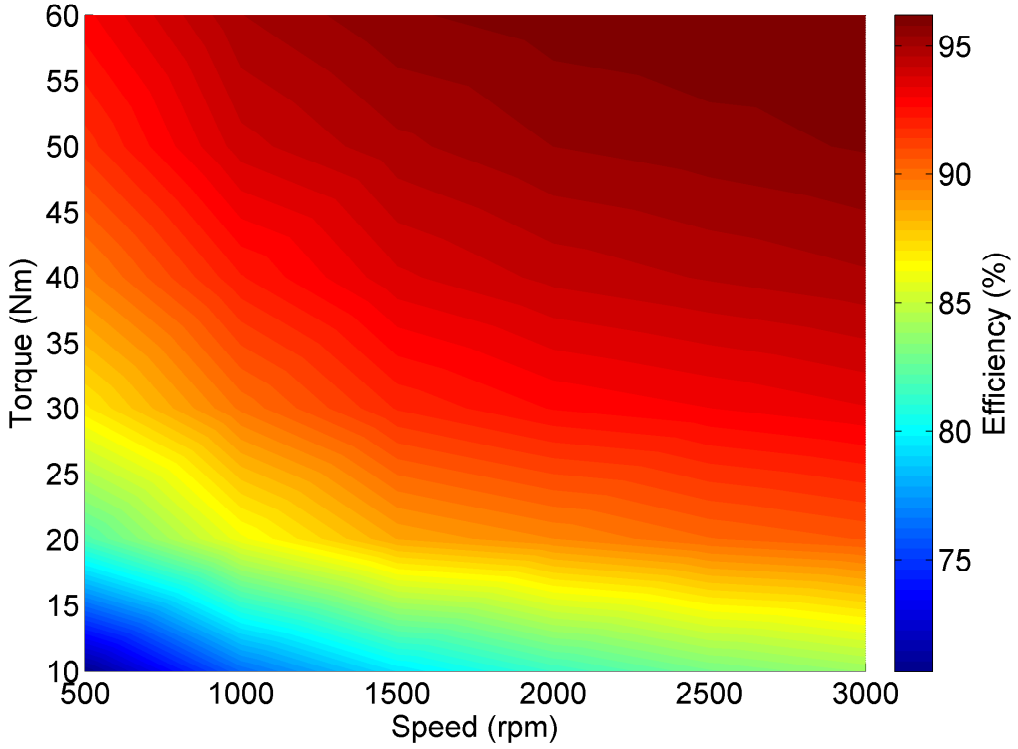

Figure 5.2: Inverter efficiency calculated from loss model and measured currents

5.1.2 Motor Loss Model

5.1.2.1 Copper Loss

Estimation of the copper loss is based on the equations presented in Section 2.4.2.1.

The stator copper loss is calculated from the stator current magnitude and the resistance of the stator windings:

$$P_{\text{cu-s}} = R_s(i_{ds}^2 + i_{qs}^2) \tag{5.11}$$

The measured stator currents are already available as they are used in the controller, and the stator winding resistance is easily measured. The stator resistance is dependent on temperature but as elsewhere in the model, this is ignored.

The rotor copper loss is calculated in a similar way, but the currents must be derived from stator quantities:

$$P_{\text{cu-r}} = R'_r i_r^2 \tag{5.12}$$

$$i_r = \frac{R'_f}{R'_f + R'_r} i_{qs} - \omega \frac{L'_m}{R'_f + R'_r} i_{ds} \tag{5.13}$$

$$L'_m = \frac{L_m^2}{L_r}. \tag{5.14}$$

The referred rotor resistance and rotor leakage inductance are obtained from a locked rotor test, whilst the magnetising inductance and iron loss resistance are obtained from a no load test. These tests and their results are explained in Section 4.3.7.

5.1.2.2 Iron Loss

The iron loss calculation is based on measurements of the iron loss during the no load test in Section 4.3.7. This test gives an equivalent iron loss resistance of 9.23 Ω. The iron loss is then calculated from the magnetising current according to (5.15).

$$P_{\text{iron}} = R'_f (i_{qs} - i_r)^2 \tag{5.15}$$

The rotor current is calculated as before, in (5.12).

This is a very coarse approximation of the iron loss since it does not take into account the frequency of the motor currents or their harmonic content, which have a very significant effect on the iron loss. The iron loss resistance used is valid for a 50 Hz sinusoidal supply. However, since the modulation scheme cannot ultimately effect the fundamental frequency, it is unnecessary to calculate the effect of this in the loss model.

5.1.2.3 Other Losses

Anything other than a very simple approximation of stray load loss, by assuming a constant proportion of the loss, is impractical in a real-time application. Thus stray load loss is not considered in this model.

The friction and windage losses cannot be affected by the modulation scheme whilst maintaining torque control and so are also ignored.

5.1.2.4 Validation

The motor loss model is validated in a similar way to the inverter loss model. The measured currents from the experimental data used to produce Figure 4.12 were used

to find the d- and q-axis RMS currents. The motor loss was then calculated according to (5.11), (5.12) and (5.15), using the motor parameters determined in Section 4.3.7. The results are plotted in Figure 5.3. Generally, there is good agreement with the practical results in Figure 4.12, but there are also some discrepancies. The range of efficiency encountered is similar in both the measured and calculated results, but there is divergence between the results in some areas, particularly at high speeds and low torque, where the calculated efficiency increases faster with increasing torque than the practical results suggest.

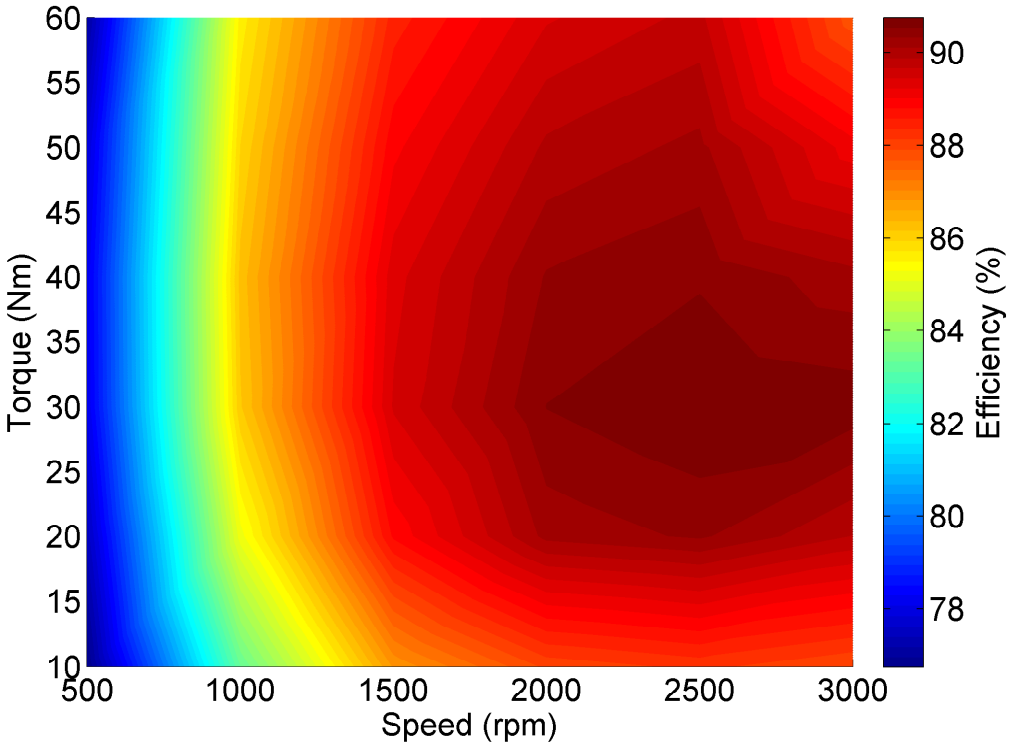

Figure 5.3: Motor efficiency calculated from loss model and measured currents

5.2 Calculation of Loss Reduction

With certain assumptions, the loss reduction possible under the proposed loss minimisation scheme can be calculated.

5.2.1 Conduction Loss

The reduction in conduction loss is caused by the IGBTs and diodes—which have different conduction losses for a given current—conducting each phase current for a different proportion of the switching period due to the loss minimisation scheme. The conduction loss for both types of device is shown against device current in

Figure 5.4, calculated from (5.2) and (5.5). The difference in loss ΔP_{Q-D}, also shown in Figure 5.4, is given as a function of current by (5.16).

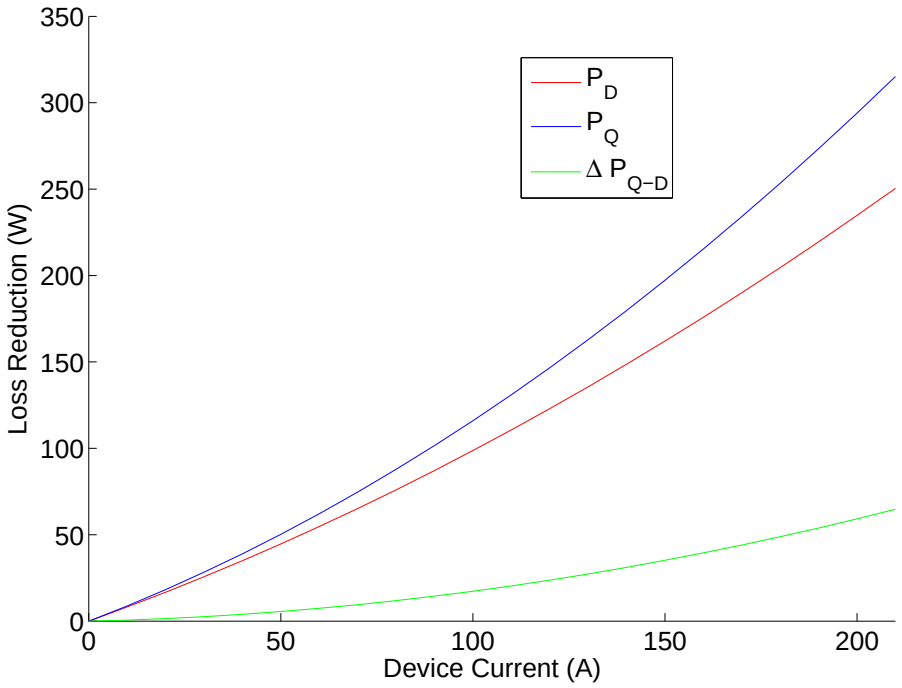

Figure 5.4: Relationship between conduction loss and current in an IGBT and a diode

$$\Delta P_{Q-D} = 0.05i + 0.00123i^2 \tag{5.16}$$

The maximum reduction in conduction loss in one phase is obtained when the discretised reference value causes the diode to conduct for a period T_{pwm}/r longer than in the conventional scheme (and the IGBT a period shorter by the same amount). The maximum difference between the pure demanded reference value and the discretised value is one discretisation level. Since the process is applied to all three phases, the mean conduction loss reduction is given by:

$$\Delta P_{\text{cond}} = \frac{\Delta P_{Q-D}}{r} \cdot 3. \tag{5.17}$$

The reduction in conduction loss calculated from (5.17) is shown for resolutions in the range 5–75 for values of current up to the full rating in Figure 5.5.

The estimated conduction loss reduction for each point on the tested operating range for a resolution of 75 is calculated using the same experimental data as was used to produce Figure 4.11 and (5.16). This is shown in Figure 5.6.

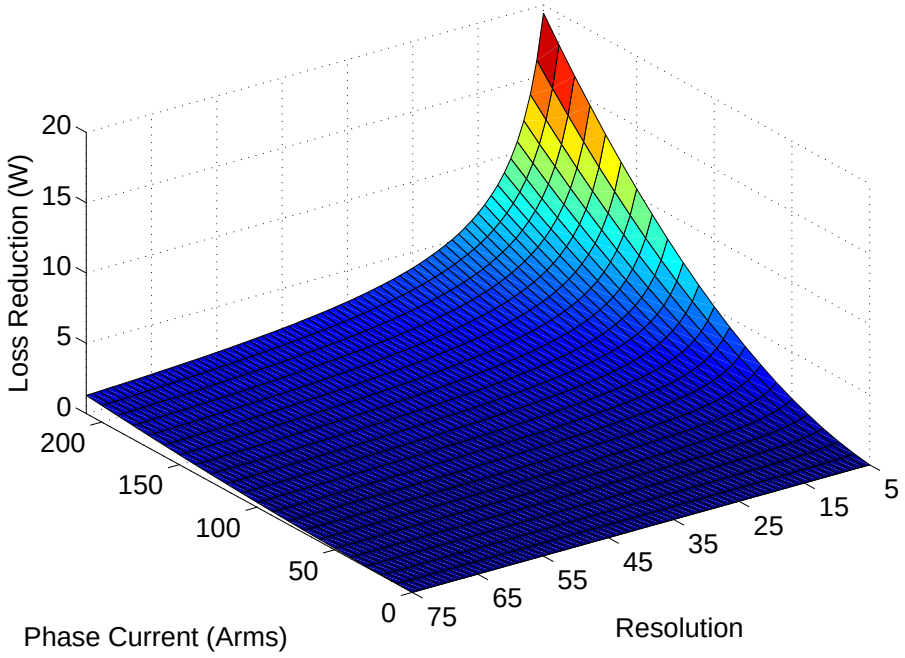

Figure 5.5: Relationship between conduction loss reduction, current and resolution

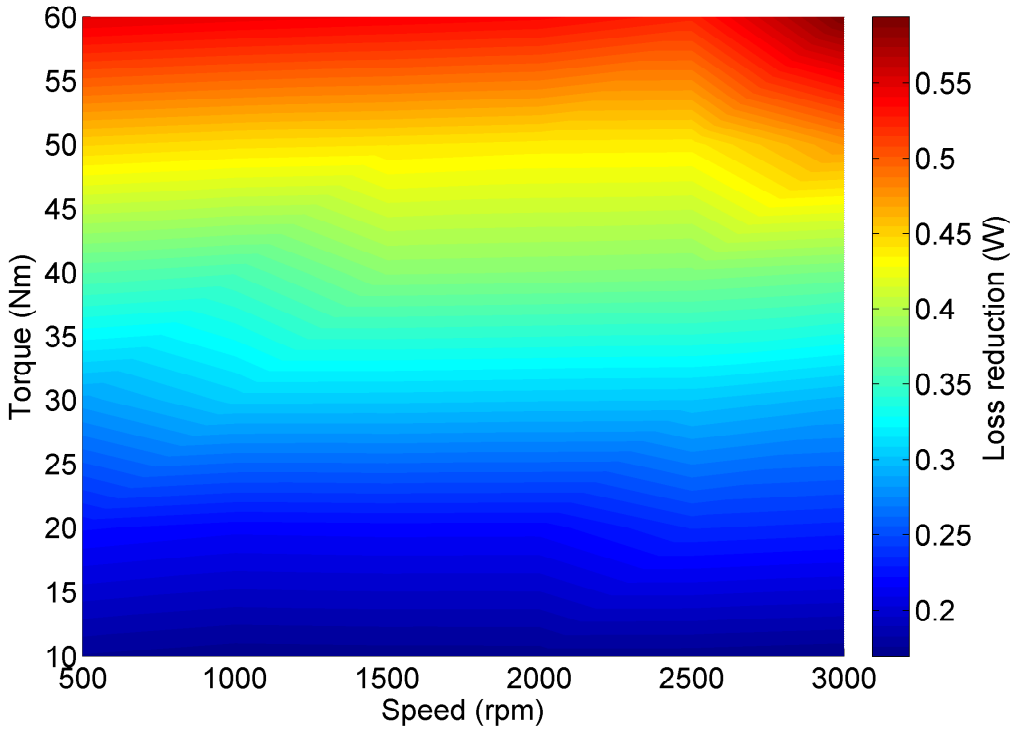

Figure 5.6: Estimated conduction loss reduction calculated from measured currents

5.2.2 Switching Loss

Unlike the conduction loss calculations, where the current is assumed to remain constant over the switching period, to determine a difference in switching loss, the change in current within the switching period must be considered, since the switching loss is dependent on the current at the point of switching. To identify a difference in switching energy caused by moving the point of switching in time, the variation of the current over time must be considered. In order to simplify these calculations, the current is assumed to vary linearly over the switching period.

The reduction in switching loss achievable with the proposed modulation scheme is determined by the rate of change of current, which determines the difference in current at the switching points between the pure and discretised reference values. This is determined by the fundamental frequency of the current waveform. The rate of change is given by the time derivative. Assuming a sinusoidal waveform, the angle derivative is a cosinusoid, and when the absolute value of this is taken, the mean value can be calculated as

$$\overline{\left|\frac{di}{d\theta}\right|} = \frac{2}{\pi}\ \mathrm{A}\cdot\mathrm{rad}^{-1}. \tag{5.18}$$

The mean time derivative is then determined by the frequency, giving

$$\overline{\frac{di}{dt}} = \frac{2}{\pi}\cdot 2\pi f = 4f\ \mathrm{A}\cdot\mathrm{s}^{-1}. \tag{5.19}$$

The switching losses for each device are given in (5.9) and (5.10). With 240 V across the devices, the switching energies in terms of device current are:

$$E_{\mathrm{sw}Q} = 3.92\times 10^{-5}\cdot i \tag{5.20}$$

$$E_{\mathrm{sw}D} = 5.17\times 10^{-3}\left(\frac{i}{400}\right)^{0.6}. \tag{5.21}$$

The reduction in switching loss over a range of currents and fundamental frequencies is given in Figure 5.7 for a resolution of 75 and a switching frequency of 10 kHz. The calculation is based on three switching events for each type of device in each switching period, which is valid since there are two switching events—one on, one off—for each phase in each period, and on average half will involve IGBTs and half diodes.

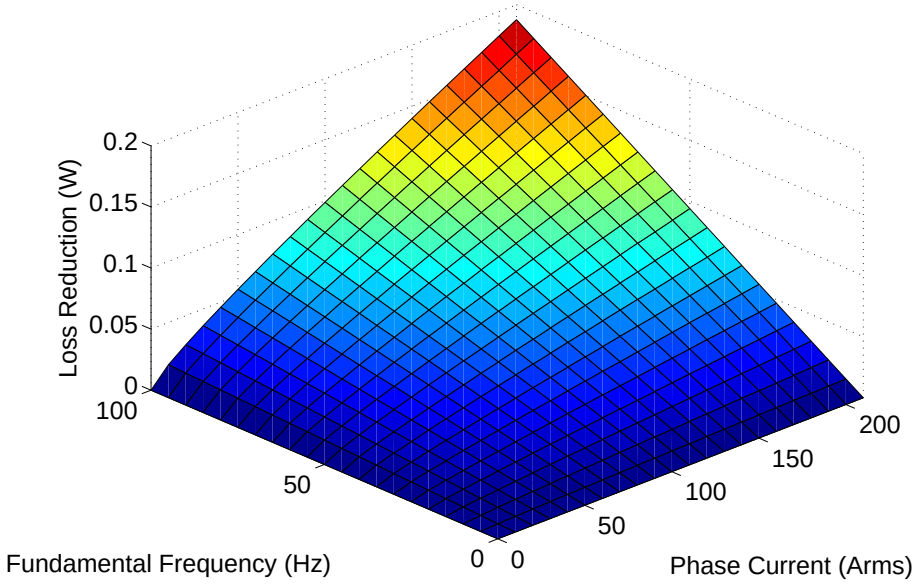

Figure 5.7: Relationship between switching loss reduction, current and fundamental frequency

The variation of the reduction in switching loss with resolution is shown in Figure 5.8 for a fundamental frequency of 50 Hz.

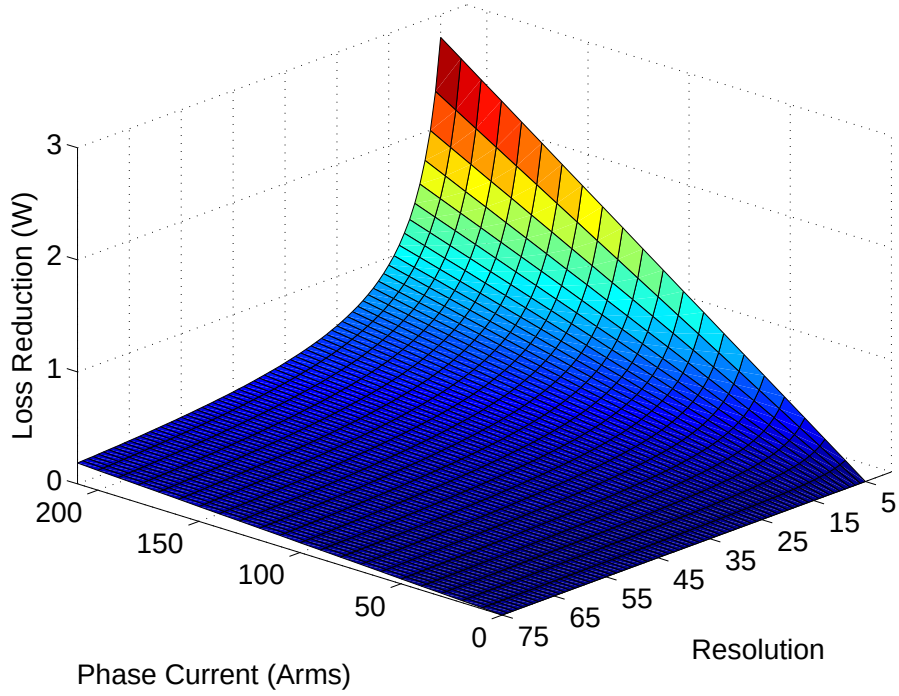

Figure 5.8: Relationship between switching loss reduction, current and resolution

The potential switching loss reduction over the tested operating range was found using the same experimental data used to produce Figure 4.12. At each point the rate of change of current over the switching period is determined. The IGBT and diode switching losses are calculated using the measured current at the start of the switching period, and again at a point one discretisation period, i.e. T_{pwm}/r, away, with the current calculated from its value at the start of the period and its rate of change. For the diode switching loss, this is an approximation, since the switching point may not be at the beginning of the period, and the absolute value of the current is relevant since the diode switching loss expression contains a term in which the current is raised to a power. The IGBT switching loss is directly proportional to current, and so only the difference in current between the two points is strictly necessary.

This is repeated for all phases and the minimum loss found for each phase. This is compared with the maximum for all phases to give the potential loss reduction. This may not be achievable in practice, since the switching combination associated with the calculated loss reduction may not be available, but it serves to illustrate the scale of potential gains. The calculated loss reduction is plotted over the tested operating area in Figure 5.9 for a resolution $r = 75$.

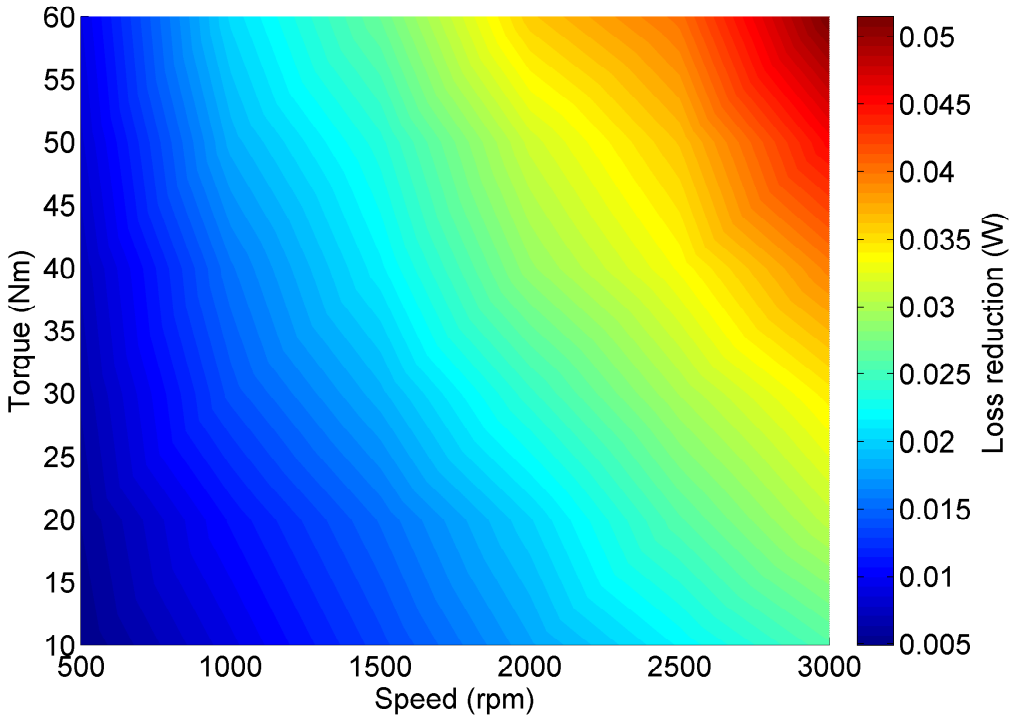

Figure 5.9: Estimated switching loss reduction calculated from measured currents

In order to implement a scheme which distinguishes the variation in switching loss between different voltage vectors in real time, the controller must predict the change in current over the switching period. The simplest way to do this is to take the

previous two points and extrapolate this linearly to the end of the next switching period, as demonstrated by Figure 5.10. During the period between samples $N = 0$ and $N = 1$, the current for the period $N = 1$ to $N = 2$ is calculated by finding the rate of change of current between $N = -1$ and $N = 0$. This is extrapolated forward two sample periods, assuming a constant rate of change of current, giving the current indicated by the dashed blue line. In the next sample period, the samples $N = 0$ and $N = 1$ will be used and the red dashed line extended to $N = 3$. Since the period of the PWM cycle is much smaller than the time constant of the motor, this is a good approximation.

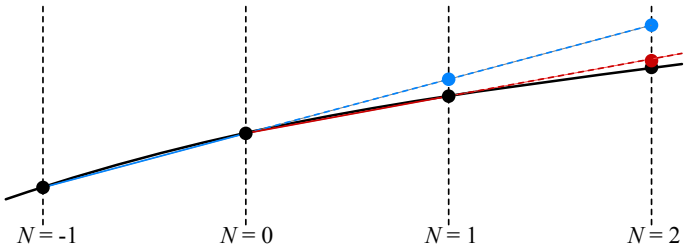

Figure 5.10: Linear extrapolation of current

5.2.3 Motor Loss

It has been described previously how the assumption of constant current over a switching cycle is sufficient to discern a difference in conduction loss between two possible voltage vectors, whereas to discern a difference in switching loss, the variation of the current over the period is required. In order to determine the effect of choosing between two different voltage vectors on the motor losses, the variation of current within the period must be determined, such that the difference in mean current over a single cycle is revealed.

To achieve this would require the solving, in real-time, of an equivalent circuit model to determine the effect of the voltage on the current, which would need to be performed for each possible voltage vector. This would require a large computational effort, and would likely require a more powerful processor than would usually be found in this application, which may offset the advantages somewhat.

For this reason it is concluded that it is not practical to include the motor losses in a real-time model for the purposes of detecting changes resulting from the application of different voltage vectors.

5.2.4 Conclusion

Whilst the proposed motor loss model can be demonstrated to represent the actual loss reasonably well, as discussed above, the complexity of its implementation compared to the potential reduction in loss precludes it from being realised in the practical loss model.

It is clear from Figure 5.7 and Figure 5.8 that even at very low resolutions, high currents and high fundamental frequencies, the gain in efficiency achieved by considering the switching loss is minimal. Although still small in absolute terms, Figure 5.6, when compared with Figure 5.9, shows that the possible reduction in conduction loss is an order of magnitude greater than the possible reduction in switching loss.

Therefore, the final practical implementation of the loss model-based control scheme will consider only the effects of inverter conduction loss in order to realise the optimum balance between efficiency gain and control complexity.

5.3 Practical Results

The discretised modulation scheme developed in Chapter 3 is implemented on the test platform described in Chapter 4. The loss model is implemented with the components as described above, with the result fed into the cost function along with the vector errors.

The efficiency of the scheme under each set of test conditions is measured over the same operating range as previous results: 500–3 000 rpm and 10–60 Nm. The results presented here show the efficiency of the whole drive system, from the DC link to the motor shaft.

As has been demonstrated previously, it is possible to measure the power at the inverter output/motor input, giving separate inverter and motor efficiencies. In an ideal situation, these efficiencies would be measured to show the effect of the new control scheme on both, as the inverter efficiency is of particular interest since it is the inverter conduction loss which is being minimised. However, as can be seen from the analysis of the measurement uncertainty in Section 4.4.3, the measurement of the drive efficiency directly from the DC and mechanical power is much more accurate than the measurement of the inverter and motor efficiency, which rely on the inherently inaccurate 3-phase power measurement. Therefore, all efficiency measurements in this section show total drive efficiency.

Using the total drive efficiency has the advantage of showing any effect the control scheme may have on the motor efficiency, as well as the inverter efficiency. For example, if reducing the loss in the inverter causes an unforeseen rise in motor loss,

this will be seen. Unfortunately, it will not be able to distinguish this from a change in inverter efficiency.

The measured efficiency under the discretised modulation scheme with a resolution of 750 with the conduction loss model is shown in Figure 5.11.

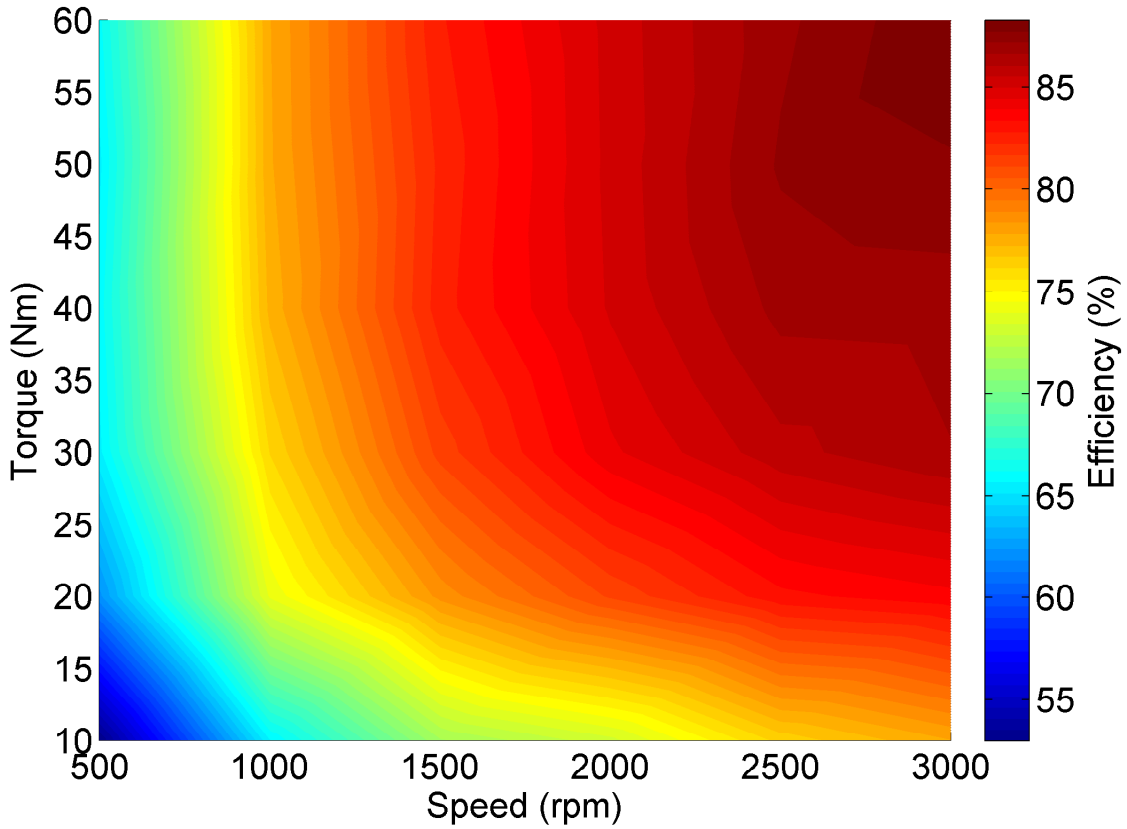

Figure 5.11: Map of drive efficiency with discretised modulation scheme with $r = 750$ and conduction loss model

Comparing Figure 5.11 and Figure 4.13 does not reveal any major differences or trends in efficiency. The difference in efficiency between the new control scheme under these conditions and the conventional control scheme, the efficiency of which is shown in Figure 4.13, is plotted in Figure 5.12, in order to more clearly show the differences. Note that the figures shown are percentage points, relative to the original efficiency. In other words, the difference between an efficiency of 70% and 80% is shown as 10%, rather than $\frac{10}{70} = 14.3\%$. A positive value represents a gain in efficiency relative to the condition being compared, in this case the conventional control scheme.

Figure 5.12 shows a fairly consistent gain in efficiency of up to almost 3% at lower torques. This falls off with increasing torque, eventually showing a reduction in efficiency at higher torques, particularly at low speed.

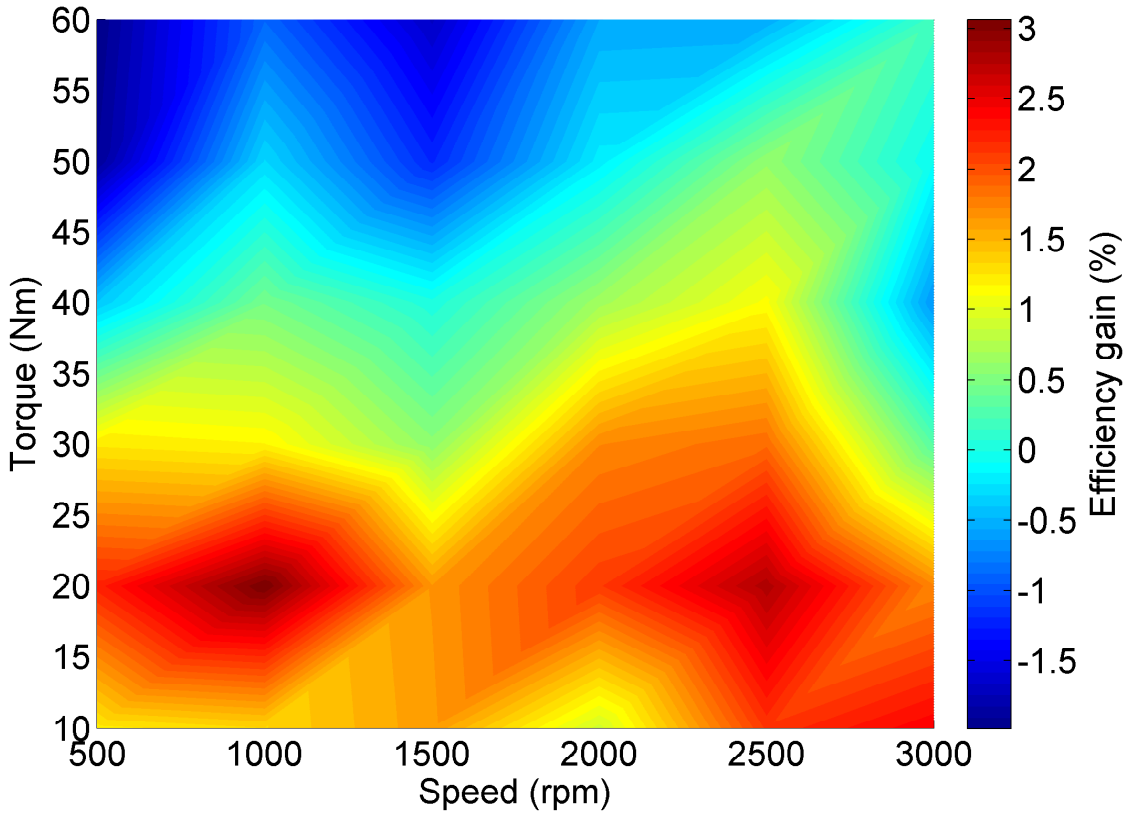

Figure 5.12: Map of drive efficiency gain with discretised modulation scheme with $r = 750$ and conduction loss model, relative to conventional control

The torque ripple at each operating point under this control scheme is illustrated in Figure 5.13, and shows largely the same trend as the conventional scheme, as shown in Figure 4.14. However, the absolute values are, on the whole, slightly lower with the discretised modulation scheme. The variation with speed is also smoother.

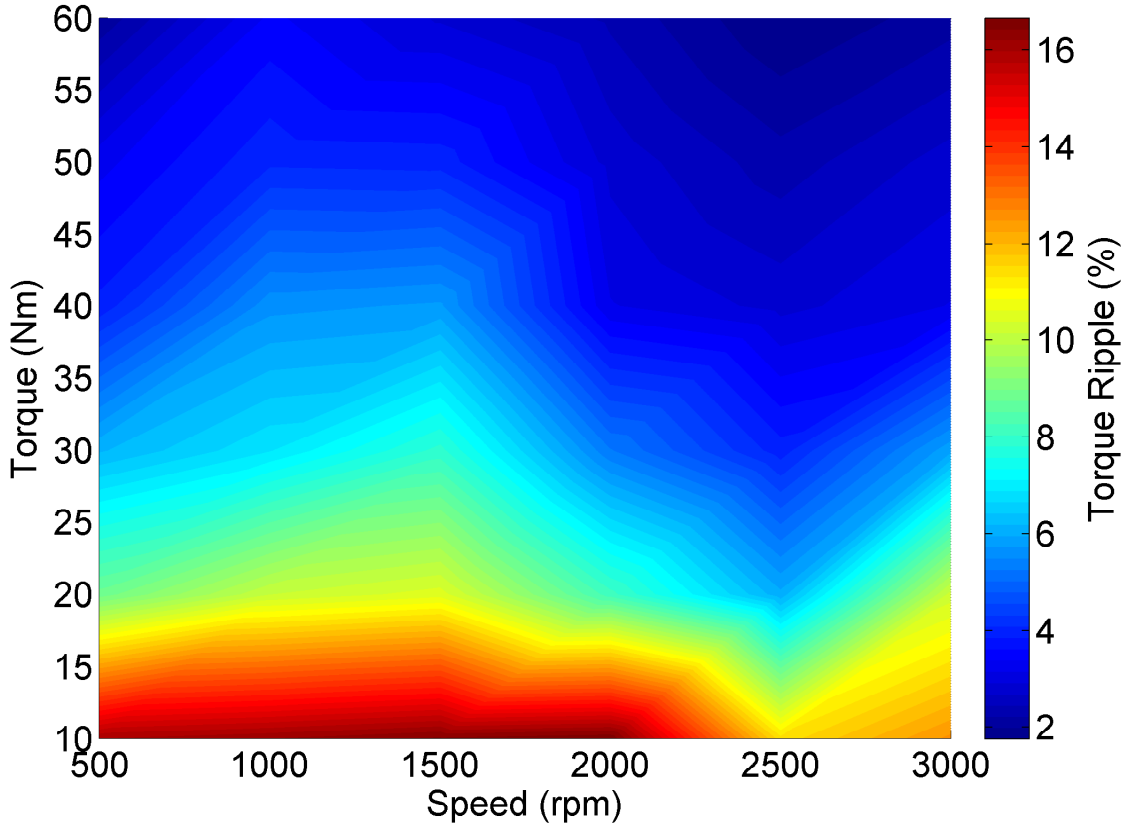

Figure 5.13: Map of torque ripple with discretised modulation scheme with $r = 750$ and conduction loss model

5.3.1 Validation

After examination of the practical results presented above, further simulation work was undertaken to explain the effects seen.

Switching loss dominates over conduction loss at low current levels since the energy lost in the devices during commutation is significant even at low current. At higher currents, the conduction loss becomes more significant. Since the proposed loss reduction scheme targets conduction loss, it would be logical to expect greater efficiency gains at high torque, but this is not the case: in fact the opposite is true.

Additional simulation results show minimal change in conduction loss between the conventional and proposed controllers, in agreement with the predictions in Sec-

tion 5.2. They also show minimal change in overall drive efficiency. However, the machine model used does not account for iron loss. It is suggested that a change in waveform quality may account for the difference in efficiency between the two schemes.

An increase in the harmonic distortion of the current causes additional iron losses [78]. Figure 5.14 shows a simulated comparison of the THD of the phase current with discretised modulation ($r = 750$) and the loss model ($A = 1$) over the whole tested operating region. As with the efficiency maps, the difference in THD is shown in percentage points. A gain in THD is shown at high torque, whilst a reduction is shown at low torque. This correlates with the efficiency changes seen, and suggests that these are caused by variation in iron loss due to distortion of the current waveform.

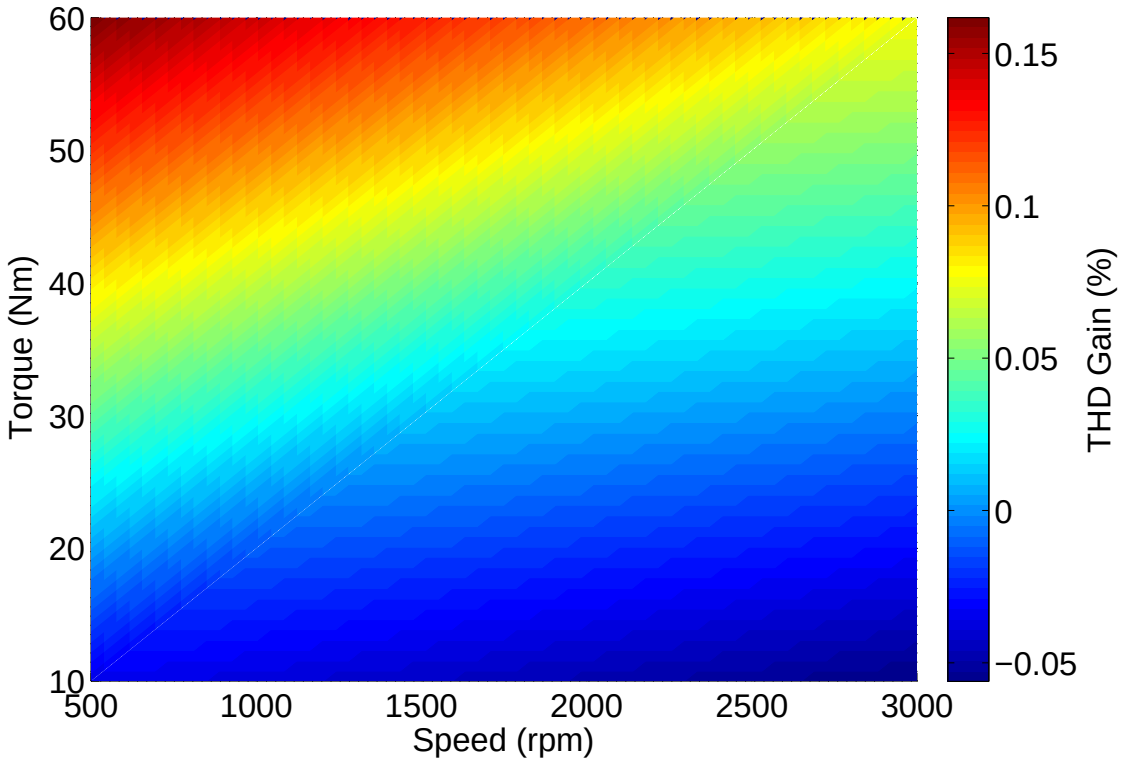

Figure 5.14: Map of simulated THD of discretised modulation scheme with $r = 750$ and conduction loss model, relative to conventional control

The simulation results presented in Chapter 3 show the relationship between resolution and THD; distortion is increased with decreasing resolution. Since increasing the priority of the loss model in the cost function (i.e. decreasing the weighting factor) will result in the applied voltage deviating further from the demand, this will increase THD and further decrease efficiency by causing increased iron loss.

Detailed loss analysis is shown for four different operating points in Table 5.2. These results show that the conduction loss is reduced by a small amount in line with the predictions presented in Section 5.2. The switching loss and motor losses are

approximately the same, since the controller does not have a significant impact on these. The motor losses do not include the iron loss, as noted above, and thus these results do not provide a true picture of the whole drive efficiency.

Table 5.2: Simulated loss breakdown at different operating points under conventional and discretised modulation

Modulation Scheme	Speed (rpm)	Torque (Nm)	Motor loss (W)	Cond. loss (W)	Switch. loss (W)	Inverter loss (W)	Total loss (W)
Conventional	500	10	141.98	199.14	809.27	1008.41	1150.40
Discretised[1]			145.28	199.10	809.52	1008.62	1153.90
Conventional		60	615.59	300.02	1110.04	1410.07	2025.65
Discretised[1]			618.09	299.93	1109.80	1409.73	2027.82
Conventional	3000	10	56.10	122.86	529.15	652.02	708.11
Discretised[1]			66.28	122.79	529.16	651.95	718.23
Conventional		60	602.48	365.14	1251.41	1616.55	2219.03
Discretised[1]			645.55	364.69	1251.37	1616.07	2261.62

[1] $r = 750$, $A = 1$

The simulated torque ripple under conventional modulation (Figure 5.15) and discretised modulation with the loss model (Figure 5.16) show broadly similar results to the practical tests. Comparing the two schemes shows that the inclusion of the loss model increases torque ripple by a mean of 1.8 %-points over the whole operating area. The only major difference is at high speed and low torque, where the simulation shows that the torque ripple reduces with increasing speed under conventional modulation, but not under the loss minimisation scheme.

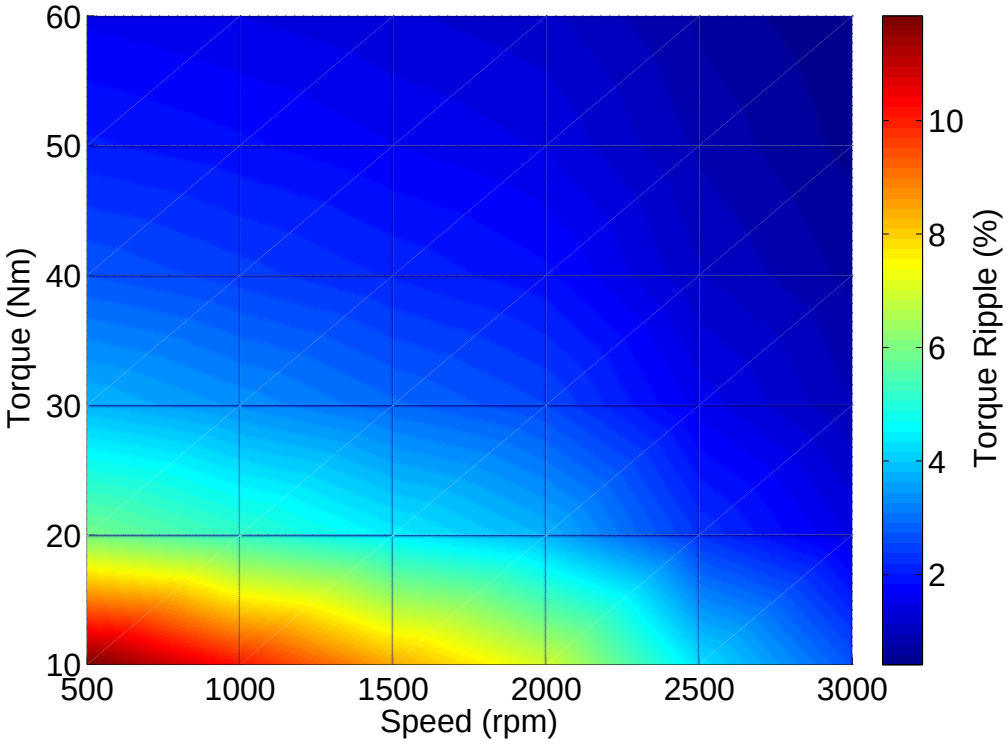

Figure 5.15: Map of simulated torque ripple under conventional control

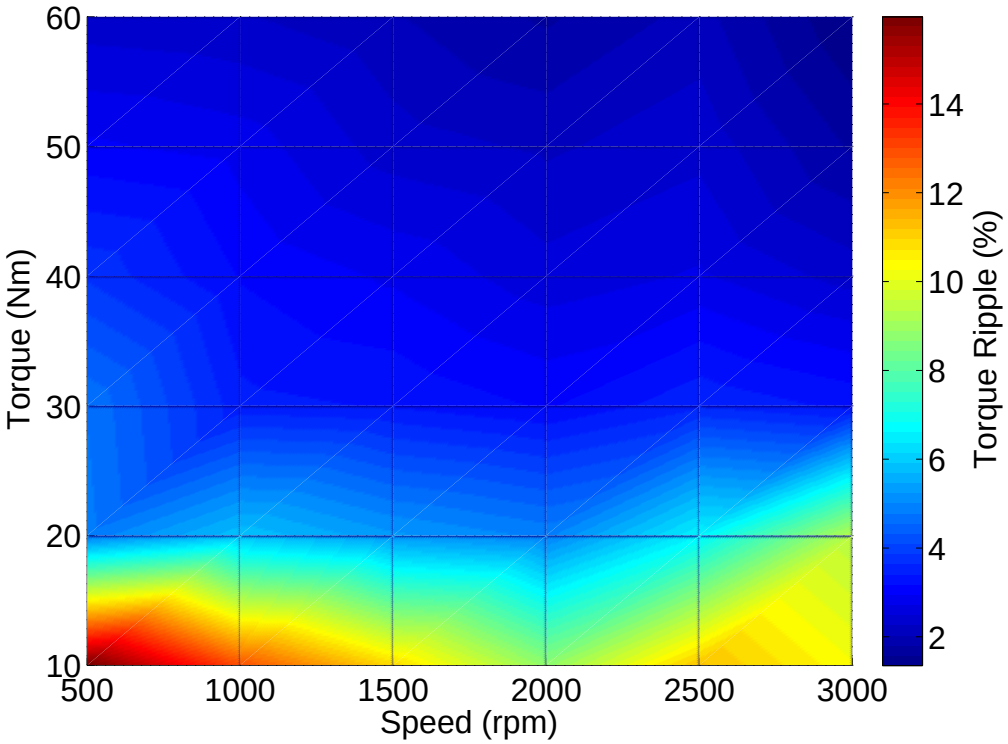

Figure 5.16: Map of simulated torque ripple under discretised modulation scheme with $r = 750$ and conduction loss model

5.3.2 Effect of Resolution

A comparison of efficiency under discretised modulation with a resolution of 75 with the conduction loss model against the conventional control scheme is shown in Figure 5.17. This shows broadly the same relationship exists between the discretised and conventional control schemes as with a resolution of 750. However, it also shows a small but consistent reduction in efficiency compared with the results for a resolution of 750, contrary to the expected behaviour.

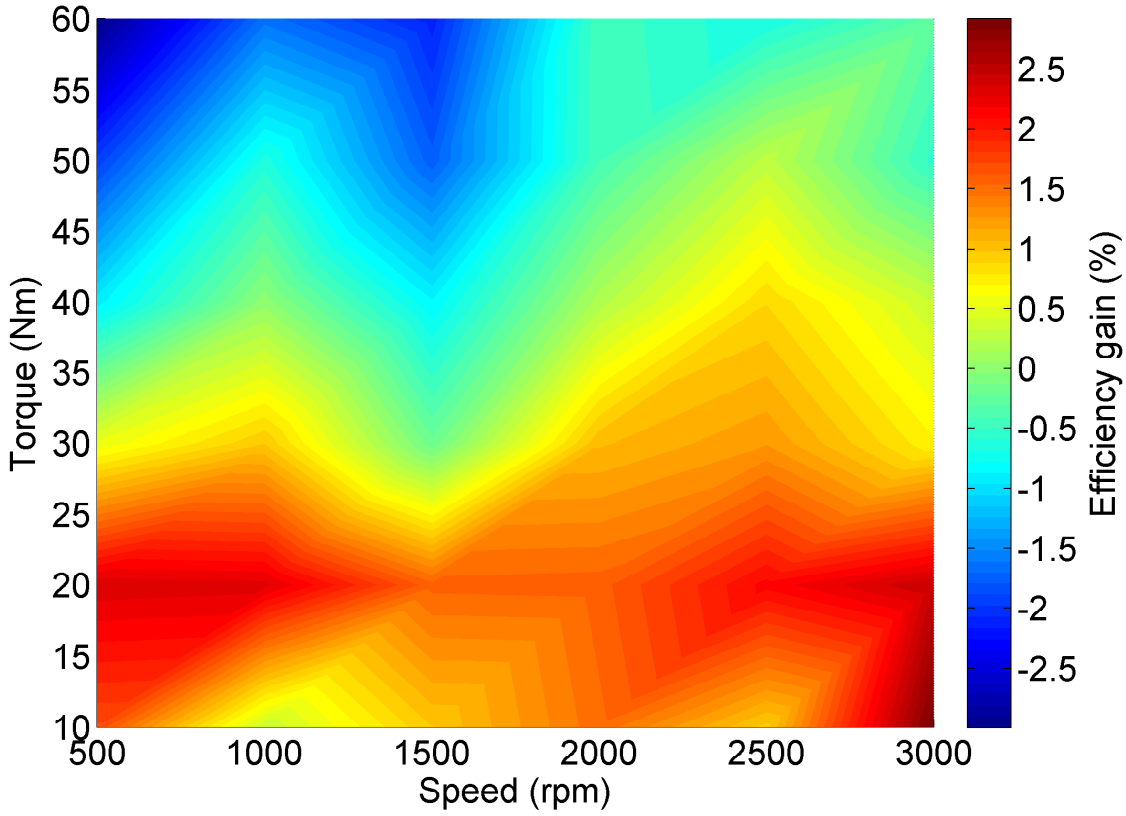

Figure 5.17: Map of drive efficiency gain with discretised modulation scheme with $r = 75$ and conduction loss model, relative to conventional control

Torque ripple under this scheme is shown in Figure 5.18. These results show a very similar torque ripple characteristic to the conventional and higher resolution cases, but there is a significant increase towards the lower torque region where the highest torque ripple is found, of a few percentage points. This corresponds to what is suggested by the simulation results in Section 3.2.6.

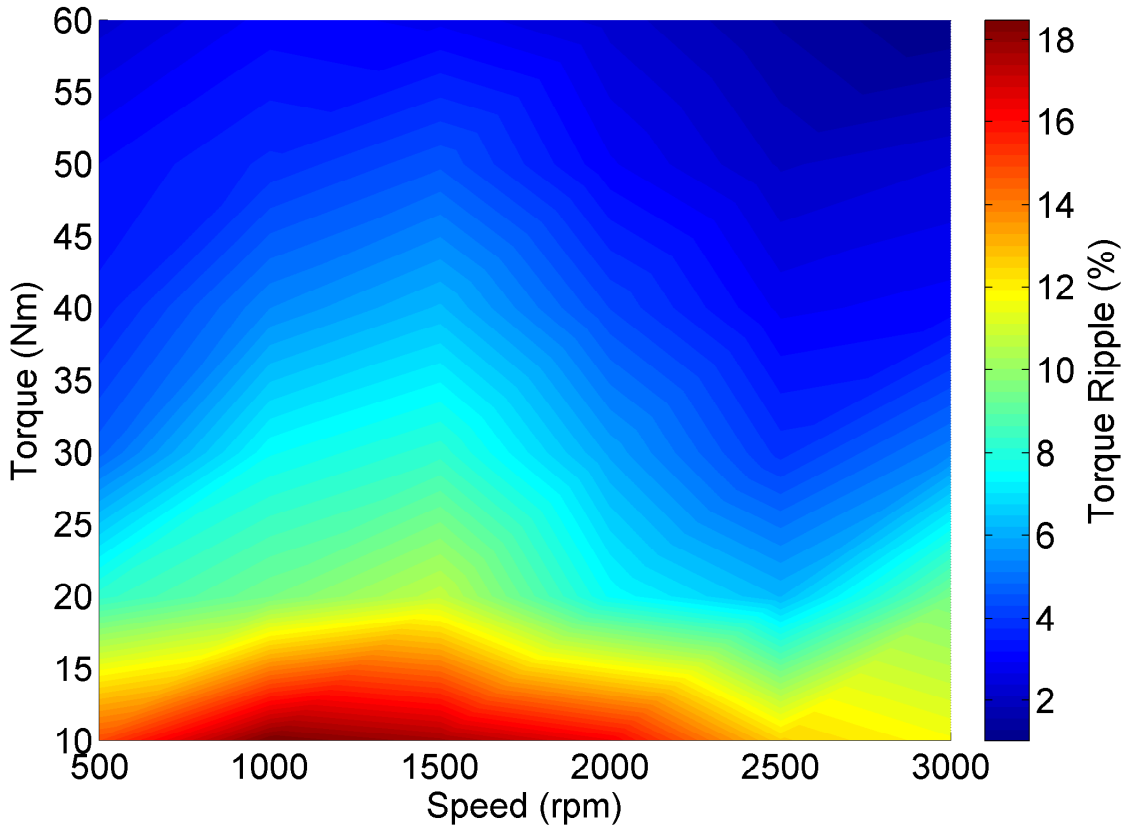

Figure 5.18: Map of torque ripple with discretised modulation scheme $r = 75$ and conduction loss model

Figure 5.19 shows the effect of resolution on efficiency gain at three different operating points with $A = 1$. All operating points show a steep drop in efficiency gain for resolutions lower than 75. This is contrary to the theoretical conduction loss reduction for low resolutions calculated and shown in Figure 5.5, where the efficiency gain increases with decreasing resolution indefinitely.

However, the theoretical calculation does not take into account any efficiency reductions that may occur in the motor when very low resolutions are used. It is inevitable that there is a point below which these efficiency reductions outweigh the reductions made in conduction loss, and these results show that point to be around $r = 75$.

Above this point, the relationship between efficiency gain and resolution is unclear, and varies between operating points.

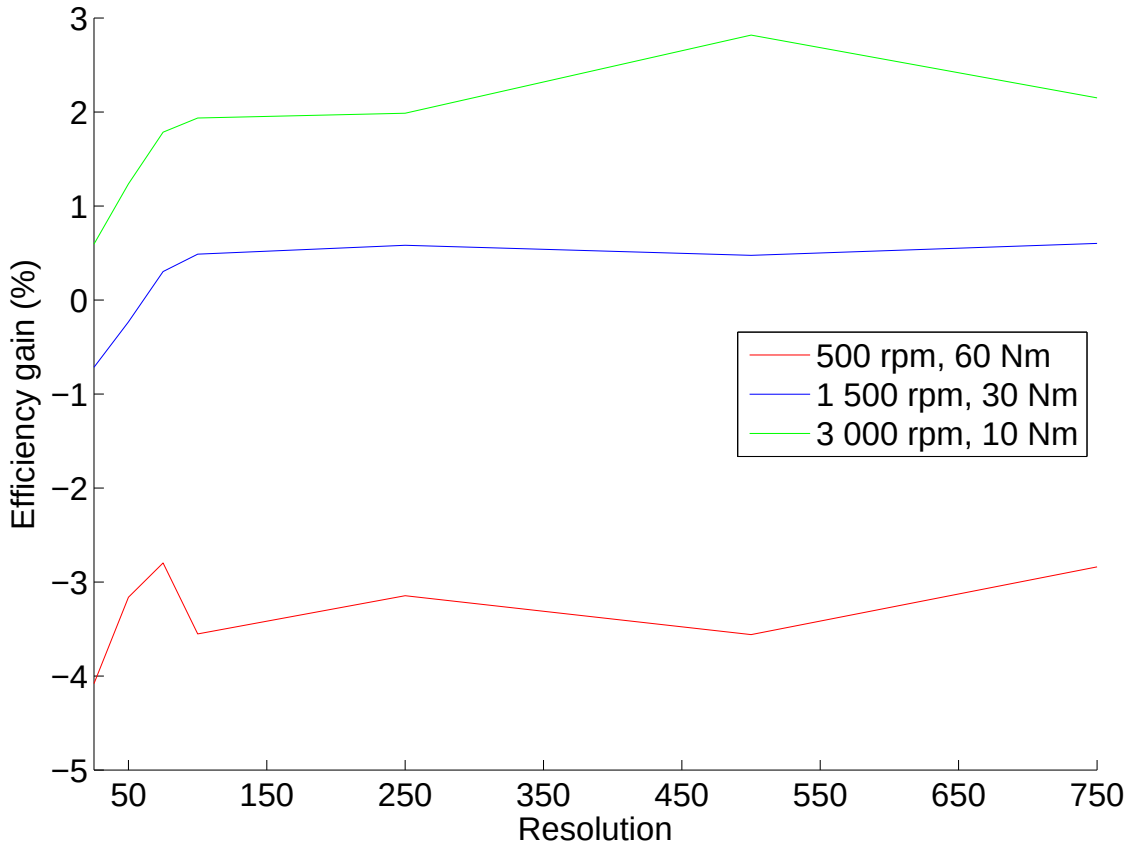

Figure 5.19: Relationship between resolution and efficiency gain for various operating points

5.3.3 Effect of Weighting Factor

The comparison of efficiency between discretised modulation with $r = 75$ and weight $A = 0.5$ and conventional modulation in Figure 5.20 shows a similar trend to the other discretised modulation tests, with efficiency gains at low torque and reductions at higher torque.

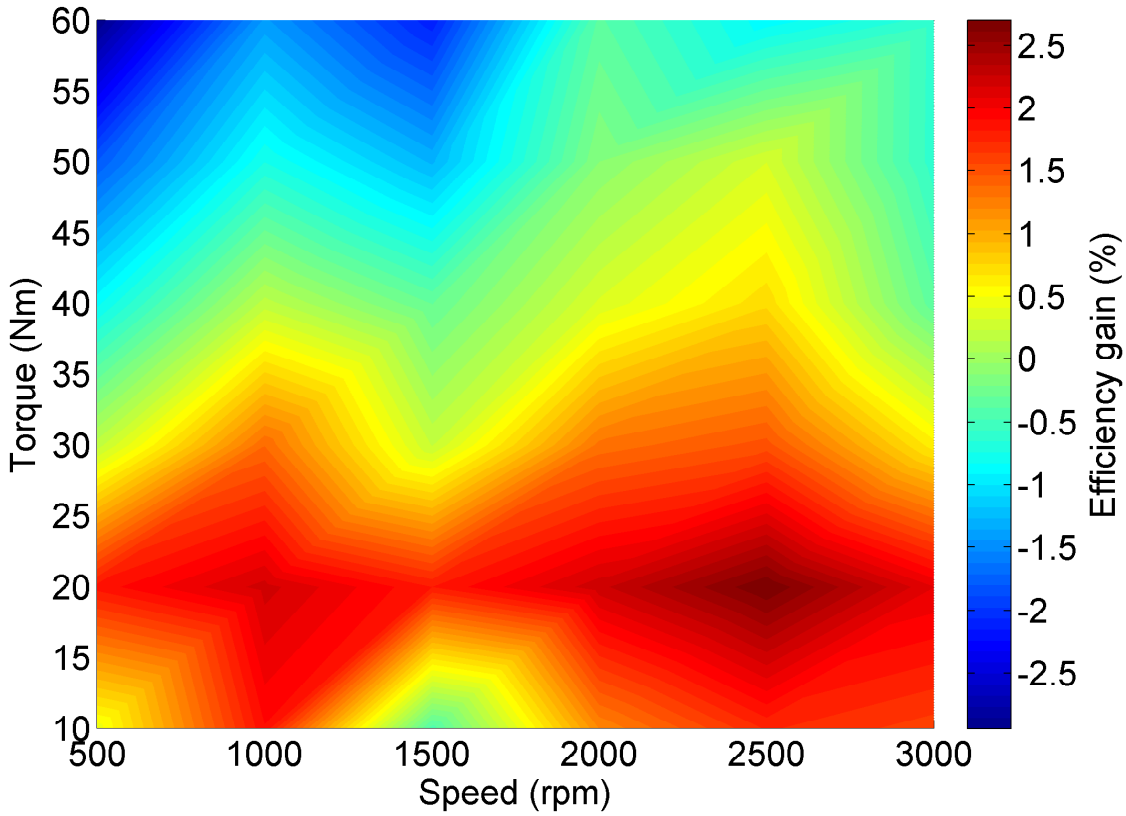

Figure 5.20: Map of drive efficiency gain with discretised modulation scheme with $r = 75$, $A = 0.5$ and conduction loss model, relative to conventional control

Comparing the efficiency of the discretised modulation scheme with $r = 75$, $A = 0.5$ with the same modulation scheme with $r = 75$, $A = 1$, as in Figure 5.21, shows that increasing the weighting of the loss model in the cost function has, overall, a neutral effect on the drive efficiency. Most points on the efficiency map show between 0.5 and -0.5 % points of change, with a few outlying points showing greater gains and losses. It is difficult to draw any firm conclusions about the effect of weighting factor from these results and so a more in-depth study is presented below.

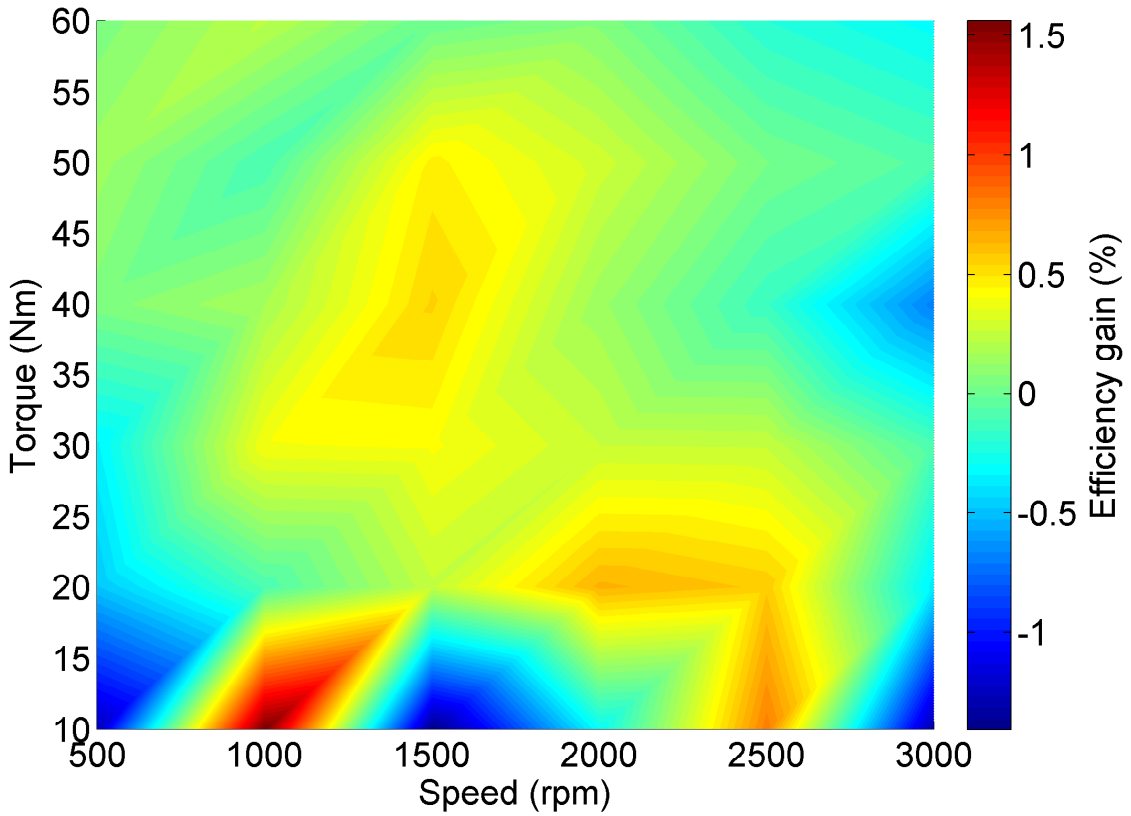

Figure 5.21: Map of drive efficiency gain with discretised modulation scheme with $r = 75$, $A = 0.5$ and conduction loss model, relative to the same scheme with $A = 1$

Comparing the map of torque ripple in Figure 5.22 with the same results for $A = 1$ shows that the increase in weighting factor causes a slight decrease in torque ripple which is particularly noticeable towards the area with greater torque ripple, but shows the same profile as the other schemes.

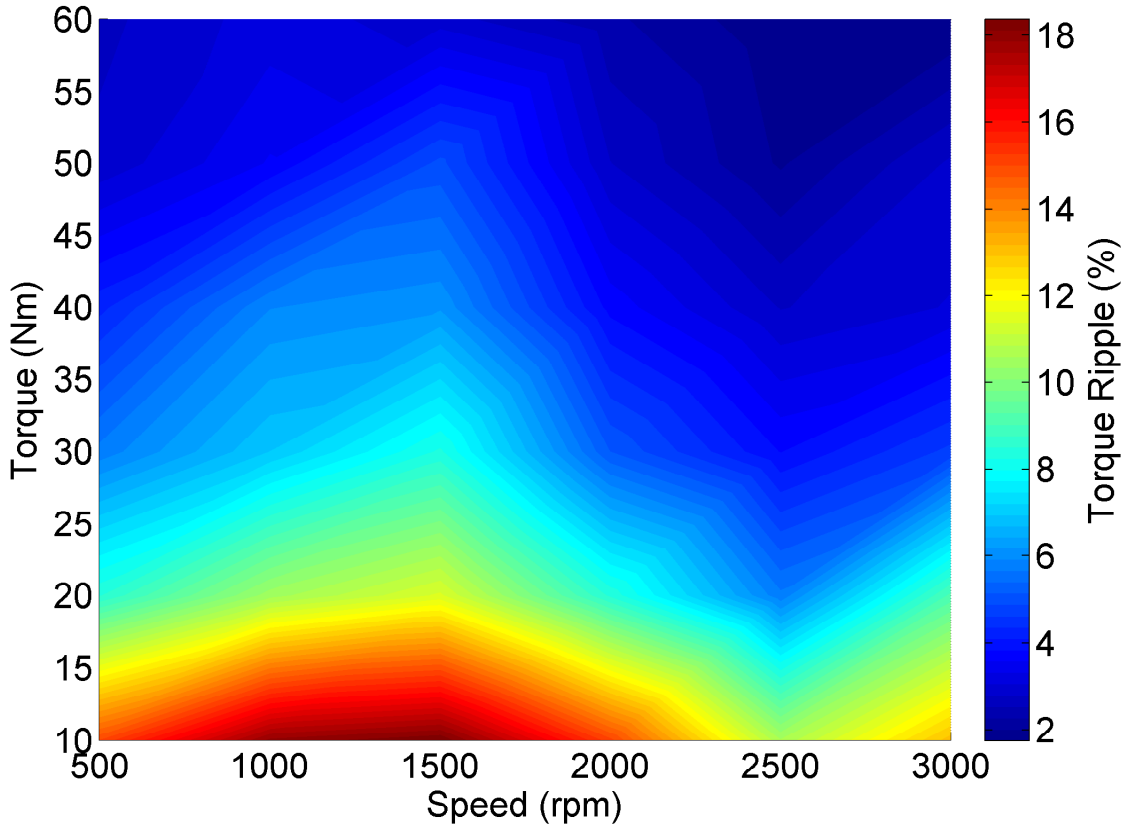

Figure 5.22: Map of torque ripple with discretised modulation scheme with $r = 75$, $A = 0.5$ and conduction loss model

The study of the effect of weighting factor on the efficiency gain in Figure 5.23 shows that maximum efficiency gain is achieved using a weighting factor in the range 0.5–2 across all three tested operating points, with a resolution of 75. This appears to vary with the absolute efficiency gain, and hence the operating point on the speed-torque plane, having a proportional relationship with torque, and hence current. Two of the three tests exhibit a significant dip in efficiency gain for a weighting factor of 5, compared with both the maximum at lower weighting factors and the local maximum at 10. This implies that including the loss model with insufficient priority (a weighting factor of 5) actually reduces the efficiency of the system, compared with conditions where it has very low priority (a weighting factor of 10), before an overall improvement is seen at lower weighting factors. This trend is echoed in the third test—500 rpm, 60 Nm—albeit with the salient points shifted to lower weighting factors.

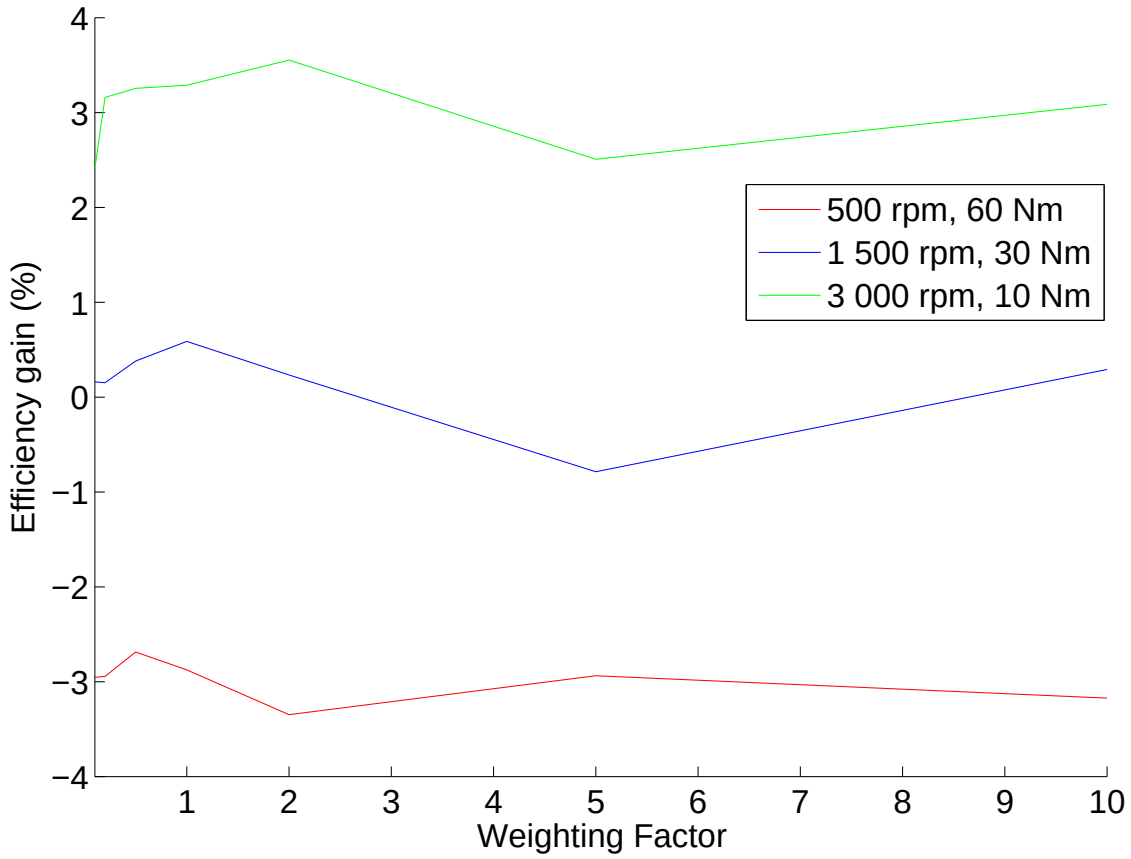

Figure 5.23: Relationship between weighting factor and efficiency gain for various operating points

5.4 Computation Time

As stated in the full description of the test platform in Section 4.3.3, the control algorithm is implemented on 16-bit floating point DSP with a clock speed of 150 MHz. At the PWM frequency of 10 kHz used during experimentation, the processor has 100 µs in which to perform all the required processing: read the currents from the ADCs, implement FOC to produce a voltage demand vector, discretise and perturb this demand, calculate vector errors and losses associated with each permutation, find the optimum and update the registers with the corresponding reference values.

Sufficient overhead must also be allowed for the performance of supplementary operations, such as communication with the PC for control and feedback purposes, which occur asynchronously, but nevertheless, must be performed without disrupting the operation of the main control routine.

Execution times for each process in the discretised control scheme with the loss model implemented were measured using one of the CPU timer peripherals on the DSP which allow timing of operations by recording the number of clock cycles elapsed. These times, and selected combinations thereof, are shown in Table 5.3 in order of the task's execution in the PWM interrupt service routine.

As can be seen, the conventional approach, using only FOC and SVM requires only 25.3 µs to execute on this processor, including reading the ADCs and loading the registers: approximately a quarter of the time available for the specified switching frequency. Apart from the supplementary tasks as mentioned above, the processor would otherwise be idle for the rest of the PWM period.

The discretised modulation method effectively adds 27.8 µs to the required time, plus some time to execute the cost function. With a deviation of one, giving eight permutations, the cost function takes 3.1 µs to execute, but this would be less if the results of the loss model were not included.

The loss model itself requires 21.8 µs to execute. A large part of this, 15.9 µs, is due to the need to scale the results to apportion the weight of the losses and vector errors, which is time-consuming since it inevitably involves floating point division, but is unavoidable.

Hence the use of the discretised modulation scheme and the inclusion of the loss model requires a total of 49.6 µs of additional processing time in each switching period. This is twice as long as the conventional scheme, and represents half the total switching period, but still fits comfortably within the time available and leaves approximately 25% of that time free for overheads. Thus the proposed control scheme is considered to be suitable for implementation with switching frequencies up to at least 10 kHz.

Table 5.3: Computation times for tasks in the algorithm on the chosen test platform

Task		Time (µs)			
Read ADCs		7.8			
2 phase to 2 axis current calculations					
Read rotor angle					
Calculate slip angle and scale					
Calculate field angle					
Transformation: s-frame to e-frame	FOC	14.1	21.9	25	
Current error calculations					
2-axis PI control					
Transformation: e-frame to s-frame					
Limit voltages					
Transformation: Cartesian to polar					
SVM		3.1			74.9
Discretise and perturb					
Formulate permutations		10.1			
Calculate effective switching times	Discretised	3.1	27.8		
Saturate negative switching times	Modulation	2.7			
Calculate effective vectors		7.1			
Calculate vector errors		4.8		49.6	
Calculate device powers		0.7			
Calculate conduction losses					
Formulate loss permutations	Loss Model	2.1	21.8		
Find minimum and mean, scale losses		15.9			
Evaluate cost function		3.1			
Load registers		0.3			

Chapter 6

Conclusions

This chapter represents a critical analysis of the work undertaken in this project, resulting in a set of conclusions on the validity and impact of the ideas proposed. Suggestions for future work to take these ideas forward are also given.

6.1 Multi-objective Controller Structure

Chapter 3 outlines the reasons why a fundamental change in the structure of the controller is needed in order to achieve control which is influenced by real-time evaluation of a loss model in the manner proposed.

Three possible methods of achieving this by alteration of the modulator design are proposed, termed single-state modulation, cascaded polar modulation and discretised modulation. These are evaluated separately from the loss model, and each is shown to be capable of driving an IM.

The performance of each scheme is assessed relative to each other and to the conventional scheme, which is taken throughout to be FOC with SVM. It is important to note that the evaluation of these methods in isolation here is not a definitive measure of their performance since they are designed with the sole purpose of being used in combination with another objective (specifically the minimisation of loss). The consideration of another objective may augment the action of the filter in the schemes where this is required, and so may improve their performance compared with operating with vector accuracy as the only goal.

The metrics used for comparison are THD and torque ripple. These represent a good gauge of the quality of the control as they are direct measures of the quality of the electrical output from the inverter and the mechanical output from the motor, with respect to demanded values.

6.1.1 Single-state and Cascaded Polar Modulation

Single-state modulation offers very simple implementation, the filter being the most complex element. The scheme also offers an inherent reduction in switching loss, with a maximum of one switching event per phase in each switching period. Unfortunately, it also produces the most torque ripple and THD. As suggested above, the inclusion of the loss model may have a positive impact on these measures, but this is unlikely to be significant enough to make this scheme suitable for a high performance application such as an EV drive.

Sharing many features with single-state modulation, cascaded polar modulation is also simple to implement, but offers much reduced THD and torque ripple. However, this is still significantly higher than the conventional case. Since only the angle is varied from the conventional scheme, there is limited scope for loss reduction.

Since these two schemes showed lower performance in isolation, they were not investigated in combination with another objective. Whether there is merit in pursuing these further with another control objective could be the subject of future work. In particular, a more thorough analysis of the filter and the different possible design parameters of it, could reveal potential in these techniques. The existence of the filters in itself is a disadvantage of these methods, since the optimal design may well be unique to an individual motor.

6.1.2 Discretised Modulation

The comparison between discretised modulation, without the loss model, and SVM from measurements taken on the practical test platform in Section 4.5.3 shows favourable torque ripple results for discretised modulation at two resolution levels, although there is a small increase in THD.

Whilst there is limited benefit in implementing this control scheme over the conventional, the primary aim is to create a structure which can incorporate additional control objectives, and these results prove that such a structure can be achieved without significantly affecting the quality of control.

6.2 Loss Model Controller

6.2.1 Loss Model

A study of various existing approaches to the monitoring of losses in IM drives has revealed a simple but sufficiently detailed model which is presented in Section 5.1. Validation of this and comparison to practical test results show that it accurately reflects the losses.

An analysis of the potential reduction in loss resulting from including the inverter switching loss in the loss model with discretised modulation suggests that, due to the small variation in the switching point, even at values of resolution which result in high THD and torque ripple, the reduction in loss is negligible.

Furthermore, the rate of change of current needs to be estimated in each switching cycle for the correct switching loss to be calculated, and so the switching loss is not included in the practical implementation of the loss model.

Consideration of the motor loss model shows that even more detailed estimation of the current within the switching period is needed to enable the difference in loss between voltage vectors to be discerned. Therefore it is concluded to be impractical to include the motor loss in the model.

An estimate of the difference in conduction loss resulting from two voltage vectors can be made based on a constant value of current for the switching period, which can be assumed to be the current measured at the start of the previous period. The predicted conduction loss reduction for realistic current and resolution is much larger than the predicted switching loss reduction, but is still small compared with the system output power.

6.2.2 Practical Implementation and Results

In order to achieve a practically-implementable solution which has the potential to reduce loss, some sources of loss are omitted from the final model for the reasons described. Whilst this deviates from the original principle behind the work of considering the whole system to ensure reductions in loss from one source are not cancelled by gains from another, this still offers an opportunity to increase the efficiency of the drive. Hence the model of inverter conduction loss was implemented for with discretised modulation for testing on the practical system.

In practical testing of the discretised modulation scheme with a resolution of 750 and the conduction loss model (see Section 5.3) both increases and decreases are seen in different areas of the operating region; broadly speaking efficiency gain decreases with torque, ranging from a 2%-point decrease to a 3%-point increase, approximately. The results of each test are summarised in Table 6.1 and Table 6.2. Taking a simple mean across the tested operating region shows that on average the efficiency is increased by 0.63%-points, a modest but not insignificant increase.

Table 6.1: Statistical analysis of efficiency and torque ripple measurements for conventional and discretised modulation schemes under different conditions

Modulation scheme	Efficiency (%)		
	Mean	Max	Min
Conventional	77.5	88.1	51.7
Discretised, $r = 750$, $A = 1$	78.3	88.3	52.9
Discretised, $r = 75$, $A = 1$	78.3	87.4	53.4
Discretised, $r = 75$, $A = 0.5$	78.0	87.5	52.1

Table 6.2: Statistical analysis of comparisons of efficiency for conventional and discretised modulation schemes under different conditions

Modulation scheme	Compared with	Efficiency gain (%-points)		
		Mean	Max	Min
Discretised, $r = 750$, $A = 1$	Conventional	0.63	3.06	-1.98
Discretised, $r = 75$, $A = 1$	Conventional	0.32	2.92	-2.99
Discretised, $r = 75$, $A = 1$	Discretised, $r = 750$, $A = 1$	-0.31	0.99	-1.19
Discretised, $r = 75$, $A = 0.5$	Conventional	0.33	2.70	-2.95
Discretised, $r = 75$, $A = 0.5$	Discretised, $r = 75$, $A = 1$	0	1.56	-1.41

All the comparisons of the discretised loss model controller and the conventional controller in Section 5.3 show greater efficiency gains at lower torque. Figure 6.1 shows the distribution of operating points in a typical HEV application over the UDDS, an American standard drive cycle for testing light duty vehicles representing urban driving conditions [79]. This shows that the operating points encountered in this application are concentrated in the region where the greatest efficiency improvements are demonstrated by the new scheme.

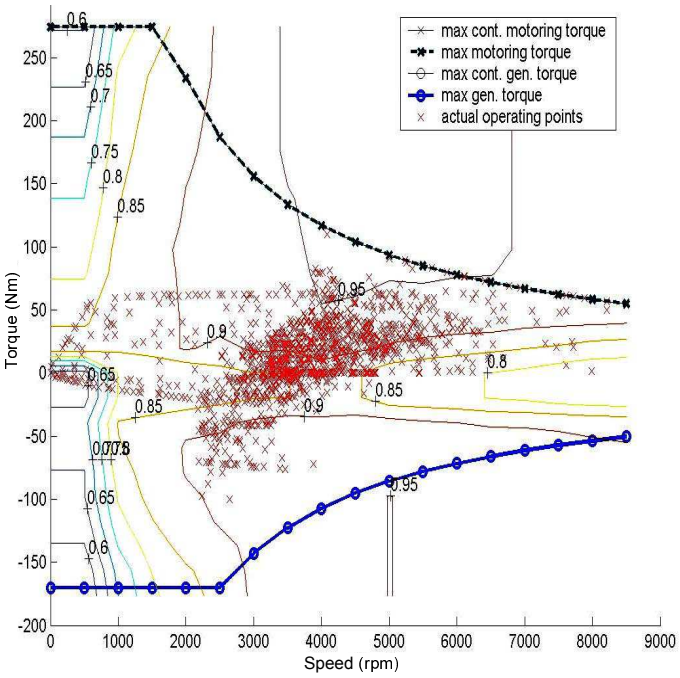

Figure 6.1: Typical operating point distribution in an HEV application [79]

6.2.2.1 Effect of Resolution

Reducing the resolution to 75 reduces the efficiency gain to approximately half that seen at $r = 750$, but still shows a small mean gain in efficiency. The range of effects seen is similar, with maximum gain and loss both at around 3% points. This is in line with the study of various resolution values in Figure 5.19, which shows the optimum resolution to be between 500 and 750, though this varies with operating point, and so requires a compromise.

6.2.2.2 Effect of Weighting Factor

The optimum weighting factor appears to be around 0.5 to 2 from the study shown in Figure 5.23. However, this also varies with operating point and the results in Table 6.2 show that when averaged over the whole tested operating range, the weighting factor does not significantly alter the efficiency.

6.2.2.3 Torque Ripple

Torque ripple measurements, shown in Table 6.3, indicate that the discretised modulation scheme has a significant positive effect on torque ripple, reducing it on average by 20%. There is little variation with resolution or weighting factor. The reasons for this are unclear, but since the goal was simply not to negatively impact performance, an improvement in this metric is welcome.

Table 6.3: Statistical analysis of efficiency measurements for conventional and discretised modulation schemes under different conditions

Modulation scheme	Torque ripple (%)		
	Mean	Max	Min
Conventional	8.3	20.5	2.2
Discretised, $r = 750$, $A = 1$	6.6	16.7	1.8
Discretised, $r = 75$, $A = 1$	6.6	18.5	1.0
Discretised, $r = 75$, $A = 0.5$	6.6	18.4	1.8

6.2.2.4 Discrepancy Between Analytical and Practical Results

The loss reduction calculations in Section 5.2 forecast negligible difference between the discretised and conventional modulation schemes, whereas the practical results presented in Section 5.3 show significant differences at certain points. Two possible reasons are suggested for this:

1. The measured data is inaccurate;
2. The loss is reduced by a mechanism not accounted for in the analytical calculations.

The measurement uncertainty of the total drive efficiency is approximately 4% as calculated in Section 4.4.3.1. This absolute value of typical error therefore varies with the absolute efficiency; lower efficiencies have lower absolute errors. The maximum efficiency of the conventional modulation scheme will typically vary by approximately 3.5%-points in either direction from its measured value of 88.1%, and the minimum efficiency by approximately 2.1%-points. Therefore the uncertainties are comparable to the differences seen, which brings the accuracy of the measurements into question. However, the appearance of the same trends in efficiency under different conditions suggests that the measurements are more reliable than this uncertainty—calculated from stated errors of equipment—indicates.

One mechanism through which the efficiency may be affected by the new control scheme, other than the influence of the loss model, is the difference in the harmonic

content of the supply. Whilst the discretised modulation scheme shows a small increase in THD when tested in Chapter 4, those results were without the loss model and at a single operating point. Further simulation results which include the loss model, presented in Chapter 5, show that the THD trends over the operating region broadly correspond to the trends in efficiency, with a higher THD in areas of decreased efficiency and vice versa. This suggests that the differences in efficiency seen under the new control scheme are caused by variation in waveform harmonic content and hence iron loss.

6.2.3 Computation Times

A thorough breakdown of the processing time required to complete each part of the control algorithm on the device used was generated. This shows that it is feasible to implement the control scheme as proposed on this device, which runs at a clock speed of 150 MHz, with a switching frequency of 10 kHz. This control scheme increases the processing time needed by a factor of approximately three: from 25.3 µs to 74.9 µs. Of the additional time, just over half—27.8 µs—is taken by the discretised modulation scheme, and the remainder by the evaluation of the loss model.

Increasing the deviation D, as described in Section 3.1.5, will significantly affect the computation time, as the number of permutations is equal to $(2D)^3$. Hence for $D = 2$ there are 64 permutations (eight times as many as with $D = 1$), and the discretised modulation and loss model execution times are expected to increase by eight times. To implement a higher deviation, or other changes to the control scheme which require more execution time, a faster processor would be needed, or a lower switching frequency could be used.

6.2.4 Limitations

The inverter loss model, both the conduction and switching components, is reliant on knowledge of some parameters of the power devices. Since these are values which are commonly available from device manufacturers, this does not represent a significant disadvantage, though the scheme is reliant on the accuracy of this data, unless it is independently verified. However, as described in Section 5.1, the parameters vary with temperature, and the datasheet value may not reflect the value at the temperature of the intended application, or give a relationship. The manufacturer may also fit curves to experimental data which gives the best fit over a range of voltages and currents, but is not necessarily accurate at a given point.

Although not included in the implementation, if the motor loss model were included, it would present a similar issue: the motor parameters required, although obtainable

from a standard test, are not readily available for a given machine. These parameters also vary with temperature, although techniques exist to predict these changes [80].

6.3 Future Work

A number of features of the proposed control scheme present opportunities for future research to gain a better understanding of the current system or expand it.

As mentioned previously, the optimum resolution varies with the operating point. This presents the possibility of varying the resolution in real time to track the optimum, perhaps by performing a more detailed study of the optimum at each operating point offline and using a lookup table or fitted function to optimise online. This should attempt to relate the optimum resolution to current, or current and speed, rather than torque, since a torque sensor is not normally used in a real application. A similar approach could be taken with the weighting factor, though since this has less effect on the efficiency, to do so may be considered adding unnecessary complexity.

The concept of deviation in the discretised modulation scheme, the number of discretised levels either side of the demanded reference values which are evaluated, was introduced in Section 3.1.5. Only a deviation value of $D = 1$ was tested, due to the limitation on computation time, as explained above. An investigation of this would be of interest, though it is possible that an analysis would show that the discretised values furthest from the demand minimise the cost function in a large majority of, or even all, cases, indicating that the "compromise" values offered by the intermediate values offer no benefits.

There are also a number of wider points concerning the implementation of the control scheme which could be studied.

Due to practical constraints explained in Section 4.4.1, the full rated operating range of the motor was not tested. It would be useful to determine the effect of the new control scheme at higher speed and torque, and also at negative torque, i.e. in generator mode, since this mode is often used in EVs.

The variation in efficiency gain, and indeed loss, offered by the new scheme over the operating region raises the question of whether the scheme could be implemented on a selected region. The efficiency loss in certain pre-defined regions, which should again be determined by speed and current, could be negated by switching to the conventional control scheme, whilst using the new control scheme where it is shown to give an efficiency gain. A suitable method of switching between the two control schemes would need to be found to ensure smooth operation across the boundary.

Whilst this work has concentrated on induction machines, the methods developed could also be applied to PMSMs, and it would be interesting to see how the results would compare between implementation on the two machine types.

Appendix A

Simulink Models

This appendix shows the proposed modulation schemes described in Chapter 3 as they were simulated in the Simulink software.

Figure A.1 shows the block diagram of the single-state modulation scheme simulation, as described in Section 3.2.3. This structure fits within the FOC scheme illustrated in Figure A.2.

The d- and q-axis voltage demands `v_d_e'*` and `v_q_e'*` and field angle `theta_e'`, all expressed in the e'-frame, are inputs to this block. The constants `theta_v(1-6)s` define the angles of the six active fundamental vectors in the s-frame, where 1024 represents 360° corresponding to a lookup table used for sine and cosine functions. These angles are transformed to the e'-frame by the 's ->e" block. The 'Demand Cart -> Pol' block transforms the voltage demand from Cartesian co-ordinates to polar form.

The 'Select Adjacent Vectors' block takes the angle demand and selects the two active fundamental vectors adjacent it in each direction. These are transmitted to the 'Cost Function' block by the signal `AB` to enable it to apply the chosen state.

The angles of the adjacent vectors are passed to the 'Select Applied Vector' block which determines the vector applied in the previous switching period.

This vector is passed to the filter in polar form (the MI is determined as 0 or 1 based on whether a zero or active vector was applied) along with the angles of the demand and adjacent vectors. Section 3.1.3.1 describes the principles of the filter and Section 3.2.3.3 shows its effect on the signals.

The 'Calculate Errors' block calculates errors for each possible next state and passes these to the 'Cost Function' block which finds the minimum.

Values are then selected by the 'Look up Switching States' block which saturate the modulator to enforce the state which minimises the cost function for the next switching period, and these are output in the signal `cmpABC`.

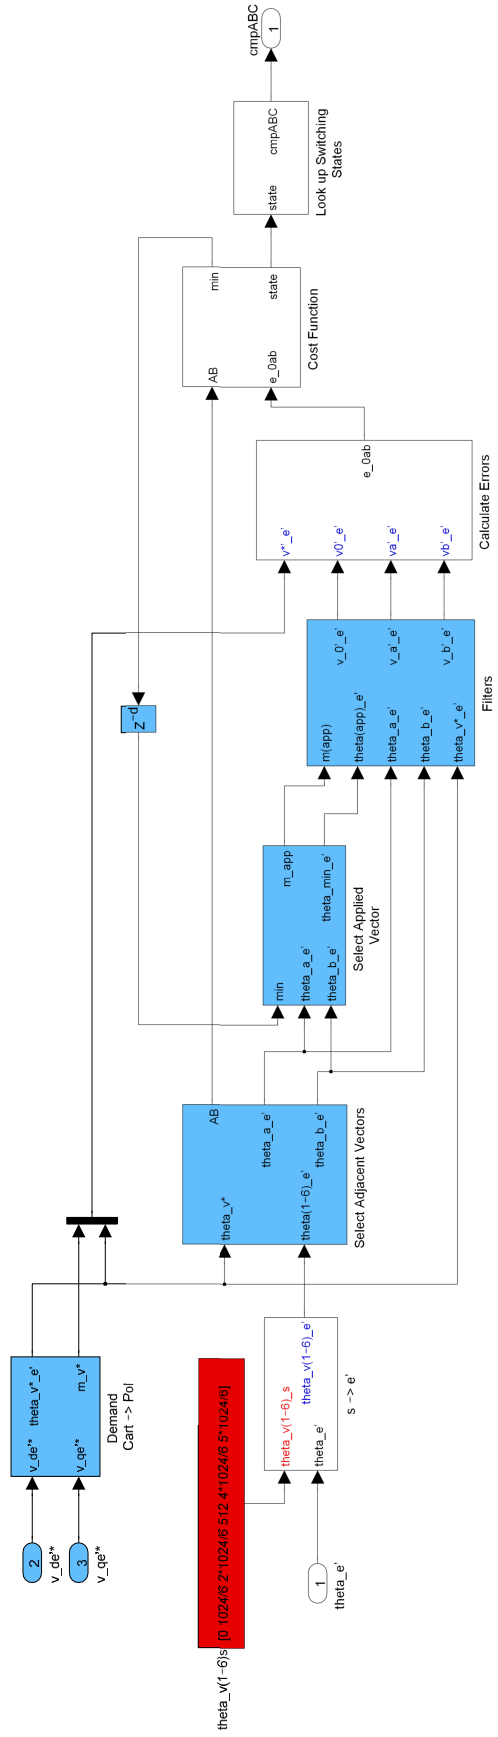

Figure A.1: Outline structure of the single-state modulation scheme as simulated

Figure A.2 shows the block diagram of the FOC scheme as simulated, external to the modulation schemes. An implementation using the phase-shifted reference frame discussed in Section 3.1.3.2 is shown, which is used with the single-state and cascaded modulation schemes. The discretised modulation scheme uses the standard e-frame but is structurally very similar. The operation of this scheme is described in Section 2.1.2.2.

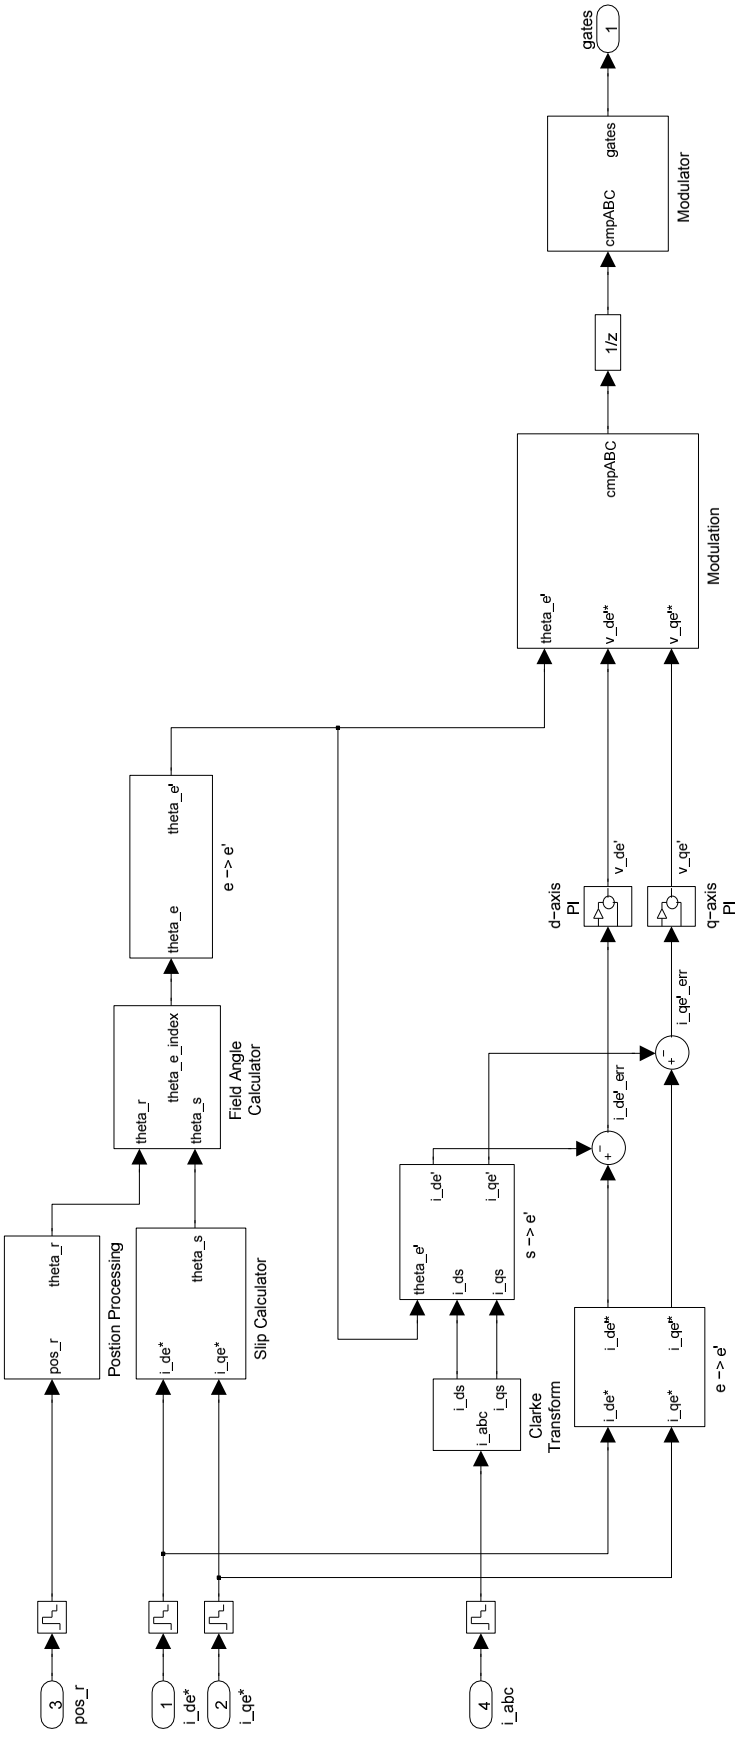

Figure A.2: The FOC scheme in simulation software

Figure A.3 shows the block diagram of the cascaded modulation scheme simulation, as described in Section 3.2.4.

The inputs to the block, vector angle constants and the 'Demand Cart -> Pol,' 's ->e',' and 'Select Adjacent Vectors' blocks are the same as the single-state modulation scheme shown in Figure A.1.

The 'Calculate Times' block calculates the times `T1`, for which the chosen active vector will be applied for, and `T0`, for which the zero vector will be applied for.

Since only the angle of the vector applied in the previous switching period is required, this is the only value found by the 'Select Applied Vector Angle' block.

The previous values of `T0` and `T1` are used to calculate the effective angle applied in the previous period, according to (3.10).

The action of the 'Filter' block is described fully in Section 3.2.4.1.

The 'Calculate errors' and 'Cost Function' blocks are similar to the corresponding blocks in the single-state modulation simulation, except that only active vectors are considered.

The final block generates modulator reference values which will apply the active vector that minimises the cost function for a period `T_1` and the zero vector for the rest of the switching period.

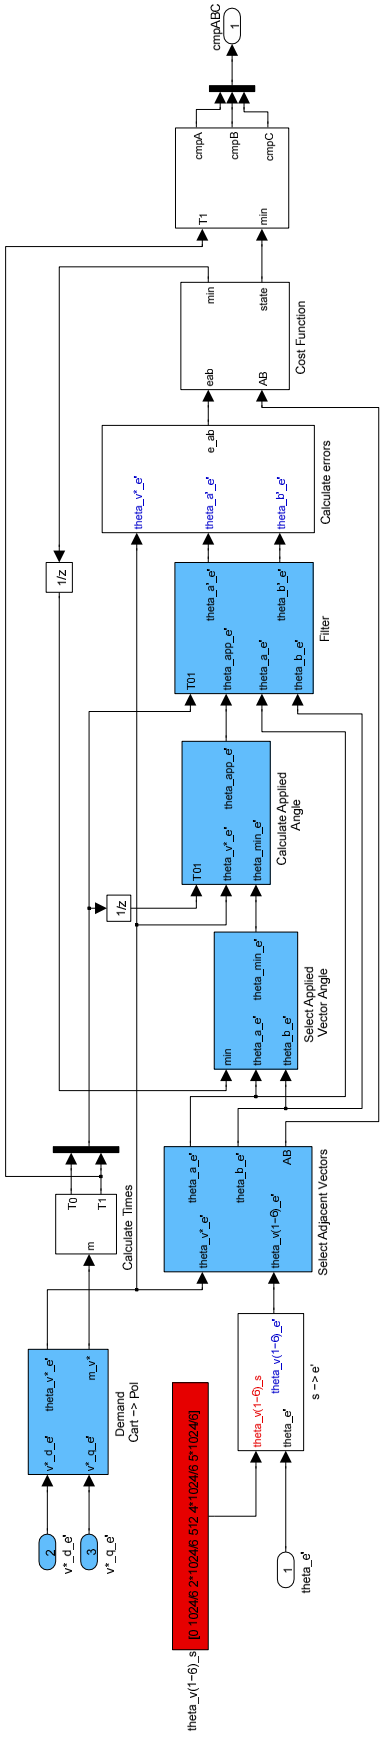

Figure A.3: Outline structure of the cascaded modulation scheme as simulated

Figure A.4 shows the block diagram of the discretised modulation scheme simulation, as described in Section 3.2.5.

This block takes the same inputs as the other two modulation schemes, but expressed in the e-frame.

The demand vector is transformed to the s-frame and expressed in polar form by the blocks 'e ->s' and 'Demand Cart ->Pol' respectively.

The sector in which the demand vector lies and the switching times `T_0`, `T_a`, and `T_b` needed to reproduce it are calculated in the 'Switching Time Calculations' block, according to the method described in Section 2.1.1.2. From these values, the corresponding modulator reference values are calculated by the next block. Up to this point, the simulation represents conventional SVM.

The 'Discretise and Perturb' block produces a set of possible values of each of the three reference values according to the deviation and resolution chosen. For each of these, the switching times, and hence effective vectors, are calculated, in the reverse process of the previous two blocks. If the perturbation of a value causes it to go negative, the `sat` signal for that permutation is set.

The 'Calculate Errors' block calculates the error between each permutation and the demand, both in MI and angle. The error value for a permutation is saturated if its corresponding `sat` signal is set, preventing it from being chosen. The error values are passed to the cost function where the minimum is found, and the corresponding reference values selected.

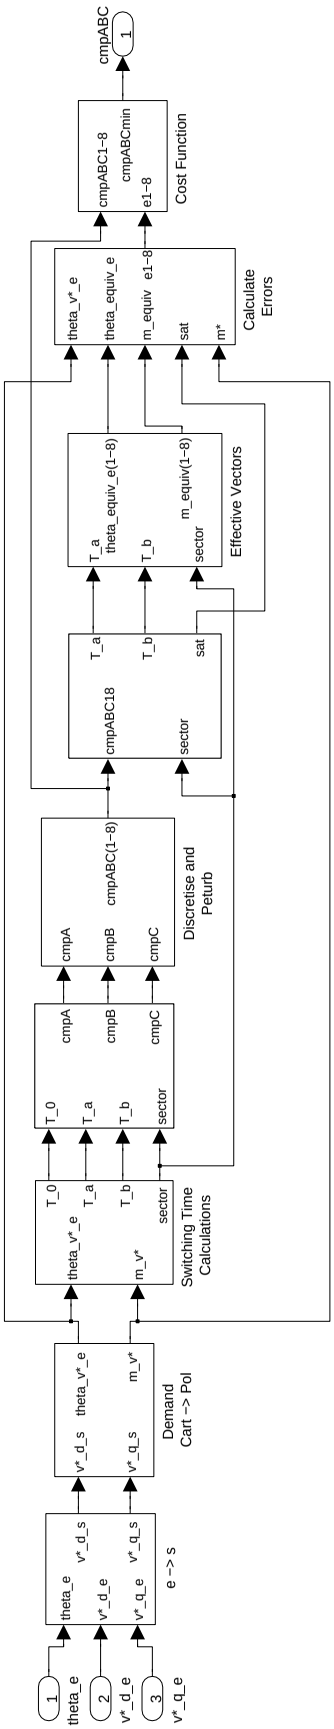

Figure A.4: Discretised modulation scheme as simulated

Appendix B

Loss Modelling

This appendix provides further details of the loss modelling techniques discussed in Section 2.4.

B.1 Inverter Loss Modelling

The models studied are again summarised:

- Physical model [81]
- Simple analytical model [56]
- Complex empirical model [57]

B.1.1 Physical Model

The physical model in [81] uses constants that describe the physical attributes of the devices, as well as the state variables (voltage, current) to calculate the loss.

IGBT Switching Loss

In this model, the IGBT turn-on and turn-off losses are expressed as

$$P_{sw,Q} = \frac{Df_s}{2\pi}\left[\frac{C_1}{\sqrt{C_1^2 - J'^2}}\left(\pi + 2\tan^{-1}\left(\frac{J'}{\sqrt{C_1^2 - J'^2}}\right)\right) + \frac{C_2}{\sqrt{C_2^2 - J'^2}}\left(-\pi + 2\tan^{-1}\left(\frac{J'}{\sqrt{1C_2^2 - J'^2}}\right)\right)\right] \tag{B.1}$$

where D, C_1 and C_2 are constants which take into account physical aspects of the IGBT, such as transconductance, breakdown voltage and semiconductor permittivity, and the current is represented by

$$J = \frac{I}{A}\sin(\theta - \phi) = J^{'}\sin(\theta - \phi). \tag{B.2}$$

The two summed terms correspond to the turn-on and turn-off losses respectively.

IGBT Static Loss

The IGBT conduction loss is given as

$$P_{cond,Q} = I^2 R_{CE,on} \left(\frac{1}{8} + \frac{1}{3\pi} M \cos\phi \right). \tag{B.3}$$

where I is the peak line current and M is the modulation index.

The IGBT blocking loss is not considered.

Diode Switching Loss

The diode switching loss is given as

$$P_{sw,D} = \frac{f_s V_R}{2S} \left(\frac{dI_F}{dt} \right) \left(\frac{S t_{rr}}{S+1} \right)^2 \tag{B.4}$$

where S is the snappiness factor; V_R is the reverse voltage applied to the diode; $\frac{dI_F}{dt}$ is the rate of fall of the forward current; and t_{rr} is the reverse recovery time.

Diode Static Loss

The diode conduction loss is given as

$$P_{cond,D} = I^2 R_D \left(\frac{1}{8} - \frac{1}{3\pi} M \cos\phi \right) + I \cdot V_D \left(\frac{1}{2\pi} - \frac{1}{8} M \cos\phi \right) \tag{B.5}$$

where R_D is the diode forward resistance and V_D is the forward voltage drop.

B.1.2 Empirical Model

The empirical model in [57] parameterises the current and voltage waveforms in order to calculate the loss.

IGBT Switching Loss

The IGBT turn-on loss is given as

$$\begin{aligned} W_{on} &= \int_{t_{1(on)}}^{t_{4(on)}} V_{ce}(t) I_c(t) dt \\ &\approx (V_{ce0} - V_{1(on)}) I_{1(on)} \left(1.25 t_r + \frac{1 - e^{-1.25\alpha_{1(on)} t_r}}{\alpha_{1(on)}} \right) \\ &+ V_{1(on)} I_{1(on)} \left(\frac{\left(1 - e^{-1.25\lambda_{1(on)} t_r}\right)}{\lambda_{1(on)}} + \frac{\left(1 - e^{-1.25(\alpha_{1(on)} + \lambda_{1(on)}) t_r}\right)}{\alpha_{1(on)} + \lambda_{1(on)}} \right) \\ &+ V_{cep} \left(I_{c0} k_D t_{rr} + \frac{I'_{1(rec)} \left(1 - e^{-\alpha_{2(rec)} (k_D t_{rr})^2}\right)}{2\alpha_{2(rec)}} \right) \\ &+ V_{2(on)} \left(I_{c0} \sqrt{\frac{\pi}{4\alpha_{2(rec)}}} + I_{rrm} \sqrt{\frac{\pi}{8\alpha_{2(rec)}}} \right) \\ &+ V_{3(on)} \left(I_{c0} \sqrt{\frac{\pi}{4\lambda_{3(on)}}} + I_{rrm} \sqrt{\frac{\pi}{4(\alpha_{2(rec)} + \lambda_{3(on)})}} \right). \end{aligned} \tag{B.6}$$

The parameters in this equation are shown in Figure 2.14.

Details of the turn-off loss calculation can be found in the paper [57].

IGBT Static Loss

The IGBT conduction loss in this model is simply the product of device current and forward saturation voltage.

The IGBT blocking loss is given as the product of blocking voltage and leakage current.

Diode Switching Loss

The diode turn-off loss is given by

$$W_{\text{rec}} \approx \frac{I_{d0}^2 V_{ds}}{2 I_{1(\text{rec})}} + \frac{I_{\text{rrm}}}{2} \sqrt{\frac{\pi}{2\alpha_{2(\text{rec})}}} \left(\sqrt{2} V_{d0} - V_{1(\text{rec})} \right) \tag{B.7}$$

where I_{d0} is the on-state current; V_{ds} is the on-state forward voltage drop; I_{rrm} is the peak reverse recovery current; V_{d0} is the post-switching blocking voltage; $V_{1(\text{rec})} = V_{d0} + V_{ds}$ and

$$\alpha_{2(\text{rec})} = \frac{\ln |10|}{t_{\text{rrb}}^2} \tag{B.8}$$

where t_{rrb} is the time taken for the reverse recovery current to decay from its peak value to 10% of that value.

The diode turn-on loss is ignored.

Diode Static Loss

The static losses in the diode are calculated as the product of current and forward voltage.

B.1.3 Analytical Model

The analytical model in [56] uses simple analytical expressions to represent empirical data provided by the device manufacturer.

IGBT Switching Loss

The IGBT switching loss is the sum of the turn-on and turn-off energies, given as

$$E_{\mathrm{onT}}(i_c) = B_{\mathrm{onT}} i_c + C_{\mathrm{onT}} i_c^2 \tag{B.9}$$

$$E_{\mathrm{offT}}(i_c) = B_{\mathrm{offT}} i_c + C_{\mathrm{offT}} i_c^2 \tag{B.10}$$

where B and C are the linear and quadratic coefficients of current which best fit the manufacturer's characteristics.

IGBT Static Loss

The IGBT conduction loss is calculated as

$$p_{\mathrm{fwT}} = v_{ce}(i_c) i_c. \tag{B.11}$$

where the forward characteristic is given by

$$v_{ce}(i_c) = A_{\mathrm{fwT}} + B_{\mathrm{fwT}} i_c. \tag{B.12}$$

in which A and B are coefficients of best fit.

Diode Switching Loss

The diode switching loss is given by

$$E_{\mathrm{recD}}(i_f) = B_{\mathrm{recD}} i_f + C_{\mathrm{recD}} i_f^2 \tag{B.13}$$

where B and C are coefficients of best fit.

The diode turn-on loss is ignored.

Diode Static Loss

The diode conduction loss model is calculated as

$$p_{\text{fwD}} = v_{ak}(i_f)i_f \tag{B.14}$$

$$v_{ak}(i_f) = A_{\text{fwD}} + B_{\text{fwD}}i_f \tag{B.15}$$

where A and B are coefficients of best fit.

B.2 Motor Losses

Like most materials, copper's resistivity varies with temperature. An expression for correcting the resistance according to temperature is given in [82]:

$$R_{T2} = R_{T1}(1 + \alpha_{T1}(T_2 - T_1)) \tag{B.16}$$

where α_{T1} is the temperature coefficient measured at T_1 (usually 25 °C). The stator temperature is often available but the rotor temperature is much more difficult either to measure or to predict precisely. However, a first order approximation for the increase in machine temperature can be obtained from the total loss, P_{tl}, as in (B.17), where θ is the steady state thermal resistance and τ is the thermal time constant which can both be determined experimentally [82].

$$\Delta T = \frac{P_{tl}}{\theta(1 + \tau s)}. \tag{B.17}$$

Klingshirn and Jordan [83] introduce an additional term to account for loss in the stator windings due to nonsinusoidal excitation:

$$W_1 = mr_1^s[(I_1^s)^2 + I_{har}^2]. \tag{B.18}$$

m is simply the number of phases, and r_1^s the stator resistance at fundamental frequency. I_1^s is the fundamental component of stator current and I_{har} represents the harmonic components of stator current, excluding the fundamental, as given by

$$I_{har} = \sqrt{(I_5^s)^2 + (I_7^s)^2 + \cdots + (I_k^s)^2}. \tag{B.19}$$

The harmonic components can in turn be calculated from the stator and rotor leakage inductances at each harmonic frequency, L_k^s and L_k^r respectively, using (B.20).

$$I_k = \frac{V_k}{2\pi f_1 k(L_k^s + L_k^r)}. \tag{B.20}$$

As the components included in (B.19) imply, the current contains no even harmonics or harmonics which are integer multiples of three (in a three-phase system).

An expression is also given for loss in the rotor due to harmonic currents. This loss must be calculated for each harmonic identified as significant and the contributions summed. The loss due to the kth harmonic is given as

$$W_{2k} = m(I_k^r)^2 r_k^r \tag{B.21}$$

where r_k^r is the resistance of the rotor at the kth harmonic. This variation of the rotor resistance with frequency is due to the skin effect, which is sometimes termed the deep-bar effect in cage IMs due to its significance in rotor bars deeper than 0.5 inches. The result is that harmonic currents see a different effective resistance from that measured at DC [84]. A similar approach is taken in [82] where an expression for rotor resistance at a harmonic frequency f_n is given in terms of its resistance at DC, R_{rdc}:

$$R_{rn} = R_{rdc}(1 + c_1 d f_n^{0.5}) \tag{B.22}$$

where d is the bar depth and c_1 is a constant determined by the bar shape and material.

B.2.1 Iron Loss

In the majority of models, the iron loss is separated into two components: that caused by magnetic hysteresis and that caused by eddy currents [41, 58, 64, 66, 81, 82, 85–92]. There is also a third, improperly understood component that is rarely considered separately in these models, which is referred to as anomalous loss [58,92].

A model for iron loss in an IM based on equivalent circuit parameters is given in [41]. It is dependent on air gap flux and frequency and contains hysteresis and eddy current coefficients. The terms depending on per unit slip, s, represent the rotor iron loss, and the other terms the stator iron loss:

$$P_{Fe} = \left[k_e(1+s^2)a^2 + k_h(1+s)a\right]\Phi_m^2. \tag{B.23}$$

The air gap flux can be approximated from

$$\Phi_m \approx \frac{r_r' I_r'}{sa} \tag{B.24}$$

in which r_r' and i_r' are the referred rotor resistance and current, respectively, and a is the per unit frequency ($a = \frac{\omega}{1-s}$).

A dynamic model for iron loss in an IM using loss separation is developed in [88]. An equivalent circuit approach is still used, but the core loss resistance is defined as a nonlinear function which takes into account the effects of hysteresis and eddy currents. The core loss resistance is given as a function of the voltage across it, $u = ||\mathbf{u_s} - R_s\mathbf{i_s}||$ and the stator flux, ψ_s:

$$R_{Fe}(u, \psi_s) = \frac{R_{Ft}}{1 + \frac{k\psi_s^{n-1}}{u}} \tag{B.25}$$

where the constants R_{Ft} and k and the constant n are found by fitting a curve to a FEA of the iron losses.

Anomalous Loss

Although most models do not consider anomalous loss separately it is sometimes calculated independently. In [91] the magnetic core is modelled as an equivalent circuit and includes a shunt resistance which represents the anomalous loss. An empirically-derived expression for the power loss is given by

$$P_a(t) = C_a \left| \frac{dB}{dt} \right|^{\frac{3}{2}} \rho A l_m \tag{B.26}$$

where C_a is a constant, B is the flux density, ρ is the density, A is the cross-sectional area and l_m is the mean length of the core. The non-linear resistance for the equivalent circuit is derived from this:

$$R_a = \frac{N^{\frac{3}{2}}}{C_a \rho l_m} \sqrt{\frac{|V_L(t)|}{A}}. \tag{B.27}$$

This is a very general model intended to be applied to any kind of device with a magnetic core, and so could be applied to any type of machine, as well as transformers, inductors, etc. Anomalous iron loss is considered with specific regard to IMs in [92]. Again, a relation to the $^3/_2{}^{\text{th}}$ power of the flux derivative is given:

$$P_{anom}^{density} = k_{anom} \frac{1}{T} \int^T \left| \frac{\partial B}{\partial t} \right|^{\frac{3}{2}} dt. \tag{B.28}$$

where k_{anom} is a constant.

B.2.1.1 Hysteresis Models

The dichotomy of hysteresis and eddy current loss is refuted in [93] on the basis that these supposedly separate effects originate from the same physical process of remagnetisation, and that models based on the loss separation approach are highly inaccurate, resulting in the necessary addition of anomalous loss. Furthermore, it is noted that hysteresis and anomalous loss cannot be calculated and must be determined experimentally. The impracticality of this approach is also recognised in [94].

Two other general approaches to iron loss calculation are identified in [93]: theoretical hysteresis models and empirical expressions.

Hysteresis models are broadly based on the work of either Jiles and Atherton [95] or Preisach [96]. These models are more directly related to the physical cause of the loss, but require many parameters to be determined empirically and so are of limited practicality.

B.2.1.2 Empirical Expressions

In 1892, Steinmetz presented an empirical equation for the energy loss in terms of the magnetic flux density [97]:

$$H = \eta B^{1.6}. \tag{B.29}$$

This is commonly expanded to include a dependence on frequency, f and to account for losses due to rapid flux changes [62, 93] as in (B.30). The two terms in this equation are commonly attributed to the hysteresis and eddy current losses respectively.

$$W_{core} = k_h B^x f + k_e B^2 f^2 \tag{B.30}$$

A major disadvantage of the Steinmetz equation is that it only applies to sinusoidally excited cores, and so does not apply accurately to inverter-fed machines. Attempts have been made to account for this by calculating the frequency components of the waveform using a Fourier expansion, solving the Steinmetz equation for each and superimposing the results [98, 99]. The validity of this is contested by [93] due to the nonlinearity of most magnetic materials.

A 'modified Steinmetz equation' (MSE) is developed in [93] to describe the effects of nonsinusoidal currents whilst maintaining the practicality of an empirical model:

$$p_v = (C_m f_{eq}^{\alpha-1} \hat{B}^{\beta}) f_r \tag{B.31}$$

where C_m, α and β are constants and the average flux density, $\hat{B} = \frac{\Delta B}{2}$. The equivalent frequency, f_{eq}, is calculated by averaging the remagnetisation rate over the remagnetisation period and normalising it with respect to a sinusoidal case:

$$f_{eq} = \frac{2}{\Delta B^2 \pi^2} \int_0^T \left(\frac{dB}{dt}\right)^2 dt. \tag{B.32}$$

The accuracy and validity of (B.31) is challenged by [94]. An expression is developed to address the issues with the MSE. This new expression is referred to as a 'generalised Steinmetz equation' (GSE) and is given in (B.33).

$$\bar{P_v} = \frac{1}{T} \int_0^T k_1 \left|\frac{dB}{dt}\right|^{\alpha} |B(t)|^{\beta-\alpha} dt = k f_e^{\alpha} \hat{B}_e^{\beta}. \tag{B.33}$$

The first expression in (B.33) can be applied to any nonsinusoidal waveform in order to accurately[1] predict the core loss. When applied to a pure sine wave, it agrees with the conventional Steinmetz equation.

It can be helpful to express the nonsinusoidal waveform in terms of an effective frequency, f_e, and effective amplitude, $\hat{B}_e$, and apply the second expression in (B.33),

[1] The GSE typically gives less than 5% error, although certain situations result in errors of up to 40% due to the lack of explicit modelling of minor hysteresis loops.

in a similar way to (B.32) and (B.31). f_e and $\hat{B}_e$ must be calculated such that (B.33) holds true.

There are a number of ways of doing this, but [94] presents two options which attempt to achieve an intuitive description of the waveform. Either the effective flux magnitude can be calculated using (B.34) or the effective frequency can be calculated using (B.35) and then the other parameter calculated by solving (B.33) as appropriate.

$$\hat{B}_e = \frac{1}{4}\int_0^T \left|\frac{dB}{dt}\right| dt \tag{B.34}$$

$$f_e = \frac{\int_0^T \left|\frac{d^2B}{dt^2}\right| dt}{2\pi \int_0^T \left|\frac{dB}{dt}\right| dt} \tag{B.35}$$

B.2.2 Stray Load Loss

[63] divides the stray losses into six components:

- Eddy current loss in the stator copper due to slot leakage flux
- Loss in the motor end structure due to end leakage flux
- High frequency surface losses in both the rotor and stator due to 'zigzag' leakage flux
- High frequency tooth pulsation and rotor I^2R losses due to 'zigzag' leakage flux
- The 'six-times frequency' rotor I^2R losses due to circulating currents induced by the stator belt leakage flux
- The extra iron loss in motors with skewed slots due to the skew leakage flux.

Equations are presented to calculate each component, but these are intended to be used as an aid to motor design by indicating the effects of changing the machine geometry. In any case, it is clearly impractical and unnecessary to consider the stray load loss in this much detail in a control context.

Harmonics

In the examination of the effect of inverter-induced harmonics in [83], each component of the six given above is treated separately. Since these are disregarded here as being too detailed, the expression from [82] for per phase stray load loss at a given harmonic frequency f_n given in (B.36) is more relevant. This model uses the idea that stray load loss is fundamentally due to hysteresis and eddy currents, and so can be modelled in a similar way to the iron loss, and using the same coefficients, k_h

and k_e. In the equation, k_{sln} is a stray loss constant and V_{sln} is the voltage across the stator leakage inductance.

$$P_{sln} = k_{sln}\left(\frac{k_h}{f_n} + k_e\right)V_{sln}^2 \tag{B.36}$$

B.2.3 Mechanical Losses

A simple equation for mechanical losses is given in [66]:

$$P_{fw} = a_w V_{air}\omega^q. \tag{B.37}$$

where V_{air} is the volume of air in the motor, a_w is a windage coefficient and q is a curve-fitting parameter. The last two parameters are determined from a deceleration test. The motor is run at high speed and then excitation removed, resulting in the loss being purely mechanical. The power can then be inferred from the rate of deceleration if the motor inertia is known. An example of the resulting loss characteristic is shown in Figure B.1.

A simplified version of this model is given in [82] where the friction and windage loss is assumed to be proportional to the third power of speed:

$$P_{fw} = k_{fw}\omega_r^3. \tag{B.38}$$

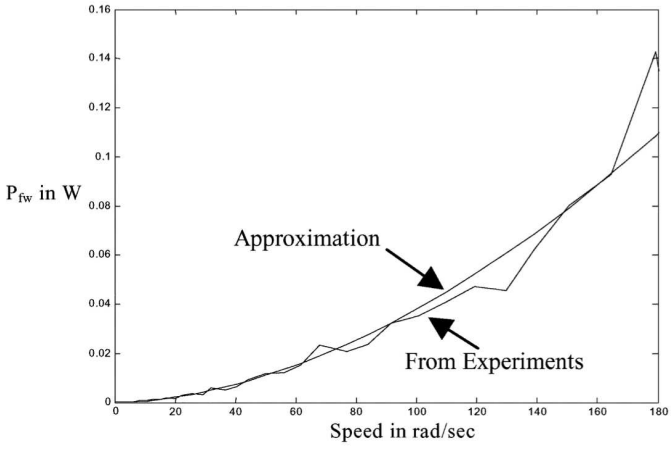

Figure B.1: Friction and windage losses against speed as determined from a deceleration test, with fitted curve

References

[1] "Climate change 2013—the physical science basis: Summary for policymakers," Intergovernmental Panel on Climate Change, Switzerland, Tech. Rep., Oct. 2013. [Cited on p. 1].

[2] "Key world energy statistics," International Energy Agency, Report, 2013. [Cited on p. 1].

[3] "National travel survey 2013," Department for Transport, UK, Tech. Rep., July 2014. [Cited on p. 1].

[4] "International energy outlook 2013," United States Energy Information Administration, Washington, DC, Tech. Rep., July 2013. [Cited on p. 2].

[5] T. Hutchinson, S. Burgess, and G. Herrmann, "Current hybrid-electric powertrain architectures: Applying empirical design data to life cycle assessment and whole-life cost analysis," *Applied Energy*, vol. 119, pp. 314–329, 2014. [Cited on pp. 3, 4].

[6] M. Zeraoulia, M. E. H. Benbouzid, and D. Diallo, "Electric motor drive selection issues for hev propulsion systems: A comparative study," *IEEE Trans. Veh. Technol.*, vol. 55, no. 6, pp. 1756–1764, Nov 2006. [Cited on p. 4].

[7] Y. Sato, S. Ishikawa, T. Okubo, M. Abe, and K. Tamai, "Development of high response motor and inverter system for the nissan leaf electric vehicle," in *SAE World Congress & Exhibition*, Detroit, MI, 2011. [Cited on p. 4].

[8] K. Iyer, X. Lu, K. Mukherjee, and N. Kar, "Online stator and rotor resistance estimation scheme using swarm intelligence for induction motor drive in EV/HEV," in *Int. Electric Drives Production Conf.*, Nuremberg, Sept. 2011, pp. 202–207. [Cited on p. 4].

[9] L. Kumar and S. Jain, "Electric propulsion system for electric vehicular technology: A review," *Renewable and Sustainable Energy Reviews*, vol. 29, pp. 924–940, 2014. [Cited on p. 4].

[10] S. Massari and M. Ruberti, "Rare earth elements as critical raw materials: Focus on international markets and future strategies," *Resources Policy*, vol. 38, no. 1, pp. 36 – 43, Mar. 2013. [Cited on p. 5].

[11] M. H. Rashid, Ed., *Power Electronics Handbook*, 3rd ed. Elsevier, 2011. [Cited on p. 10].

[12] R. Valentine, Ed., *Motor Control Electronics*. New York: McGraw-Hill, 1998. [Cited on pp. 10, 12, 15, 16, 48].

[13] A. Hava, R. Kerkman, and T. Lipo, "A high-performance generalized discontinuous PWM algorithm," *IEEE Trans. Ind. Applicat.*, vol. 34, no. 5, pp. 1059–1071, Sep 1998. [Cited on pp. 11, 12].

[14] J. Kolar, H. Ertl, and F. C. Zach, "Influence of the modulation method on the conduction and switching losses of a PWM converter system," *IEEE Trans. Ind. Applicat.*, vol. 27, pp. 1063–1075, Nov./Dec. 1991. [Cited on p. 11].

[15] M. Depenbrock, "Pulse width control of a 3-phase inverter with nonsinusoidal phase voltages," in *Conf. Rec. IEEE Int. Semiconductor Power Conversion Conf.*, 1977, pp. 399–403. [Cited on p. 11].

[16] Y. Wu, M. Shafi, A. Knight, and R. McMahon, "Comparison of the effects of continuous and discontinuous PWM schemes on power losses of voltage-sourced inverters for induction motor drives," *IEEE Trans. Power Electron.*, vol. 26, no. 1, pp. 182–191, Jan 2011. [Cited on p. 11].

[17] S. Huang, D. Pham, K. Huang, and S. Cheng, "Space vector PWM techniques for current and voltage source converters: A short review," in *Int. Conf. Electrical Machines and Systems*, Sapporo, Japan, Oct. 2012, pp. 1–6. [Cited on p. 12].

[18] A. Trzynadlowski, *Control of Induction Motors*. Academic Press, 2001. [Cited on p. 12].

[19] V. H. Prasad, "Analysis and comparision of space vector modulation schemes for three-leg and four-leg voltage source inverters," Master's Thesis, Electrical Engineering Department, Virginia Polytechnic Institute and State University, Blacksburg, VA, May 1997. [Cited on pp. 15, 54].

[20] J. Holtz, "Pulsewidth modulation—A survey," *IEEE Trans. Ind. Electron.*, vol. 39, pp. 410–420, Oct. 1992. [Cited on p. 15].

[21] C. Chan, "The state of the art of electric and hybrid vehicles," *Proc. IEEE*, vol. 90, pp. 247–275, Feb. 2002. [Cited on p. 16].

[22] A. Trzynadlowski, *The field orientation principle in control of induction motors*. Boston: Kluwer Academic, 1994. [Cited on p. 16].

[23] M. Koyama, M. Yano, I. Kamiyama, and S. Yano, "Microprocessor-based vector control system for induction motor drives with rotor time constant identification

function," *IEEE Trans. Ind. Applicat.*, vol. IA-22, pp. 453–459, May 1986. [Cited on p. 17].

[24] I. Takahashi and T. Noguchi, "A new quick-response and high-efficiency control strategy of an induction motor," *IEEE Trans. Ind. Applicat.*, vol. IA-22, pp. 820–827, Sept./Oct. 1986. [Cited on pp. 18, 19, 26].

[25] D. Casadei, F. Profumo, G. Serra, and A. Tani, "FOC and DTC: two viable schemes for induction motors torque control," *IEEE Trans. Power Electron.*, vol. 17, no. 5, pp. 779–787, Sep 2002. [Cited on p. 20].

[26] M. Aguirre, C. Calleja, A. Lopez-de Heredia, J. Poza, A. Aranburu, and T. Nieva, "FOC and DTC comparison in PMSM for railway traction application," in *Proc. 14th European Conf. Power Electron. and Applicat.*, Aug 2011, pp. 1–10. [Cited on p. 20].

[27] F. Korkmaz, I. Topaloglu, M. Cakir, and R. Gurbuz, "Comparative performance evaluation of FOC and DTC controlled PMSM drives," in *Fourth Int. Conf. Power Eng., Energy and Elect. Drives*, May 2013, pp. 705–708. [Cited on p. 20].

[28] N. Bianchi, S. Bolognani, and G. Grezzani, "Pm motors for very high dynamic applications," in *IEEE 36th Power Electron. Specialists Conf.*, June 2005, pp. 1332–1338. [Cited on p. 20].

[29] A. M. Bazzi and P. T. Krein, "Review of methods for real-time loss minimization in induction machines," *IEEE Trans. Ind. Applicat.*, vol. 46, pp. 2319–2328, Nov.-Dec. 2010. [Cited on pp. 20, 21, 23, 24].

[30] A. M. Knight, P. D. Milliband, C. Y. Leong, and R. A. McMahon, "Power losses in small inverter–fed induction motors," in *IEEE Int. Conf. Electric Machines and Drives*, San Antonio, TX, May 2005, pp. 601–607. [Cited on p. 21].

[31] Y. Wu, R. A. McMahon, Y. Zhan, and A. M. Knight, "Impact of PWM schemes on induction motor losses," in *IEEE Industrial Applications Conf.*, Tampa, FL, Oct. 2006, pp. 813–818. [Cited on p. 21].

[32] D. Neacsu, K. Rajashekara, and F. Gunawan, "Linear control of PWM inverter by avoiding the pulse dropping," in *Power Electronics in Transportation*, Oct 2002, pp. 31–38. [Cited on p. 21].

[33] S. Williamson and C. I. McClay, "Optimization of the geometry of closed rotor slots for cage induction motors," *IEEE Trans. Ind. Applicat.*, vol. 32, pp. 650–568, May/June 1996. [Cited on p. 21].

[34] M. Park and S. Sul, "Microprocessor-based optimal-efficiency drive of an induction motor," *IEEE Transactions on Industrial Electronics*, vol. IE-31, no. 1, pp. 69–73, Feb 1984. [Cited on pp. 21, 22].

[35] I. Kioskeridis and N. Margaris, "Loss minimization in scalar-controlled induction motor drives with search controllers," *IEEE Trans. Power Electron.*, vol. 11, pp. 213–220, Mar. 1996. [Cited on pp. 22, 23, 24].

[36] S. K. Sul and M. H. Park, "A novel technique for optimal efficiency control of a current-source inverter-fed induction motor," *IEEE Trans. Power Electron.*, vol. 3, pp. 192–199, Apr. 1988. [Cited on pp. 22, 23].

[37] R. Turner, R. Duke, and S. Walton, "Dynamic frequency scaling to improve converter efficiency," in *Australasian Universities Power Engineering Conference*, Perth, WA, Dec. 2007, pp. 1–5. [Cited on pp. 23, 55].

[38] M. N. Uddin and S. W. Nam, "New online loss-minimization-based control of an induction motor drive," *IEEE Trans. Power Electron.*, vol. 23, pp. 926–933, Mar. 2008. [Cited on p. 23].

[39] D. S. Kirschen, D. W. Novotny, and T. A. Lipo, "On-line efficiency optimization of a variable frequency induction motor drive," *IEEE Trans. Ind. Applicat.*, vol. 1A-21, pp. 610–616, May/June 1985. [Cited on p. 23].

[40] M. Waheedabeevi, A. Sukeshkumar, and N. Nair, "New online loss-minimization-based control of scalar and vector-controlled induction motor drives," in *IEEE Int. Conf. Power Electronics, Drives and Energy Systems*, Bengaluru, India, Dec. 2012, pp. 1–7. [Cited on p. 23].

[41] I. Kioskeridis and N. Margaris, "Loss minimization in induction motor adjustable-speed drives," *IEEE Trans. Ind. Electron.*, vol. 43, pp. 226–231, Feb. 1996. [Cited on pp. 24, 36, 149].

[42] J. R. Wells, P. L. Chapman, and P. T. Krein, "Applications of ripple correlation control of electric machinery," in *IEEE Int. Conf. Electric Machines and Drives*, vol. 3, Madison, WI, June 2003, pp. 1498–1503. [Cited on p. 24].

[43] P. Midya, P. T. Krein, R. J. Turnbull, R. Reppa, and J. Kimball, "Dynamic maximum power point tracker for photovoltaic applications," in *IEEE Power Electronics Specialists Conf.*, vol. 2, Baveno, Italy, June 1996, pp. 1710–1716. [Cited on p. 24].

[44] R. Yanamshetti, S. Bharatkar, D. Chatterjee, and A. Ganguli, "A hybrid fuzzy based loss minimization technique for fast efficiency optimization for variable speed induction machine," in *Int. Conf. Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology*, vol. 01, Pattaya, Chonburi, Thailand, May 2009, pp. 318–321. [Cited on p. 24].

[45] J.-F. Stumper, A. Dotlinger, and R. Kennel, "Loss minimization of induction machines in dynamic operation," *IEEE Trans. Energy Conversion*, vol. 28, pp. 726–735, Sept. 2013. [Cited on p. 24].

[46] A. Kusko and D. Galler, "Control means for minimization of losses in AC and DC motor drives," in *IEEE Trans. Ind. Applicat.*, vol. IA-19, July 1983, pp. 561–570. [Cited on p. 25].

[47] P. Cortes, M. Kazmierkowski, R. Kennel, D. Quevedo, and J. Rodriguez, "Predictive control in power electronics and drives," *IEEE Trans. Ind. Electron.*, vol. 55, pp. 4312–4324, Dec. 2008. [Cited on p. 25].

[48] Z. Ibrahim and E. Levi, "A comparative analysis of fuzzy logic and PI speed control in high-performance ac drives using experimental approach," *IEEE Trans. Ind. Applicat.*, vol. 38, pp. 1210–1218, Sept. 2002. [Cited on p. 25].

[49] J. Holtz and S. Stadtfeld, "A predictive controller for the stator current vector of AC machines fed from a switched voltage source," in *IEEE Int. Power Electronics Conf.*, vol. 2, Tokyo, Japan, Oct. 1983, pp. 1665–1675. [Cited on p. 25].

[50] S.-M. Yang and C.-H. Lee, "A deadbeat current controller for field oriented induction motor drives," *IEEE Trans. Power Electron.*, vol. 17, no. 5, pp. 772–778, Sep 2002. [Cited on p. 26].

[51] P. Mutschler, "A new speed-control method for induction motors," in *Proc. PCIM Europe: Int. Intelligent Motion Conf.*, Nuremberg, Germany, Jan. 1998, pp. 131–136. [Cited on p. 26].

[52] S. Kouro, P. Cortes, R. Vargas, U. Ammann, and J. Rodriguez, "Model predictive control—a simple and powerful method to control power converters," *IEEE Trans. Ind. Electron.*, vol. 56, no. 6, pp. 1826–1838, June 2009. [Cited on pp. 27, 41].

[53] S. Saeidi and R. Kennel, "Continuous set nonlinear model predictive control for an induction motor," in *European Conf. Power Electronics and Applications*, Lille, France, Sept. 2013, pp. 1–10. [Cited on pp. 27, 41].

[54] S. Mukherjee, R. G. Hoft, and J. A. McCormick, "Digital measurement of the efficiency of inverter-induction machines," *IEEE Trans. Ind. Applicat.*, vol. 26, pp. 872–879, Sept./Oct. 1990. [Cited on pp. 28, 34].

[55] S. S. Williamson, A. Emadi, and K. Rajashekara, "Comprehensive efficiency modeling of electric traction motor drives for hybrid electric vehicle propulsion applications," *IEEE Trans. Veh. Technol.*, vol. 56, pp. 1561–1562, May 2006. [Cited on pp. 29, 30].

[56] F. L. . Mapelli, D. Tarsitano, and M. Mauri, "Plug-in hybrid electric vehicle: modeling, prototype realization, and inverter losses reduction analysis," *IEEE Trans. Ind. Electron.*, vol. 57, pp. 598–607, Feb. 2010. [Cited on pp. 29, 30, 33, 144, 147].

[57] A. D. Rajapakse, A. M. Gole, and P. L. Wilson, "Electromagnetic transients simulation models for accurate representation of switching losses and thermal performance in power electronic systems," in *IEEE Trans. Power Delivery*, vol. 20, Jan. 2005, pp. 319–327. [Cited on pp. 29, 30, 31, 32, 144, 145, 146].

[58] M. J. Melfi, "Quantifying the energy efficiency of motors fed by adjustable frequency inverters," in *Petroleum and Chemical Industry Conf.* Anaheim, CA: IEEE Ind. Applicat. Soc., Sept. 2009. [Cited on pp. 34, 149].

[59] S. W. Nam and M. N. Uddin, "Model-based loss minimization control of an induction motor drive," in *IEEE Int. Symp. Industrial Electronics*, vol. 3, Montreal, QC, July 2006, pp. 2367–2372. [Cited on pp. 34, 36].

[60] P. Vas, *Parameter estimation, condition monitoring, and diagnosis of electrical machines*, ser. Monographs in electrical and electronic engineering. Oxford University Press, 1993. [Cited on pp. 35, 82].

[61] E. Dlala, A. Belahcen, J. Pippuri, and A. Arkkio, "Interdependence of hysteresis and eddy-current losses in laminated magnetic cores of electrical machines," *IEEE Trans. Magn.*, vol. 46, no. 2, pp. 306–309, Feb. 2010. [Cited on p. 35].

[62] R. E. Steven, *Electrical Machines and Power Electronics.* Wokingham, UK: Van Nostrand Reinhold, 1983, vol. 59. [Cited on pp. 35, 151].

[63] P. L. Alger, G. Angst, and E. J. Davies, "Stray-load losses in polyphase induction machines," *Trans. Amer. Inst. Elect. Engineers - Power Applicat. and Syst.*, vol. 78, no. 3, pp. 349–355, Apr. 1959. [Cited on pp. 36, 152].

[64] P. K. Sen and H. A. Landa, "Derating of induction motors due to waveform distortion," *IEEE Trans. Ind. Applicat.*, vol. 26, pp. 1102–1107, Nov./Dec. 1990. [Cited on pp. 36, 149].

[65] IEEE Power Engineering Society, "IEEE standard test procedure for polyphase induction motors and generators," *IEEE Std. 112-2004*, Nov. 2004. [Cited on pp. 36, 82].

[66] V. Raulin, A. Radun, and I. Husain, "Modeling of losses in switched reluctance machines," in *IEEE Trans. Ind. Applicat.*, vol. 40, Nov./Dec. 2004, pp. 1560–1569. [Cited on pp. 37, 149, 153].

[67] B. K. Bose, *Power Electronics And Motor Drives.* Burlington, VT: Academic Press, 2006. [Cited on p. 38].

[68] M. Masten and I. Panahi, "Digital signal processors for modern control systems," *Control Engineering Practice*, vol. 5, no. 4, pp. 449 – 458, 1997. [Cited on p. 38].

[69] P. Cortes, J. Rodriguez, R. Vargas, and U. Ammann, "Cost function-based predictive control for power converters," in *IEEE Industial Electroncs Conference*, Paris, France, Nov. 2006, pp. 2268–1173. [Cited on p. 39].

[70] A. K. Laha and K. Mahesh, "Sb-robustness of directional mean for circular distributions," *Journal of Statistical Planning and Inference*, vol. 141, no. 3, pp. 1269 – 1276, 2011. [Cited on p. 44].

[71] N. I. Fisher, *Statistical Analysis of Circular Data*. Cambridge University Press, 1993. [Cited on p. 44].

[72] J. Conway and R. Guy, *The Book of Numbers*, ser. Copernicus Series. Springer, 1996. [Cited on p. 51].

[73] IEEE Industry Applications Society, "IEEE recommended practices and requirements for harmonic control in electrical power systems," *IEEE Std. 519-1992*, pp. 1–112, April 1993. [Cited on p. 56].

[74] Y. Ito, S. Tomura, and S. Sasaki, "Development of vibration reduction motor control for hybrid vehicles," in *Annu. Conf. IEEE Industrial Electronics Society*, Taipei, Taiwan, Nov. 2007, pp. 516–521. [Cited on p. 71].

[75] F. W. Fuchs, J. Schroder, and B. Wittig, "State of the technology of power loss determination in power converters," in *European Conf. Power Electronics and Applications*, Lille, France, Sept. 2013, pp. 1–10. [Cited on p. 86].

[76] S. Meyer, *Data analysis for scientists and engineers*, ser. Wiley series in probability and mathematical statistics: Applied probability and statistics. Wiley, 1975. [Cited on p. 88].

[77] D. Gosden, B. Chalmers, and L. Musaba, "Drive system design for an electric vehicle based on alternative motor types," in *Int. Conf. Power Electronics and Variable-Speed Drives*, London, UK, Oct. 1994, pp. 710–715. [Cited on p. 92].

[78] E. Nicol Hildebrand and H. Roehrdanz, "Losses in three-phase induction machines fed by PWM converter," *IEEE Trans. Energy Conversion*, vol. 16, no. 3, pp. 228–233, Sep 2001. [Cited on p. 115].

[79] X. Li and S. Williamson, "Comparative investigation of series and parallel hybrid electric vehicle (HEV) efficiencies based on comprehensive parametric analysis," in *IEEE Vehicle Power and Propulsion Conf.*, Arlington, TX, Sept. 2007, pp. 499–505. [Cited on p. 131].

[80] H. Toliyat, E. Levi, and M. Raina, "A review of RFO induction motor parameter estimation techniques," *IEEE Trans. Energy Conversion*, vol. 18, pp. 271–283, June 2003. [Cited on p. 134].

[81] S. S. Williamson, A. Emadi, and K. Rajashekara, "Comprehensive efficiency modeling of electric traction motor drives for hybrid electric vehicle propulsion applications," *IEEE Trans. Veh. Technol.*, vol. 56, pp. 1561–1562, May 2006. [Cited on pp. 144, 149].

[82] G. Sousa, B. Bose, J. Cleland, R. Spiegel, and P. Chappell, "Loss modeling of converter induction machine system for variable speed drive," in *Proc. Int. Conf. Power Electronics and Motion Control*, vol. 1, San Diego, CA, Nov. 1992, pp. 114 –120. [Cited on pp. 148, 149, 152, 153].

[83] E. A. Klingshirn and H. E. Jordan, "Polyphase induction motor performance and losses on nonsinusoidal voltage sources," *IEEE Trans. Power Apparatus and Systems*, vol. 87, no. 3, pp. 624–631, Mar. 1968. [Cited on pp. 148, 152].

[84] ——, "Simulation of polyphase induction machines with deep rotor bars," *IEEE Trans. Power Apparatus and Systems*, vol. 89, no. 6, pp. 1038–1043, Jul./Aug. 1970. [Cited on p. 149].

[85] N. Mohan, T. E. Undeland, and W. P. Robbins, *Power Electronics: Converters, Applications and Design*, 3rd ed. Hoboken, NJ: Wiley, 2003. [Cited on p. 149].

[86] S. Morimoto, Y. Tong, Y. Takeda, and T. Hirasa, "Loss minimization control of permanent magnet synchronous motor drives," in *IEEE Trans. Ind. Electron.*, vol. 41, Oct. 1994, pp. 511–517. [Cited on p. 149].

[87] N. Margaris, T. Goutas, Z. Doulgeri, and A. Paschali, "Loss minimization in DC drives," *IEEE Trans. Ind. Electron.*, vol. 38, pp. 328–336, Oct. 1991. [Cited on p. 149].

[88] M. Ranta, M. Nikkanen, E. Dlala, A.-K. Repo, and J. Luomi, "Inclusion of hysteresis and eddy current losses in dynamic induction machine models," in *IEEE Int. Elect. Machines and Drives Conf.*, Miami, FL, May 2009, pp. 1387–1392. [Cited on p. 149].

[89] A. H. Wijenayake and P. B. Schmidt, "Modeling and analysis of permanent magnet synchronous motor by taking saturation and core loss into account," in *Proc. Int. Conf. Power Electronics and Drive Systems*, vol. 2, Singapore, May 1997, pp. 530–534. [Cited on p. 149].

[90] E. S. Sergaki, G. S. Stavrakakis, and A. D. Pouliezos, "Optimal robot speed trajectory by minimization of the actuator motor electromechanical losses," *J. Intell. Robotic Syst.*, vol. 33, no. 2, pp. 187–207, Feb. 2002. [Cited on p. 149].

[91] J. G. Zhu, S. Y. R. Hui, and V. S. Ramsden, "Discrete modelling of magnetic cores including hysteresis eddy current and anomalous losses," *IEE Proc. A - Science, Measurement and Technology*, vol. 140, no. 4, pp. 317–322, July 1993. [Cited on pp. 149, 150].

[92] T. C. Green, C. A. Hernandez-Aramburo, and A. C. Smith, "Losses in grid and inverter supplied induction machine drives," *IEEE Proc. - Elect. Power Applicat.*, vol. 150, no. 6, pp. 712–724, Nov. 2003. [Cited on pp. 149, 150].

[93] J. Reinert, A. Brockmeyer, and R. W. A. A. De Doncker, "Calculation of losses in ferro- and ferrimagnetic materials based on the modified steinmetz equation," *IEEE Trans. Ind. Applicat.*, vol. 37, pp. 1055–1061, July/Aug. 2001. [Cited on pp. 150, 151].

[94] J. Li, T. Abdallah, and C. R. Sullivan, "Improved calculation of core loss with nonsinusoidal waveforms," in *IEEE Industrial Applications Conf.*, vol. 4, Chicago, IL, Sept./Oct. 2001, pp. 2203–2210. [Cited on pp. 150, 151, 152].

[95] D. Jiles and D. Atherton, "Theory of ferromagnetic hysteresis," *Journal of Magnetism and Magnetic Materials*, vol. 61, no. 1-2, pp. 48 – 60, Sept. 1986. [Cited on p. 150].

[96] F. Preisach, "Über die magnetische nachwirkung," *Zeitschrift für Physik A Hadrons and Nuclei*, vol. 94, pp. 277–302, 1935. [Cited on p. 150].

[97] C. P. Steinmetz, "On the law of hysteresis," *Trans. American Institute of Electrical Engineers*, vol. 9, no. 1, pp. 1–64, Jan. 1892. [Cited on p. 151].

[98] P. Gradzki, M. Jovanovic, and F. Lee, "Computer-aided design for high-frequency power transformers," in *Proc. IEEE Applied Power Electronics Conf. and Expo.*, Los Angeles, CA, Mar. 1990, pp. 336 –343. [Cited on p. 151].

[99] R. Severns, "HF-core losses for nonsinusoidal waveforms," in *Int. High Frequency Power Conversion Conf.*, Toronto, Canada, June 1991, pp. 140–148. [Cited on p. 151].